EARNINGS VOLATILITY: CONCEPTS, MEASUREMENT, AND ACCOUNTING FOR ITS INCREASE IN THE UNITED STATES, 1971–2009

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Christopher Michael Handy

August 2013

EARNINGS VOLATILITY: CONCEPTS, MEASUREMENT, AND ACCOUNTING FOR ITS INCREASE IN THE UNITED STATES, 1971–2009

Christopher Michael Handy, Ph.D.
Cornell University 2013

Earnings volatility—variability over time in a worker's earnings—is interesting for its potential welfare consequences and as a labor market outcome. In the presence of liquidity constraints, households may not be able to smooth consumption when faced with volatile earnings. And rising earnings volatility may indicate that workers are changing jobs more often or that implicit contracts governing compensation within jobs are smoothing pay less, to name two possibilities. This dissertation shows that earnings volatility has increased in the U.S., and assesses, for a specific concept and measure of volatility, the reasons for the increase.

After an introduction, Chapter 2 introduces my preferred volatility measure, the volatility of residual earnings after estimating life cycle earnings profiles. Previous literature covers an array of concepts and measures of volatility. I outline differences among these approaches and argue that the primary motivation for studying volatility—potential welfare losses—has implications for the specification of the life cycle profiles and measures of earnings volatility.

Chapter 3 describes the data I use, a sample of male heads of household from the PSID. In Chapter 4, I show that earnings volatility has increased in the U.S., and that about 70 percent of the increase is explained by volatility of wages rather than hours worked. I describe how earnings volatility differs across groups, and show that it has increased among almost all groups. Finally,

I consider whether these findings depend on some measurement choices. All measures show increasing earnings volatility in the U.S., but the amount of the increase and comparisons of volatility across groups are often sensitive to methodological choices.

Chapter 5 addresses why earnings volatility has increased in the U.S. I develop a decomposition approach to attribute changes in economy-wide volatility to various factors. I create a demand shock index that measures workers' predicted exposure to labor demand shocks, using national changes in the occupation-industry distribution of hours worked. My major finding is that larger or more frequent labor demand shocks explain about half of the increase in economy-wide earnings volatility between 1975 and 2005.

BIOGRAPHICAL SKETCH

Christopher Handy was born on March 14, 1985, in Panama City, Florida. He completed high school at Mississippi Country Christian Academy in Dell, Arkansas in 2003, and was a National Merit Scholar. He obtained a B.A., *summa cum laude*, from Vanderbilt University in Nashville, Tennessee in May 2007 with majors in Economics and Mathematics. While in the Ph.D. program in Economics at Cornell University, he worked for one year as a research assistant and for three years as a teaching assistant, and received an M.A. in Economics in January 2012.

This dissertation is dedicated to my wife, Jen, who has been an unwavering source of love throughout my graduate studies. In addition to her professional accomplishments during our time in Ithaca, her support allowed me to become a better teacher, researcher, friend, and husband.

ACKNOWLEDGEMENTS

I am most grateful to my committee chair, Gary Fields, both for his patience and generosity as my advisor and for his numerous comments and suggestions, which greatly improved the focus and clarity of this dissertation. George Jakubson was especially giving of his time, and provided very helpful feedback early in the research process and throughout the job market process. Larry Kahn gave consistently excellent comments on everything I wrote, and his thoughts improved this dissertation throughout.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

This dissertation is about the meaning and measurement of earnings volatility, which is variability over time in a worker's earnings or some component of those earnings, and about the recent increase in earnings volatility among men in the U.S. In this introductory chapter, I do three things. First, I describe the motivation for studying earnings volatility. Second, I discuss the questions that I answer in this dissertation. Finally, I present a summary of my findings.

1.1 Motivation

The primary motivation for studying earnings volatility is that the welfare of individuals and households is reduced if they are not able to smooth consumption in the face of volatile income. Researchers frequently reference this motivation. For example, Dynarski and Gruber (1997) and Gorbachev (2011) study the link between earnings volatility and household consumption volatility. Studies of earnings volatility that mention as motivation risk, uncertainty, or welfare losses through consumption volatility include Shin and Solon (2011), Ziliak, Hardy, and Bollinger (2011), Dynan, Elmendorf, and Sichel (2012), and Moffitt and Gottschalk (2012).

The increase since the 1970s in earnings volatility among men in the U.S. was first documented by Gottschalk and Moffitt (1994) and has continued since then, as I show in chapter 4. This increase has sparked concern that the welfare losses associated with volatility are also growing. This does not have to be the case: government transfers or spouse's labor earnings, for example, could keep the volatility of household income from rising, or improvements in financial products or access to the financial system could prevent consumption volatility from rising even if households' income volatility increased.

However, there are reasons to believe that higher earnings volatility would, in fact, lead to more consumption volatility. Gruber (1997) reports that household food consumption drops by 7 percent on average upon job loss by male household heads, and by much more for workers eligible for less generous unemployment insurance benefits. For low-asset households, credit cards are one option for replacing lost income, but credit limits and high interest rates limit the extent to which this method can be used to smooth consumption. Sullivan (2008) finds that low-asset households—those in the second and third deciles of the asset distribution—use credit cards to replace 11.5 to 13.4 cents of each dollar of earnings lost due to unemployment. Very low-asset households—those in the bottom decile of the asset distribution—do not increase credit card borrowing in response to unemployment shocks, likely because they have very limited access to credit.

There is also direct evidence that household consumption volatility has increased in the U.S. over the same time period as the increase in earnings volatility documented in previous literature and in this dissertation. Dynarski and Gruber (1997) found that household consumption volatility rose as earnings volatility of household heads increased between 1970 and 1991. Gorbachev (2011) also documents rising consumption volatility, and finds a link between volatility of earnings and volatility of food consumption that is particularly strong among low-income households. She calculates that by 2004, "an average household would be willing to sacrifice 4.15 percent of its annual nondurable consumption to reduce consumption risk back to where it was in 1971."

A secondary motivation for studying earnings volatility is to learn about labor markets. Earnings volatility is inherently interesting because it reflects, for example, pay setting within jobs and the consequences of job change. Also, the persistent increase in earnings volatility suggests ongoing or permanent changes in the U.S. labor market. For these reasons, Moffitt and Gottschalk (2002) called rising volatility "one of the more puzzling aspects" of recent labor market trends. The empirical part of this dissertation addresses this puzzle in part by studying labor demand, job change, and earnings variability within jobs.

1.2 Questions

What is earnings volatility?

Before studying a phenomenon empirically, it is good to be clear about concepts and measures of that phenomenon. This is especially important for research on earnings volatility. Previous studies of volatility begin with a measure, sometimes discuss advantages and disadvantages of the chosen measure, then move to the empirics. But without considering underlying concepts, or at least comparing multiple measures, it is impossible to know whether one measure is more appropriate than another. Nevertheless, researchers often compare results across different measures of volatility, implicitly treating earnings volatility as a single concept. Therefore, my first task is to identify the conceptual differences among measures of volatility that have been used in previous research.

How should earnings volatility be measured?

All or almost all of the measures I discuss seem to be reasonable ways to quantify the variability of earnings, but there are many differences among them, and authors have arrived at different conclusions depending on which measure is used. Which measures should be preferred? To make progress on this issue, I identify some methodological recommendations that are consequences of the fact that risk is the primary motivation for studying volatility.

How has earnings volatility changed over time, and how does it differ across workers?

My primary empirical question is why earnings volatility has increased over time. Before addressing that issue, I first establish that volatility has, in fact, increased among men in the U.S. Another empirical issue, secondary in interest but instrumental in answering the question of why volatility increased, is how volatility differs across workers.

How important are some methodological choices for these empirical results?

I identify conceptual differences among measures of volatility and make some methodological recommendations for studying volatility. How important are these methodological choices for the different results obtained by previous researchers? I show how important these issues are within a single sample from a single data source. Of course, different data sources and sample restrictions could also explain the varying results found in previous research.

Why did earnings volatility increase?

After identifying a conceptual basis for my preferred measure of volatility, and after establishing that volatility has increased over time in my sample, I turn to the question of why earnings volatility increased. To answer this, I use a decomposition approach that makes use of differences in earnings volatility across workers, along with changes in worker characteristics over time.

1.3 Summary of findings

I begin in chapter 2 by introducing my preferred method of measuring earnings volatility. My approach is to first estimate life cycle earnings profiles, then measure the variability over time of workers' residuals from these profiles. The life cycle profiles I specify are somewhat different from the previous literature on earnings volatility, but the measure I apply to the residuals is what Moffitt and Gottschalk (2012) call the window averaging approach: a worker's volatility in a given year is the variance of the worker's residual log earnings during a nine-year window centered on that year. Economy-wide earnings volatility is the mean of workers' earnings volatilities for a given year.

The rest of chapter 2 is concerned with the first two questions listed above. First, I identify multiple concepts of earnings volatility. The most important distinction is what component of earnings is used when measuring the variability of earnings: earnings or log earnings (*volatility of total earnings*), residual log earnings from life cycle profiles (*volatility of residual earnings*), or the transitory as opposed to permanent component of residual earnings (*volatility of transitory earnings*).

Second, I argue that the risk motivation for studying earnings volatility implies that volatility of residual earnings, rather than volatility of total earnings or transitory earnings, is the most salient volatility concept, and that this motivation has implications for the measurement of the volatility of residual earnings. Life cycle earnings profiles should account for earnings changes that workers likely anticipate, but should leave unanticipated changes to contribute to measured earnings volatility. Therefore, the life cycle profiles should not absorb macroeconomic shocks, which are unlikely to be anticipated by workers, but should absorb worker-specific earnings trends, which recent research suggests are mostly anticipated. Also, the welfare loss of interest implies volatility measures should be based on squared changes or variances rather than on absolute changes or standard deviations.

After describing in chapter 3 the data that I use, I show in chapter 4 that under my preferred measure of earnings volatility, economy-wide volatility increased by over 70 percent among a sample of male heads of household in the PSID. This finding is robust to changing the length of the window over which workers' earnings volatility is computed, to including workers who do not have earnings observations for each time period in the window, and to dropping earnings volatility observations based on any imputed values of earnings. Most of the increase in economy-wide earnings volatility can be attributed to rising volatility of hourly wages, rather than rising volatility of hours worked.

I also describe how earnings volatility differs across workers. There are significant differences in earnings volatility across age groups, marital status, and many other categories, but no significant differences across education groups or by union membership. These descriptive exercises are useful in formulating theories about why economy-wide earnings volatility increased. For example, highly educated workers do not have significantly more earnings volatility than workers with little education, so a shift in the composition of the labor force toward more educated workers will not be empirically important, in an accounting sense, in explaining rising economy-wide earnings volatility. I show that mean earnings volatility has increased over time within almost all groups, no matter how the sample is partitioned. This suggests that compositional changes cannot explain the increase in economy-wide earnings volatility over time.

I show that volatility of residual earnings is largest among the lowest earnings quintile, and that from 1975 to 2005, earnings volatility rose more in the lowest earnings quintile than in any other quintile. In a simple exercise to demonstrate how the welfare losses from earnings volatility differ across earnings quintiles, I show that a worker in the highest earnings quintile would pay about 1 percent of earnings to eliminate volatility of residual earnings in 1975 and a little over 2 percent in 2005, but a worker in the lowest earnings quintile would pay about 3 percent in 1975 and almost 6 percent in 2005.

The increase in economy-wide volatility of residual earnings could be due to factors associated with job change—more frequent job change, more earnings volatility for a given rate of job change, or both—or factors associated with pay setting within jobs. I find that the rate of job change in my sample actually decreased slightly over time. The composition of job changes shifted toward more voluntary changes rather than involuntary changes, and voluntary changes are associated with less earnings volatility than involuntary changes. Therefore, job change is unlikely to have strong explanatory power as a proximate cause of economy-wide volatility of residual earnings. To look for clues as to how changes in pay setting within jobs might have affected earnings volatility, I look at the level and change over time of volatility of residual earnings in single-digit occupation and industry categories, both for all workers and for job stayers only. There is no clear pattern to the sector-specific increases in earnings volatility.

The final section of chapter 4 examines the empirical importance of the methodological recommendations made in chapter 2. I find that measuring volatility of residual earnings rather than volatility of total earnings does not change the time series of economy-wide earnings volatility. Allowing aggregate shocks to contribute to measured earnings volatility also does not change economy-wide earnings volatility. Including worker-specific earnings trends in life cycle profiles has a modest effect on the degree to which economy-wide earnings volatility increases over time, and the functional form of the earnings volatility measure has a large influence on the size of the increase over time in economy-wide earnings volatility. All of the methodological recommendations have noticeable effects on cross-worker comparisons of earnings volatility, such as comparing the volatility of job changers to the volatility of job stayers.

Chapter 5 addresses my final question: why has economy-wide volatility of residual earnings, by my preferred measure, increased in the U.S.? It provides the most extensive empirical analysis of this issue to date, using worker-level data; previous papers addressing this topic use calibration exercises or unrepresentative firm-level data. I develop a decomposition approach to attribute the change over time in economy-wide volatility to various factors.

I find that larger or more frequent shocks to labor demand are the strongest explanation for the increase in economy-wide earnings volatility. A demand shock index, which predicts workers' exposure to labor demand shocks on the basis of changes in the national employment distribution across occupation-industry sectors, has risen over time and accounts for about half of the increase in economy-wide earnings volatility. This finding is robust to the inclusion of a rich set of controls and to various ways of constructing the demand shock index. Worker and job characteristics, such as education, experience, occupation, and union membership, explain close to none of the increase in economy-wide earnings volatility.

I also test existing explanations for rising economy-wide earnings volatility and find that these theories explain very little of the increase in volatility. In particular, because of previous research on the relationship between rising earnings volatility and job and occupation mobility, I assess the role of job change and occupation change in the increase in economy-wide earnings volatility. Job change is associated with spikes in a worker's earnings volatility, but the incidence of job change did not increase over time in my sample, so patterns of job change predict no increase over time in economy-wide earnings volatility. This is confirmed when using job, occupation, or industry tenure as explanatory variables instead of job change. Even though workers with low tenure have higher earnings volatility—possibly because they are accumulating specific human capital at a faster rate—the incidence of low tenure has not increased. Finally, I find no support for the possibility that demand shocks affect earnings volatility only through job change, which might cause the earnings volatility consequences of job change to increase even if the rate of job change does not increase.

CHAPTER 2

MEASURING EARNINGS VOLATILITY

This chapter begins by introducing my preferred method of measuring earnings volatility in section 2.1. I then discuss the literature on earnings volatility in two separate sections. In section 2.2, I focus on the conceptual differences among some previous approaches. This includes both the subset of earnings that are of interest—total earnings, residual earnings from a life cycle earnings profile, or the transitory, as opposed to permanent, component of residual earnings—and the measure used to quantify earnings variability. In section 2.3, I argue that measurement choices should be informed by the fact that risk is the primary motivation for studying earnings volatility. This has implications for the type of earnings volatility that is most relevant; for the estimation of life cycle earnings profiles, which is a preliminary step in measuring the volatility of residual earnings or transitory earnings; and for the functional form of the earnings volatility measure.

2.1 My preferred measure

The primary results in this dissertation are based on a measure of the volatility of residual earnings, which is the variability of residuals from a life cycle earnings profile. This method avoids counting as volatility earnings changes that are explained by life cycle dynamics, on the assumption that these changes are generally anticipated by workers and do not reflect risk. Rapid early-career earnings growth does not, by this view, contribute to earnings volatility. The smooth portion of this earnings growth is expected; if an earnings-experience profile were completely smooth, by this concept it would have zero earnings volatility.

The life cycle earnings profiles that I estimate include worker fixed effects, worker-specific trends, and a quartic in potential experience. The coefficients on the potential experience terms are allowed to differ by education group (high school or less, some college, and college or more), and coefficients estimated at the education-group level are allowed to change smoothly over time in a cubic path. Denote log real earnings by y and potential experience by x. For worker i in education group s in year t, the profile is

$$y_{ist} = \gamma_i + \delta_i x_{ist} + \sum_{k=1}^{3} \beta_{1ks} t^k x_{ist} + \sum_{j=2}^{4} \sum_{k=0}^{3} \beta_{jks} t^k x_{ist}^j + \varepsilon_{ist}. \tag{2.1}$$

This is a very flexible life cycle earnings profile that is a continuous function of potential experience for each worker.

Denote the residuals from this profile by e_{it}. My measure of a worker's earnings volatility at time t is the worker's variance of these residuals during the nine-year interval centered at t, $[t-4, t+4]$. Conditional on estimated residuals, this is the measure termed *window averaging* by Moffitt and Gottschalk (2012). Because window averaging can be thought of as a class of measures that potentially applies to any interval containing two or more observations, I denote this nine-year version by WA9. I have biannual, not annual, earnings observations—see chapter 3 for an explanation—so there are five observations during the nine-year window. Therefore, my expression for earnings volatility is

$$V_{it}^{WA9} = \frac{1}{4} \sum_{\tau \in [t-4, t+4]} (e_{i\tau} - \bar{e}_{it})^2, \tag{2.2}$$

where $\bar{e}_{it}$ is the worker's mean residual over the window $[t-4, t+4]$.

Throughout this dissertation, *economy-wide* earnings volatility by my preferred measure refers to the cross-worker mean of earnings volatility for a given year:

$$V_t^{WA9} = \frac{1}{N} \sum_{i=1}^{N} V_{it}^{WA9}. \tag{2.3}$$

2.2 Concepts and measures of earnings volatility

A growing literature studies the magnitude of earnings changes and earnings variability among men, women, and households, and especially whether economy-wide measures of this magnitude have increased over time. However, the methods in these studies are generally not derived from a clearly stated concept of earnings volatility. The result is a proliferation of measures that sometimes reach different conclusions about the trend in economy-wide earnings volatility, while implicitly treating earnings volatility as a single concept.

Even the name of the phenomenon being studied is not settled. Most of the literature refers to *earnings volatility*. The term *earnings instability* is less frequently used, although it was coined by the most prominent researchers on this topic (Gottschalk and Moffitt, 1994), and occasionally *earnings variability* has been used. Because these are ordinarily synonymous terms, and because there is no way to partition the literature into earnings volatility studies and earnings instability studies, for clarity I adopt the term *earnings volatility* consistently throughout this dissertation, even when discussing literature that has used other terms.

Conceptually, earnings volatility is variability over time in a worker's earnings or in some component of those earnings. This section reviews several concepts and measures of earnings volatility, following the diagram in Figure 2.1. I first separate the types of volatility by the earnings component for which variability is measured: total earnings, earnings residuals from a life cycle earnings profile, or transitory earnings, which are a subset of residual earnings. I will begin with volatility of residual earnings because it is the most frequently studied type of earnings volatility.

For measures of the volatility of residual earnings, the specification of the life cycle profile is an important part of the measurement process, but is logically separate from the functions that measure earnings volatility once earnings residuals have been estimated. That is, each combination of a specification for life cycle earnings profiles and an earnings volatility measure that operates on residuals from the profiles is a valid method for quantifying the volatility of residual earnings. This section discusses the measurement of earnings volatility conditional on estimated earnings residuals (if applicable), and explains why measures may differ even for a common sample of earnings or earnings residuals. The next section discusses in detail the specification of the life cycle profiles.

Some of the earnings volatility measures apply to individual workers, in which cases a natural economy-wide measure is the mean of the worker-specific measures for a given time period. Others apply to groups of workers and are not additively decomposable into worker-specific measures. For measures of the volatility of transitory earnings, this is a consequence of the fact that the permanent and transitory components of residual earnings cannot be separately identified for individuals. For measures of the volatility of total earnings or residual earnings, authors using non-decomposable measures are not interested in obtaining worker-specific earnings volatility, perhaps because they are only interested in economy-wide earnings volatility or because they want to compare earnings volatility across groups but do not need to do further conditional analysis, such as regressing earnings volatility on worker-specific

Figure 2.1: Concepts and measures of earnings volatility

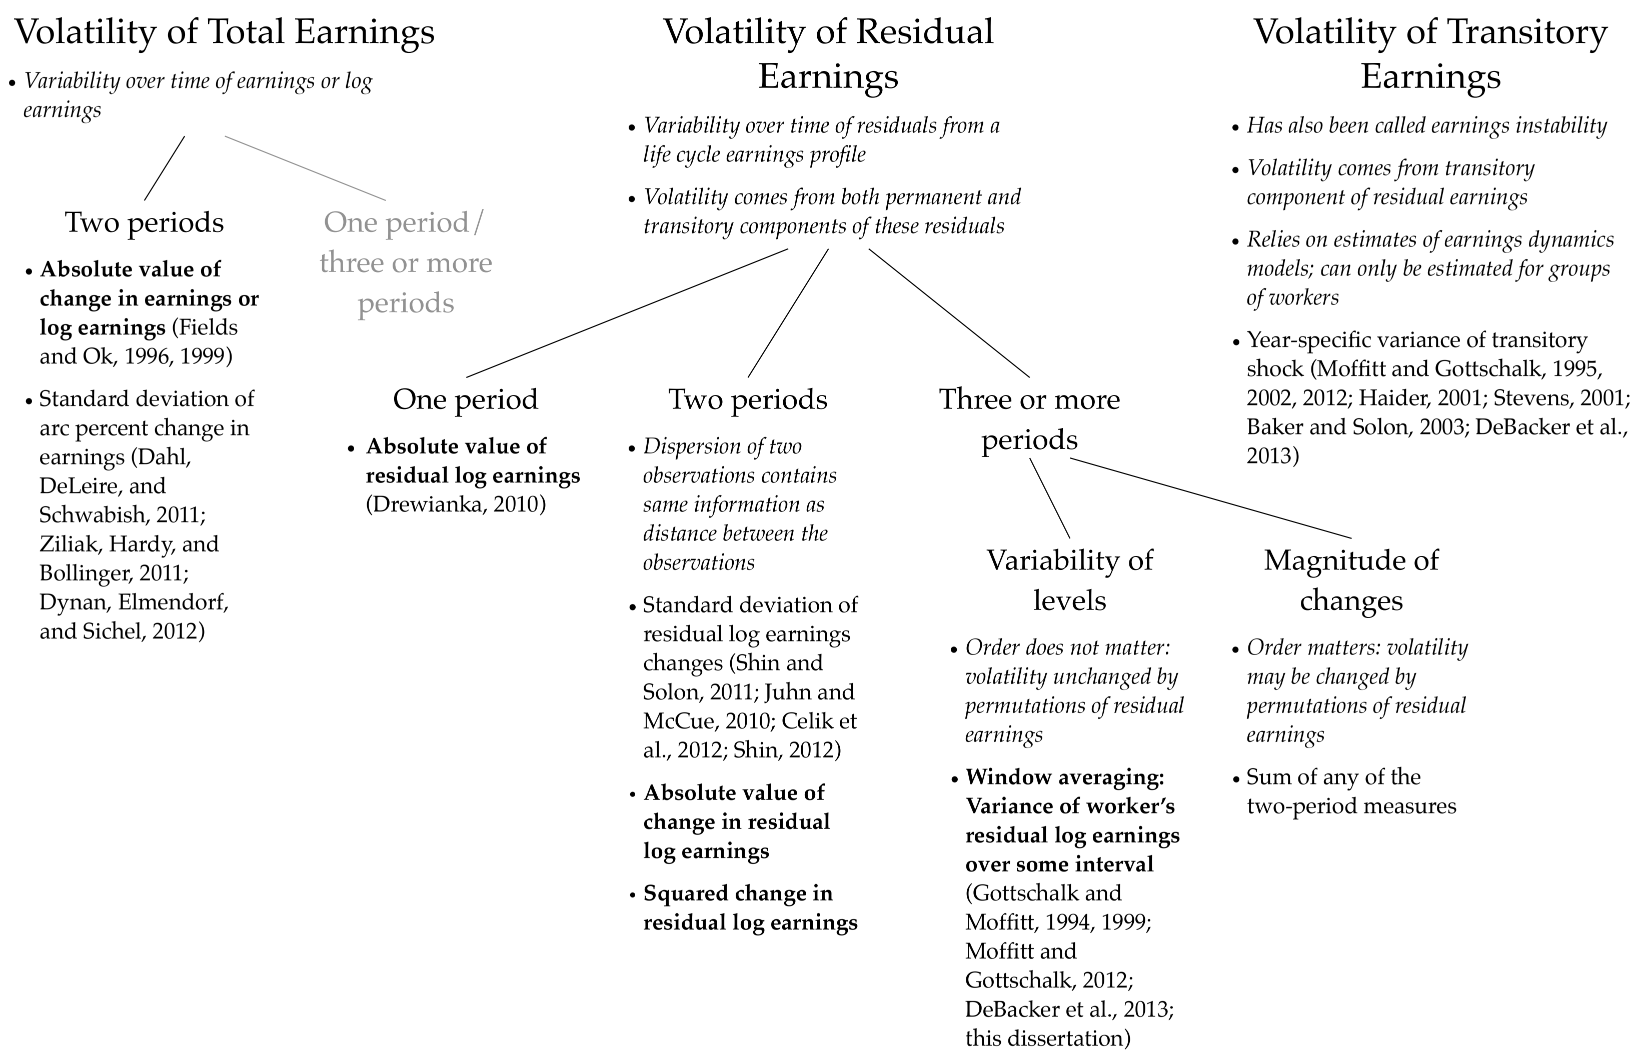

Notes: Bold text denotes earnings volatility measures that apply to individual workers. In each of these cases, the natural economy-wide measure is the mean of the worker-specific measures for a given time period. Non-bold text denotes earnings volatility measures that apply to groups of workers and are not additively decomposable into worker-specific measures.

covariates. In these cases, as I detail below, the measures can always be viewed as a function of worker-level measures and an average earnings change.

For notation, let real earnings for worker i in year t be Y_{it}, and denote log real earnings by y_{it}. A life cycle earnings profile may be estimated using the log earnings observations, and residuals from the profile are denoted e_{it}. In studies of the volatility of transitory earnings, it is typical to assume that these residuals are additively separable into a permanent component, μ_{it}, which may be multiplied by a time-varying factor loading, and a transitory component, ν_{it}. First differences are denoted by Δ.

2.2.1 Volatility of residual earnings

Conceptually, volatility of residual earnings is variability over time in a worker's earnings that is not attributable to life cycle dynamics. Measures of the volatility of residual earnings typically begin by estimating life cycle earnings profiles, then quantifying the variability over time of residuals from these profiles. However, this process can be reversed: One of the measures discussed below first estimates a life cycle profile of log earnings changes, then computes volatility using the residuals from this regression.

Volatility of residual earnings reflects changes in both the permanent and transitory components of residual earnings. I discuss measures of the volatility of residual earnings categorized by whether they apply to earnings residuals from a single time period, two time periods, or three or more time periods.

Volatility in one time period

Drewianka (2010) measures the volatility of a worker's residual earnings as the distance between log earnings and the estimated life cycle earnings profile, which is the absolute value of residual earnings:

$$V_{it}^{D} = |e_{it}|\,. \tag{2.4}$$

The natural corresponding economy-wide measure is the mean of workers' volatilities in a given year:

$$V_{t}^{D} = \frac{1}{N}\sum_{i=1}^{N} V_{it}^{D}. \tag{2.5}$$

While it may seem odd to measure volatility using a single observation, the measure can be interpreted as the unexpected component of earnings in a given year. This interpretation relies on the validity of the estimated life cycle earnings profile as a measure of the expected component of earnings, so Drewianka (2010) estimates very flexible profiles with worker-specific intercepts as well as worker-specific coefficients on experience and experience squared. If a worker's earnings always lie on the worker's estimate life cycle earnings profile, the worker will have zero earnings volatility by this measure in each year.

Still, this method can produce undesirable results. Consider a worker who has an earnings residual of -1 for ten years, followed by ten years with an earnings residual of 1. Most would say this worker has little to no earnings volatility at the beginning and end of the panel, and positive earnings volatility near the time of the earnings change. But by V_{it}^{D}, earnings volatility is positive in each year and does not change over time. Volatility measures based on earnings changes, discussed immediately below, give more intuitive results.

Volatility over two time periods

The absolute residual change (AC) and squared residual change (SC) measures One natural measure of the volatility of a worker's residual earnings is the magnitude of the worker's change in residual log earnings:

$$V_{it}^{AC} = |\Delta e_{it}| = |e_{it} - e_{i,t-1}| \,. \tag{2.6}$$

Below, I argue that another natural measure, the square of the change in residual log earnings,

$$V_{it}^{SC} = (\Delta e_{it})^2 = (e_{it} - e_{i,t-1})^2 \,, \tag{2.7}$$

is a more attractive measure when the motivation for studying earnings volatility is risk, in the sense that shocks in the labor market may pass through to consumption volatility. For both measures, the natural corresponding measure of economy-wide volatility of residual earnings is the mean of workers' volatilities:

$$V_t^{AC} = \frac{1}{N}\sum_{i=1}^{N} V_{it}^{AC}; \quad V_t^{SC} = \frac{1}{N}\sum_{i=1}^{N} V_{it}^{SC}. \tag{2.8}$$

Given the simplicity and directness of these measures, it is surprising that they have not previously been used by researchers studying volatility of residual earnings.

The Shin and Solon (SS) measure Shin and Solon (2011), Juhn and McCue (2010), Celik et al. (2012), and Shin (2012) measure economy-wide volatility of residual earnings as the standard deviation of residual log earnings changes from one period to the next. They first compute changes in log earnings, then estimate a life cycle profile of these changes, and measure volatility using residuals from this profile. Denote the residuals by $\widetilde{\Delta e}_{it}$ and the mean of these residuals in year t by $\overline{\widetilde{\Delta e}}_t$. The SS measure of economy-wide volatility of residual earnings is

$$V_t^{SS} = sd(\widetilde{\Delta e}_{it}) = \left[\frac{1}{N-1}\sum_{i=1}^{N}\left(\widetilde{\Delta e}_{it} - \overline{\widetilde{\Delta e}}_t\right)^2\right]^{\frac{1}{2}}. \tag{2.9}$$

As Shin and Solon (2011) point out, this is an easily understood and easily calculated measure, which is likely why others have adopted it in work that describes the evolution of economy-wide earnings volatility over time in different data sets.

This measure can be expressed as a function of worker-specific residual earnings changes and the average or aggregate residual earnings change:

$$V_t^{SS} = \left[\frac{1}{N-1} \sum_{i=1}^{N} \left(\widetilde{\Delta e}_{it} \right)^2 - \frac{N}{N-1} \left(\overline{\widetilde{\Delta e}}_t \right)^2 \right]^{\frac{1}{2}}. \tag{2.10}$$

Suppose that the SS approach to taking residuals—earnings changes first, then life cycle profiles—roughly matches the more common approach of estimating life cycle profiles on log earnings and using changes in the residuals from those profiles. That is, suppose $\Delta e_{it} \approx \widetilde{\Delta e}_{it}$. Then it is easy to see from (2.7) and (2.10) that the square of the SS measure is related to the mean of the SC measure:

$$(V_t^{SS})^2 \approx V_t^{SC} - \left(\overline{\Delta e}_t \right)^2. \tag{2.11}$$

In my view, the SS measure understates earnings volatility by missing the aggregate earnings change, which is unlikely to be anticipated by workers. In practice, researchers using the SS measure estimate the life cycle profile of earnings changes separately by year, implicitly setting $\overline{\widetilde{\Delta e}}_t = 0$. I further discuss this issue conceptually in section 2.3 and empirically in section 4.3.

Volatility over three or more time periods

For measures of the volatility of residual earnings that apply to three or more time periods, it is useful to distinguish between two subcategories. Measures of the variability of earnings residuals, which are defined on levels of residuals, are unaffected by any permutation of the earnings residuals for a given worker; the order of the residuals does not matter. Measures of the magnitude of changes in earnings residuals are potentially affected by permutations of the worker's earnings residuals; order matters. I will briefly discuss each type, then characterize the relationship between them for a special case.

Note that for two-period measures of the volatility of residual earnings, order does not matter, because the variability of two observations contains the same information as the distance between the two observations. Another way of seeing this is to note that each measure discussed in this chapter is unchanged if the order of the residuals is exactly reversed. For two-period measures, reversal is the only possible permutation of the observations.

Variability of earnings residuals—the window averaging (WA) measure In the window averaging (WA) measure, the volatility of a worker's residual earnings in year t is the variance of the worker's earnings residuals in some interval that contains t. This measure was introduced by Gottschalk and Moffitt (1994) and later used by those authors in Gottschalk and Moffitt (2009) and Moffitt and Gottschalk (2012), as well as by DeBacker et al. (2013). Typically, the earnings volatility in the window is assigned to the window midpoint, so that volatility at time t applies to a window centered on t. Let p be the length of the window and $w = \frac{p-1}{2}$ be the number of distinct time periods in the window both before and after t. I will treat window averaging as a class of measures, with the specific measure determined by the window length and denoted WAp. Let the worker's average earnings residual over the window be $\bar{e}_{it}$. Then the volatility of the worker's residual earnings by this measure, assuming an earnings observation for each time period during the window, is

$$V_{it}^{WAp} = \frac{1}{p-1} \sum_{\tau=t-w}^{t+w} (e_{i\tau} - \bar{e}_{it})^2, \tag{2.12}$$

and the corresponding measure of economy-wide volatility of residual earnings is the mean of the worker-specific measures:

$$V_t^{WAp} = \frac{1}{N}\sum_{i=1}^{N} V_{it}^{WAp}. \tag{2.13}$$

Note that a temporal rearrangement of the residual earnings observations within the window does not change the worker-level measure.

Typically, the chosen interval is a nine-year window centered at t, so the resulting WA9 measure at the worker level is

$$V_{it}^{WA9} = \frac{1}{8}\sum_{\tau=t-4}^{t+4} (e_{i\tau} - \bar{e}_{it})^2. \tag{2.14}$$

However, in principle the method could be applied to as few as two observations, as I discuss below.

Moffitt and Gottschalk (2012) interpret the window averaging measure using an earnings dynamics model, taking the worker's mean earnings residual over the window as the worker's permanent earnings for the window, and the deviations from that mean as transitory components. Of course, the permanent component of earnings may change over the window; for example, the worker may change jobs. Because of this, the window averaging measure will be driven by changes in both the permanent and transitory components of earnings, which is why I classify it as a measure of the volatility of residual earnings rather than of the volatility of transitory earnings.[1]

Magnitude of changes in residual earnings Volatility measures that are defined in terms of changes in earnings residuals, and that apply to three or

[1]For example, consider an earnings series that takes on a given value before year t, then a higher value beginning in year t. The window averaging measure applied to a window centered at t will estimate that permanent earnings is between the two values, and will judge the transitory component to have been negative before year t and positive beginning in year t. Therefore, estimated earnings volatility will be positive even though this earnings series contains a change in permanent earnings while the transitory component is always zero.

more time periods, are potentially affected by permutation of a worker's earnings residuals. No such measures have been studied in the literature on earnings volatility, but many combinations of two-period measures would have this property. For example, consider the sum of two consecutive observations of a worker's squared change in residual log earnings, the SC measure introduced above. This measure is defined on three observations of earnings residuals. Denote the volatility as applying to the middle of these three time periods:

$$V_{it}^{\Sigma SC} = V_{it}^{SC} + V_{i,t+1}^{SC} = (e_{it} - e_{i,t-1})^2 + (e_{i,t+1} - e_{it})^2 . \tag{2.15}$$

Relationship between variability of residuals and magnitude of changes This section provides one way of understanding how these two types of volatility of residual earnings are related. I use the special case of three time periods and the window averaging method to show that, for a constant variance of the worker's earnings residuals over the three time periods, a measure of the magnitude of changes is minimized when the two changes in residual earnings are of exactly the same sign and magnitude, and maximized when the two changes are of equal magnitude but opposite sign.

My measure of the variability of earnings residuals during the three time periods is the window averaging measure: the variance of the residuals over that interval. To remain consistent with the empirical portions of this dissertation, I treat the earnings observations as being two years apart, so that the three time periods span five years. Therefore, I denote this measure WA5. Let the three time periods be $t-2$, t, and $t+2$, and let $\bar{e}_i$ be the worker's average residual over those time periods. Then the WA5 measure is

$$V_i^{WA5} = \frac{1}{2}\left[(e_{i,t-2} - \bar{e}_i)^2 + (e_{it} - \bar{e}_i)^2 + (e_{i,t+2} - \bar{e}_i)^2\right] . \tag{2.16}$$

For my measure of the magnitude of earnings changes during those same three time periods, first note that the window averaging approach on two observations can be expressed in terms of the change in the earnings residuals. Consider the window averaging approach applied to $t-2$ and t, denoting the measure WA3 because it spans three years. Then, by simplifying the definition of the window averaging measure,

$$\begin{aligned} V_{it}^{WA3} &= \left(e_{i,t-2} - \frac{e_{i,t-2}+e_{it}}{2}\right)^2 + \left(e_{it} - \frac{e_{i,t-2}+e_{it}}{2}\right)^2 \\ &= \frac{(e_{it}-e_{i,t-2})^2}{2}. \end{aligned} \tag{2.17}$$

Therefore, one measure of the magnitude of earnings changes over the three time periods is the sum of V_{it}^{WA3} and $V_{i,t+2}^{WA3}$, which I denote $V_i^{\Sigma WA3}$. By equation (2.17), this measure is equivalent to

$$V_i^{\Sigma WA3} = \frac{1}{2}\left[(e_{it}-e_{i,t-2})^2 + (e_{i,t+2}-e_{it})^2\right]. \tag{2.18}$$

I will describe how volatility as measured by ΣWA3 may change as the variance of the worker's earnings residuals over the three-period window, which determines WA5, is held constant. As an illustration of the general result, consider the case in which $V_i^{WA5} = 3$. The ΣWA3 measure is minimized when the two changes in earnings residuals have the same sign and magnitude—in this case, the residuals are $\left(-\sqrt{3}, 0, \sqrt{3}\right)$, and $V_i^{\Sigma WA3} = \frac{3}{2}$. The ΣWA3 measure is maximized when the two changes are of equal magnitude but opposite sign—here, the residuals are $(-1, 2, -1)$, and volatility is three times as large, at $V_i^{\Sigma WA3} = \frac{9}{2}$. Again, these two residual vectors give the same amount of volatility under the WA5 measure, even though in the latter vector, the signs of the changes are different and both earnings changes are of larger magnitude compared to the former vector.

For the general result, first note that neither the WA5 measure nor the ΣWA3 measure is affected by adding the same amount to all three earnings residuals. (WA5 is defined in terms of deviations from the mean residual, and WA3 can be expressed in terms of the change in earnings residuals.) Without loss of generality, I assume the mean earnings residual over the three time periods is zero. It will be useful to denote the sum of the worker's squared earnings residuals by c. (The example above was constructed using $c = 6$.)

The WA5 measure simplifies to

$$V_i^{WA5} = \frac{c}{2}. \tag{2.19}$$

Let $d_i = (e_{i,t+2} - e_{it}) - (e_{it} - e_{i,t-2})$ denote the difference in the residual earnings changes. This is a measure of the dissimilarity of the changes; it is zero if the changes are of the same magnitude and sign. The ΣWA3 measure simplifies to

$$V_i^{\Sigma WA3} = \frac{c}{4} + \frac{d_i^2}{12}. \tag{2.20}$$

Therefore, ΣWA3 is minimized when the two earnings changes are exactly the same, so that $d_i = 0$. This occurs when $e_{it} = 0$, $|e_{i,t-2}| = |e_{i,t+2}| = \sqrt{\frac{c}{2}}$, and $e_{i,t-2} + e_{i,t+2} = 0$.

It can be shown that, subject to the c remaining constant, ΣWA3 is maximized when $|e_{it}| = \sqrt{\frac{2c}{3}}$ and $e_{i,t-2} = e_{i,t+2} = -\frac{e_{it}}{2}$. In this case, the earnings changes are of the same magnitude—the largest allowed by the constraint on the sum of squared residuals—but of opposite sign. The dissimilarity index attains its maximum value of $d_i^2 = 6c$, and the ΣWA3 measure reaches $\frac{3c}{4}$. So, while these extreme results give the same amount of volatility by the WA5 measure, one is three times as volatile as the other by the ΣWA3 measure.

2.2.2 Volatility of total earnings

Volatility of total earnings is variability over time in a worker's earnings or log earnings. Types of volatility of total earnings follow the same categorization detailed in Figure 2.1 for volatility of residual earnings. Rather than covering the same distinctions, this subsection briefly details the small literature on volatility of total earnings. All of the measures apply to two time periods.

The Fields and Ok measures Fields and Ok (1996) develop axiomatically a class of measures of non-directional mobility (see section 2.3 for a conceptual distinction between volatility and non-directional mobility). Their preferred measure of a worker's non-directional mobility is the absolute value of the worker's earnings change, $|Y_{it} - Y_{i,t-1}|$. Fields and Ok (1999) recognize that one may be concerned with proportional rather than absolute earnings changes, and they axiomatically justify using the absolute value of the worker's change in log earnings, $|y_{it} - y_{i,t-1}|$, as a measure of non-directional mobility. In both cases, economy-wide non-directional mobility is the mean of worker-specific non-directional mobility in a given time period.

The arc percent change (APC) measure Dahl, DeLeire, and Schwabish (2011), Ziliak, Hardy, and Bollinger (2011), and Dynan, Elmendorf, and Sichel (2012) measure economy-wide volatility of total earnings as the standard deviation of the arc percent change in earnings. Denote the worker's average earnings level in years $t-1$ and t by $\overline{Y}_{it} = \frac{Y_{i,t-1}+Y_{it}}{2}$ The arc percent change in earnings is

$$\Delta Y_{it}^{APC} = \frac{100 \times \Delta Y_{it}}{\overline{Y}_{it}}. \tag{2.21}$$

As Ziliak, Hardy, and Bollinger (2011) point out, the arc percent change is bounded below by -200 percent and above by 200 percent.

Denote the mean arc percent earnings change in year t by $\overline{\Delta Y}_t^{APC}$. The arc percent change measure of the economy-wide volatility of total earnings is

$$V_t^{APC} = sd\left(\Delta Y_{it}^{APC}\right) = \left[\frac{1}{N-1}\sum_{i=1}^{N}\left(\Delta Y_{it}^{APC} - \overline{\Delta Y}_t^{APC}\right)^2\right]^{\frac{1}{2}}. \tag{2.22}$$

As with the SS measure above, the arc percent change measure can be expressed as a function of worker-specific earnings changes and an aggregate earnings change:

$$V_t^{APC} = \left[\frac{1}{N-1}\sum_{i=1}^{N}\left(\Delta Y_{it}^{APC}\right)^2 - \frac{N}{N-1}\left(\overline{\Delta Y}_t^{APC}\right)^2\right]^{\frac{1}{2}}. \tag{2.23}$$

From this expression, it is easy to see that the APC measure arguably understates earnings volatility by excluding the aggregate shock, $\overline{\Delta Y}_t^{APC}$.

This measure is defined on levels of earnings rather than residual log earnings. As a result, one virtue of the measure is that workers who have positive earnings in at least one of the two years will have an earnings volatility observation, whereas measures that work with log earnings restrict the sample to workers with positive earnings in both periods.

2.2.3 Volatility of transitory earnings

Volatility of transitory earnings describes variability in earnings residuals that is due to a transitory component of earnings, rather than a permanent component. Permanent and transitory shocks cannot be distinguished at the worker-year level, so there are no worker-year-level measures of the volatility of transitory earnings. Instead, long panel data on workers is used to estimate the parameters of earnings dynamics models. One such parameter, the variance of the transitory earnings shock, is typically taken as the measure of the volatility of transitory earnings, and is often allowed to vary over time. This volatility is best interpreted as a summary measure of the possibly heterogeneous volatility faced by individual workers; many studies compute volatility of transitory earnings both on the full sample and on various subgroups. In this subsection I briefly discuss the primary method of measuring volatility of transitory earnings.

Estimating parametric earnings models Moffitt and Gottschalk (2012) separate earnings residuals into a permanent component, μ_{it}, with a time-varying factor loading, α_t, and a transitory component, ν_{it}. The permanent component is allowed to follow a random walk with innovations, ω_{it}, and a worker-specific drift, δ_i. The transitory component is modeled as an ARMA(1,1) process with autoregressive parameter ρ and moving average parameter θ. Transitory shocks are denoted by ξ_{it} and have a year-specific loading β_t, so that the combined shock, $\beta_t\xi_{it}$, has year-specific variance $\beta_t^2 \operatorname{Var}(\xi_{it})$. The earnings process is:

$$\begin{aligned} e_{it} &= \alpha_t \mu_{it} + \nu_{it}; \\ \mu_{it} &= \mu_{i,t-1} + \delta_i + \omega_{it}; \\ \nu_{it} &= \rho \nu_{i,t-1} + \beta_t \xi_{it} + \theta \beta_{t-1} \xi_{i,t-1}. \end{aligned} \tag{2.24}$$

See Moffitt and Gottschalk (2012) for a full statement of the model, including the specification of the covariances among shocks. The parameters of the model are estimated by minimizing the distance between the empirical and model-implied autocovariance matrices. The change over time in the volatility of transitory earnings can be seen in the evolution of the estimated β_t. Similar approaches, with some differences in the specification of the shock processes, have been adopted by Moffitt and Gottschalk (1995, 2002), Haider (2001), Stevens (2001), Baker and Solon (2003), and DeBacker et al. (2013).

2.3 Motivation and measurement of earnings volatility

As I discussed in the introduction, the study of earnings volatility is primarily motivated by the consequences of labor market risk. This section briefly explores four ways to better tie the measurement of earnings volatility to this primary motivation. First, volatility of residual earnings is conceptually more attractive than volatility of total earnings or volatility of transitory earnings. Second, life cycle earnings profiles should not absorb aggregate shocks, which are unlikely to be anticipated by workers. Third, life cycle earnings profiles should include worker-specific earnings trends, because research on consumption behavior suggests workers anticipate these trends. Finally, the risk motivation for studying earnings volatility suggests a particular welfare loss of interest. If one is willing to impose some structure on workers' utility functions, this insight can guide the choice of functional form in a earnings volatility measure. The empirical importance of each issue is discussed in chapter 4.

2.3.1 Volatility of residual earnings and total earnings

Of the three types of earnings volatility outlined above, volatility of residual earnings is the most appropriate when risk is the motivation for the research. Volatility of total earnings includes earnings variability due to life cycle dynamics, which is not interesting from a risk perspective. Volatility of transitory earnings attempts to isolate earnings variability due to transitory shocks rather than permanent shocks, but permanent shocks are much more important than transitory shocks when the concern is consumption volatility (Blundell, Pistaferri, and Preston, 2008).

2.3.2 Aggregate shocks

Should the parameters of the life cycle profile change over time? I argue that they should, but only as a continuous function of time. Time-invariant parameters are implausible. Suppose, for example, that the return to education increases, as it has in the U.S. (Acemoglu and Autor, 2011). Eventually, the new return will be incorporated into workers' expectations about the life cycle profiles. However, time-specific parameters are also unattractive, because year-to-year variation in the estimates will absorb aggregate shocks, which are unlikely to to be anticipated by workers, as part of the life cycle profile before earnings volatility is measured. Time-specific parameters also increase the sampling error of each estimate, creating unnecessary noise in measures of volatility.

All previous studies of the volatility of residual earnings or transitory earnings that are cited in this chapter use year-specific coefficients in the life cycle earnings profiles, so aggregate shocks that look like volatility to the worker may be missed in the volatility measurement. My preferred life cycle profile, equation (2.1), allows the parameters of the life cycle profile to change over time, but constrains the changes to follow a smooth path.

2.3.3 Heterogeneous earnings trends

Another issue is how to handle the fact that earnings grow at different rates among observationally equivalent workers. Within a cohort, the cross-sectional variance of earnings increases with experience (Deaton and Paxson, 1994; Guvenen, 2009; DeBacker et al., 2013). The key question in the context of earnings volatility is whether this "fanning out" is a symptom of labor market risk. Are heterogeneous trends the result of a series of unanticipated shocks—for example, from a permanent earnings process that contains a random walk or an AR(1) process with high persistence (MaCurdy, 1982; Abowd and Card, 1989; Hryshko, 2012)? Or do workers have good information on their own growth rate and therefore anticipate the resulting earnings changes?

Guvenen and Smith (2010) use consumption data to show that workers mostly anticipate earnings changes resulting from heterogeneous trends. If this is true, and worker-specific trends are not included in the life cycle earnings profiles, then workers with better- and worse-than-average trends will have persistently higher earnings volatility compared to a worker with an average trend. This could be empirically meaningful: Guvenen (2009) finds that worker-specific trends account for about one third of intracohort inequality at 10 years of potential experience and about two thirds at 30 years of potential experience.

My preferred life cycle profile, equation (2.1), includes heterogeneous trends. Most studies of the volatility of residual or transitory earnings that are cited in this chapter do not include these trends. Exceptions include Drewianka (2010), who includes worker-specific quadratics, and the earnings dynamics model in Moffitt and Gottschalk (2012).

2.3.4 Functional form choice

Different measures of earnings volatility may lead to different conclusions about whether economy-wide earnings volatility rose or fell from one time period to the next, and about the extent to which earnings volatility has changed over a longer time period, as I show in chapter 4. I have already highlighted two reasons that measures of earnings volatility may differ: they may quantify different concepts of earnings volatility, and they may begin with different earnings residuals as the result of different life cycle profile specifications. This subsection addresses a third reason that measures may differ, even for two measures of the same type of earnings volatility and using the same earnings residuals: the functional form of the measure.

The risk motivation for studying earnings volatility implies a natural welfare loss of interest: the loss in utility caused by not being able to smooth consumption as a result of the earnings shocks. This is the difference between utility without liquidity constraints and realized utility. If one is willing to make some assumptions about workers' utility functions, some earnings volatility measures become more attractive than others.

Consider a very basic two-period example. For simplicity, assume no discounting and zero interest rates, and to abstract from the issue of consumption and savings behavior in response to steady life cycle growth in earnings, assume the worker is on a flat portion of his life cycle earnings profile, so that it does not matter whether one works with the level of earnings or the deviation of earnings from a life cycle profile.

Suppose the worker receives earnings $I - v$ in the first period and $I + v$ in the second period, for some $v \in [0, I]$. Denote utility as a function of consumption in each period by $u(c)$, and total utility in the two periods by $U(c)$. Without liquidity constraints, the worker would consume I in each period for total utility $U(c^{\star}) = 2u(I)$. Suppose that with liquidity constraints, the worker can smooth all but a fraction $\kappa \in [0, 1]$ of the shock, so that total utility is $U(c') = u(I - \kappa v) + u(I + \kappa v)$. Of interest is the welfare loss $U(c^{\star}) - U(c')$.

Clearly, the functional form of the welfare loss depends on the worker's utility function. However, the loss can be approximated in the area of $v = 0$ using a Taylor series. With a second-order Taylor series approximation, the welfare loss depends on the square of v but not separately on v itself for many commonly used utility functions. For constant relative risk aversion (CRRA) utility with risk aversion parameter $\theta \geq 0$, the approximation to the welfare loss is $\frac{\kappa\theta v^2}{I^{\theta+1}}$; for the log utility case ($\theta = 1$) this simplifies to $\frac{\kappa v^2}{I^2}$.

The form of this welfare loss suggests that it is earnings shocks relative to smoothed earnings that matter, rather than the magnitudes of the shocks themselves, and that it is the squares of these normalized shocks that are of interest. So, an attractive two-period earnings volatility measure will use the square of changes in log earnings or residual log earnings.

Consider two measures: the two-period window averaging approach that takes the variance of the two earnings residuals, or an alternative measure that uses the standard deviation of those same residuals. The same approximation approach reveals that in the neighborhood of $v = 0$, the standard deviation approach follows $\frac{\sqrt{2}v}{I}$, while the variance approach follows $\frac{2v^2}{I^2}$. Also, it is straightforward to show that the absolute change in log earnings or residual log earnings follows $\frac{2v}{I}$, while the squared change in log earnings or residual log earnings follows $\frac{4v^2}{I^2}$. This means that when risk is the motivation for studying earnings volatility, measures such as WA and SC are preferable to the square root of WA and to AC, SS, or APC (recall that the last two are defined by cross-worker standard deviations rather than variances, so the units are log dollars instead of squared log dollars).[2]

This exercise illustrates why the most attractive measures of earnings volatility may be different from preferred measures of other concepts of earnings changes, such as some types of earnings mobility. Measures of the volatility of residual earnings should be chosen while keeping in mind the welfare loss

[2]Classical or mean-reverting measurement error will cause measured earnings volatility to overstate true earnings volatility, but will not affect these arguments about which volatility measures are proportional to the welfare loss.

caused by the earnings volatility: the difference between utility without liquidity constraints and realized utility. But measures of earnings mobility, for example, may be chosen while keeping in mind the welfare change from one period to another; Fields and Ok (1999) have shown that absolute changes in log earnings have attractive properties in the case of non-directional mobility.

CHAPTER 3

DATA

This chapter details the data sources used in this dissertation. All of the descriptive results in chapter 4, on the change over time in economy-wide earnings volatility and on how earnings volatility differs across, are from the PSID. The final chapter accounts for the rise in economy-wide earnings volatility among the PSID sample, and uses data from the CPS to measure labor demand shocks.

3.1 PSID

I use data from the Panel Study of Income Dynamics covering interview years 1970–2009.[1] The labor earnings measure is wage and salary earnings, which is the most consistently measured labor earnings series in the PSID. Therefore, earnings sources such as self-employment income and bonuses are excluded.

I include in the earnings sample all worker-years for men in the PSID cross section sample (that is, excluding the Survey of Economic Opportunity sample; see Shin and Solon (2011), footnote 11) who, in the earnings reference year, are heads of household, are age 25 to 59, are not students, are not self-employed in their main job, and have positive earnings and at least 500 hours worked. I drop observations in the top or bottom one percent of earnings in each year, which is designed to get rid of severe outliers and lessen the effect of earnings top-codes.

Following the recommendation of Brown and Light (1992), I clean the job tenure series by assigning a new job if and only if reported tenure is less than the time since the previous interview, then adjusting the following tenure ob-

[1] Interview year 2009 was the latest year available when I began the empirical portion of this dissertation.

servations within a job to be consistent with this first report. PSID sample weights are used throughout.

After annual interviews during 1968–1997, the PSID began biannual interviews in 1999. This is unfortunate for my purposes because biannual interviews yield only five observations over a nine-year window. This smaller sample size could make earnings volatility measurements noisier after 1997 compared to before 1997. To exclude this possibility, I drop all even-numbered interview years before calculating earnings volatility and performing any of the analysis. That is, I keep biannual interviews for 1971–2009.

This yields an earnings sample of 23,306 worker-year observations over 20 interview years, from 3,506 unique workers. I use this sample to estimate the life cycle earnings profiles specified in section 2.1. Finally, I calculate earnings volatility, which requires earnings data for each of the five interview years in the nine-year period.[2] The earnings volatility is assigned to the midpoint of this window, so that I have earnings volatility measures every two years beginning in 1975 (using the interval 1971–1979) and ending in 2005 (using the interval 2001–2009). In all there are 9,479 volatility observations, which come from 1,816 workers.

Descriptive statistics are in Table 3.1. Note that labor demand shocks are identified using CPS data, discussed below.

[2] As I discuss further below, trends in economy-wide earnings volatility over time are similar if the standard for calculating volatility is relaxed to having at least four or at least three observations during the window.

Table 3.1: Variable means by subgroup and time period

	1975–2005	1975–1981	1983–1989	1991–1997	1999–2005
Job change					
No job change	**0.529**	0.486	0.537	0.577	0.510
Voluntary job change	**0.323**	0.336	0.289	0.290	0.367
Involuntary job change	**0.148**	0.178	0.174	0.133	0.123
Demand shock index					
Low shock index	**0.500**	0.847	0.487	0.254	0.501
High shock index	**0.500**	0.153	0.513	0.746	0.499
Education					
High school or less	**0.396**	0.493	0.393	0.360	0.368
Some college	**0.254**	0.212	0.252	0.260	0.276
College or more	**0.350**	0.296	0.355	0.380	0.356
Experience					
6–15 years	**0.292**	0.336	0.371	0.261	0.234
16–25 years	**0.395**	0.294	0.400	0.480	0.382
26–39 years	**0.313**	0.369	0.229	0.258	0.384
Unemployment rate					
Below 6 percent	**0.525**	0.235	0.191	0.561	0.915
At or above 6 percent	**0.475**	0.765	0.809	0.439	0.085
Married					
Not married	**0.179**	0.089	0.157	0.200	0.234
Married	**0.821**	0.911	0.843	0.800	0.766
Union membership					
Non-union	**0.790**	0.737	0.750	0.795	0.847
Union	**0.210**	0.263	0.250	0.205	0.153
Occupation					
Managers	**0.206**	0.186	0.213	0.237	0.188
Professionals	**0.173**	0.145	0.166	0.179	0.191
Technicians	**0.047**	0.043	0.048	0.052	0.045
Sales	**0.043**	0.040	0.035	0.036	0.057
Office and admin.	**0.066**	0.067	0.071	0.060	0.068
Prod./craft/repair	**0.229**	0.268	0.249	0.217	0.199
Operators/laborers	**0.180**	0.209	0.171	0.160	0.185
Services	**0.055**	0.042	0.046	0.059	0.066

Table 3.1: Variable means by subgroup and time period (continued)

	1975–2005	1975–1981	1983–1989	1991–1997	1999–2005
Industry					
Ag./extr./constr.	**0.028**	0.027	0.029	0.027	0.028
Manufacturing	**0.076**	0.071	0.071	0.074	0.085
Trans./comm./util.	**0.329**	0.415	0.344	0.304	0.286
Wholesale/retail	**0.115**	0.109	0.127	0.119	0.106
FIRE/prof. svc.	**0.139**	0.126	0.127	0.145	0.150
Public administration	**0.032**	0.029	0.028	0.031	0.037
Employer tenure					
0–5 years	**0.420**	0.439	0.479	0.373	0.404
5–10 years	**0.251**	0.254	0.279	0.266	0.215
>10 years	**0.330**	0.307	0.242	0.361	0.381
Occupation tenure					
0–5 years	**0.252**	0.301	0.267	0.218	0.239
5–10 years	**0.282**	0.316	0.323	0.275	0.237
>10 years	**0.466**	0.383	0.410	0.507	0.524
Industry tenure					
0–5 years	**0.248**	0.278	0.251	0.232	0.241
5–10 years	**0.279**	0.323	0.306	0.267	0.244
>10 years	**0.473**	0.400	0.443	0.501	0.515
Observations	**9479**	2324	2372	2334	2449

Notes: Sample is male heads of household in the PSID; see text for details. Variables are defined for nine-year intervals. See Table 5.1 for details. The intervals are denoted by their midpoint, so the first window is centered at 1975 (using data for 1971–1979) and the final window is centered at 2005 (using data for 2005–2009).

3.2 CPS

I use March Current Population Survey IPUMS data for workers age 16–64 in years 1971–2009 (King et al., 2010). The CPS data is used to measure the initial distribution of occupation-industry sectors within demographic groups, and to measure the subsequent labor demand shocks for occupation-industry sectors. The initial distribution and subsequent shocks are based on hours worked. Prior to 1976, weeks worked in the previous year was an interval variable, and hours worked referred to the previous week. I multiply the interval midpoint by hours in the previous week to estimate hours worked during the previous year. Beginning in 1976, I calculate hours worked in the previous year by multiplying weeks worked in the previous year by the usual number of hours worked. I use the harmonized occupation (OCC1990) and industry (IND1950) variables available in IPUMS to classify workers into occupation-industry sectors.

CHAPTER 4

EARNINGS VOLATILITY IN THE UNITED STATES

This chapter presents empirical evidence on earnings volatility in the U.S. Section 4.1 looks at economy-wide earnings volatility and its components over time, section 4.2 is about which workers have higher earnings volatility and which workers have experienced larger increases in earnings volatility, and section 4.3 assesses the empirical importance of the methodological recommendations I made in chapter 2.

4.1 Economy-wide earnings volatility over time

In this section, I show that economy-wide volatility of residual earnings increased by my preferred measure among men in the U.S. between 1975 and 2005. I show that this finding is robust to changing the length of the window over which workers' earnings volatility is computed, to including workers who do not have earnings observations for each time period in the window, and to dropping earnings volatility observations based on any imputed values of earnings. Finally, I show that about 70 percent of the increase in economy-wide earnings volatility can be explained by rising volatility of hourly wages, with the rest explained by higher volatility of hours worked.

4.1.1 Has earnings volatility increased?

Figure 4.1 displays the trend in economy-wide volatility of residual earnings among male heads of household in the U.S. from 1975 to 2005. The upward trend is clear and there are sustained increases from 1975–1987 and 1993–2001. Although the rate of increase has declined over time, economy-wide earnings

Figure 4.1: Economy-wide volatility of residual earnings by year

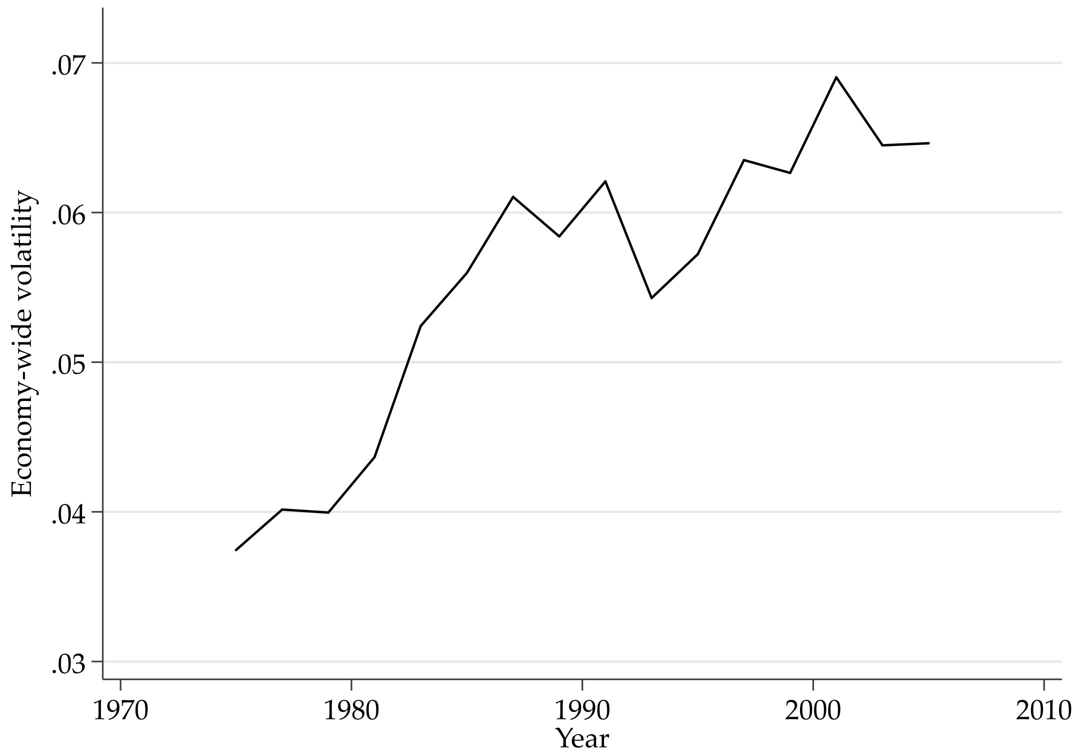

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

volatility was higher in each year during 1997–2005 than at any time in the sample before 1997.[1]

Earlier studies found that economy-wide earnings volatility for U.S. males has increased since the 1970s, with little evidence of an increase after the mid-1980s (Gottschalk and Moffitt, 1994; Haider, 2001; Stevens, 2001; Moffitt and Gottschalk, 2002). Moffitt and Gottschalk (2012) do find an uptick in the late 1990s before their sample ends in 2000; the trend in Figure 4.1 confirms this.

[1]To the extent that some unexplained earnings variation is due to measurement error, the proportional change in "true" earnings volatility was even larger.

4.1.2 Robustness of earnings volatility trend

The earnings volatility measurement process described in chapter 2 involves some arbitrary decisions, so it is useful to check whether these choices affect the finding that economy-wide volatility of residual earnings increased significantly over the sample period.

The decision to use a nine-year window for computing earnings volatility attempts to balance the need for a good estimate of the variance of residual earnings with the fact that the variance itself may change over time. Figure 4.2 shows that the trend in economy-wide volatility of residual earnings is affected very little by changing the volatility window to five or 13 years.

In computing earnings volatility over a nine-year period, I have required that all five possible earnings observations within the window are available. This disproportionately excludes workers who experience a long spell of non-employment, as well as workers who are more likely to leave the sample for a variety of reasons. Figure 4.3 shows that when the requirement for an earnings volatility observation is relaxed to at least three or at least four earnings observations within the nine-year window, the increase in economy-wide earnings volatility over the sample period is at least as large.

A final concern is that the trend in economy-wide earnings volatility has been influenced by changes in the way that PSID imputes missing or unreliable earnings observations, or by a change in the incidence of such observations. The fraction of earnings observations that were imputed rose from about 1 percent in the early 1970s to about 3 percent in the late 2000s, and the fraction of nine-year volatility windows in which at least one of the five earnings observations was imputed rose from about 3 percent to about 8 percent over the sample period. Figure 4.4 shows that the trend in economy-wide earnings

Figure 4.2: Economy-wide volatility of residual earnings by year for various window widths

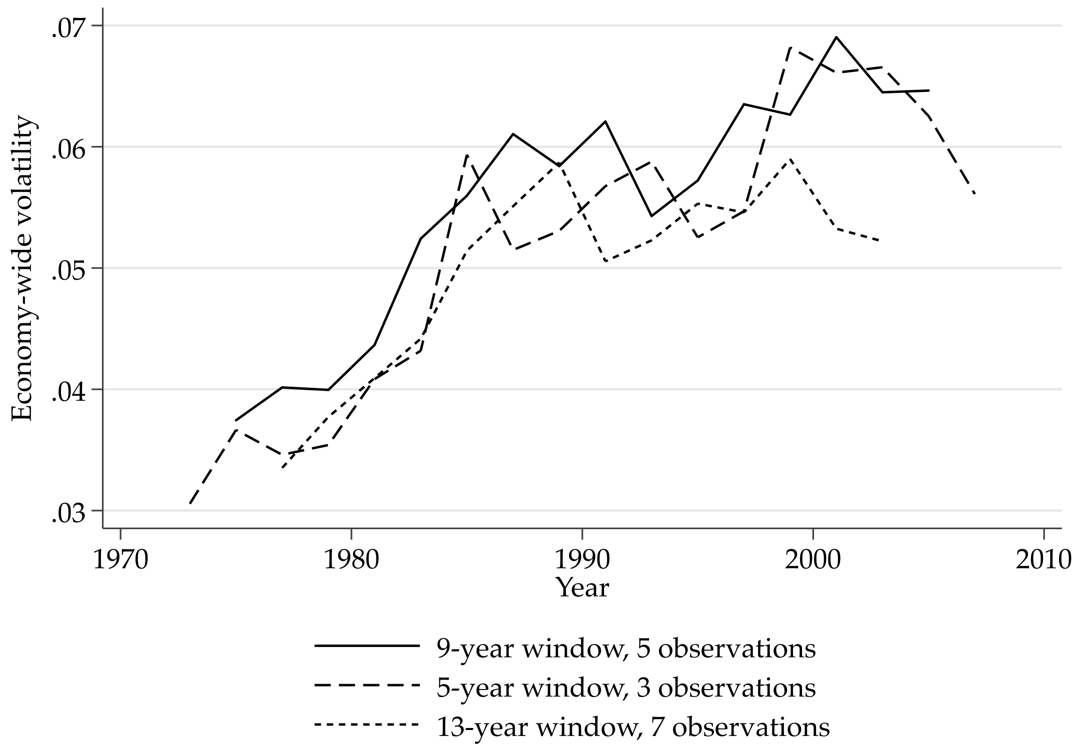

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a window of the specified width centered on that year. The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details. Earnings observations are every other year for reasons discussed in chapter 3.

volatility changes very little when imputed earnings observations are dropped.

4.1.3 Components of earnings volatility

Before analyzing the contributions of various mechanisms to the trend shown in Figure 4.1, it is useful to understand whether rising volatility of residual earnings is being driven by increases in the volatility of hourly wages, volatility

Figure 4.3: Economy-wide volatility of residual earnings by year for various minimum observation requirements

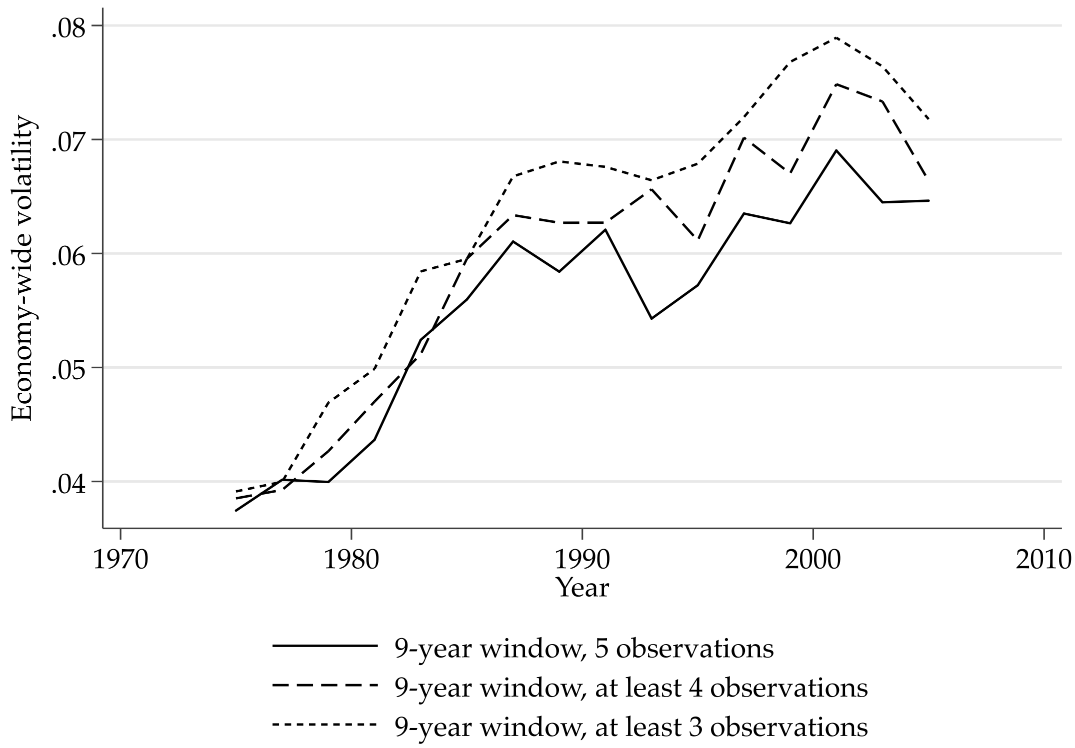

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details. Earnings observations are every other year for reasons discussed in chapter 3.

of hours worked, or both. The answer is helpful for forming hypotheses about mechanisms and also serves as a rough check on the detailed results discussed in chapter 5. If volatility of wages alone explains the rise in economy-wide earnings volatility, for example, that suggests that business cycle amplitudes are not a major explanation, and that wage-related explanations should dominate in the empirical results.

Some previous research has considered the separate contributions of volatil-

Figure 4.4: Economy-wide volatility of residual earnings by year for full and non-imputed samples

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details. The proportion of volatility observations that depend on at least one imputed earnings value rises from about 3 percent at the beginning of the sample to about 8 percent at the end of the sample.

ity of wages and hours to earnings volatility. Comparing nine-year windows centered at 1974 and 1983, Gottschalk and Moffitt (1994) computed volatility of both weekly wages and weeks worked, finding that volatility of wages explained about 60 percent of the rise in economy-wide earnings volatility. Haider (2001) found that volatility of wages and hours contributed roughly equally to the rise in economy-wide earnings volatility.

Because log earnings are the sum of log wages and log hours, the variance of a worker's log earnings over a given window is easily decomposed into three variance components: the variance of log wages, the variance of log hours, and a wages-hours covariance term. For worker i, let y be log real earnings, w be log real wages, and h be log hours. Using an excessive amount of notation to emphasize that the variances are worker-specific and apply to a window centered at time t, the decomposition is

$$\begin{aligned}\mathrm{Var}_{it}(y_{i,t-4},\ldots,y_{i,t+4}) = \mathrm{Var}_{it}(w_{i,t-4},\ldots,w_{i,t+4}) + \mathrm{Var}_{it}(h_{i,t-4},\ldots,h_{i,t+4}) \\ + 2\mathrm{Cov}_{it}\big((w_{i,t-4},\ldots,w_{i,t+4}),(h_{i,t-4},\ldots,h_{i,t+4})\big).\end{aligned} \tag{4.1}$$

This decomposition can be approximated by expressing the volatility of residual earnings as the sum of the volatilities of residual wages and residual hours, plus a covariance term. Volatilities of residual wages and residual hours are calculated by replacing the left hand side of the life cycle earnings profile, equation (2.1), with log real wages and log annual hours, respectively, then applying the same window averaging measure, equation (2.2), to the residuals.

The resulting empirical decomposition is not equivalent to the variance expression above because volatility is calculated not from raw observations, which appear in (4.1), but from residuals after removing life cycle effects. Denote volatility of residual earnings, residual wages, and residual hours for worker i in year t by V^y_{it}, V^w_{it}, and V^h_{it}, respectively. Then the decomposition to be implemented is

$$V^y_{it} = V^w_{it} + V^h_{it} + \nu_{it}, \tag{4.2}$$

where ν represents a forcing term that captures in part the covariance term in equation (4.1) and in part the fact that the variances in equation (4.1) do not exactly correspond to the volatility terms in equation (4.2), as just explained.

Taking averages across workers, economy-wide volatility of residual earn-

Table 4.1: Contributions of components to change in economy-wide volatility of residual earnings

	1975 level	2005 level	Change	Percent contribution
Residual earnings	0.0374	0.0646	0.0272	
Residual wages	0.0375	0.0571	0.0196	72.2
Residual hours	0.0271	0.0362	0.0091	33.6
Wages-hours covariance	−0.0272	−0.0287	−0.0016	−5.8

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Volatility of residual wages and volatility of residual hours are calculated using the same method applied to log real wages and log hours, respectively, instead of log real earnings. Sample is male heads of household in the PSID; see chapter 3 for details. *Percent contribution* is the ratio of the change in the component to the change in volatility of residual earnings.

ings in year t can be decomposed as

$$V_t^y = V_t^w + V_t^h + \nu_t, \tag{4.3}$$

where each term is the cross-worker mean for the specified year.

Figure 4.5 shows the trends of each component of equation (4.3). For most of the sample period, the residual hours and covariance terms are roughly constant, while the volatilities of residual earnings and residual wages rise at similar rates. Table 4.1 shows the precise contribution of each term to explaining the rise of economy-wide volatility of residual earnings between 1975 and 2005. The increase in economy-wide volatility of residual wages explains more than 70 percent of the increase in economy-wide volatility of residual earnings over this time. Increasing economy-wide volatility of residual hours worked contributes about one third, and the covariance term plays very little role.

Figure 4.5: Components of economy-wide volatility of residual earnings by year

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Volatility of residual wages and volatility of residual hours are calculated using the same method applied to log real wages and log hours, respectively, instead of log real earnings. Sample is male heads of household in the PSID; see chapter 3 for details.

4.2 Levels and trends of earnings volatility across workers

Table 4.2 shows mean volatility of residual earnings by subgroup and for five time periods: the full sample period and four sub-periods. The first column of results compares volatility of residual earnings, by my preferred measure, across workers for many different sample partitions. For the full sample period, there are significant differences in earnings volatility by earnings quintile, by occupation, and for many other cross-worker comparisons. I focus on some of these results in detail later in this section. These findings are also interesting as a starting point for considering why economy-wide earnings volatility increased over time, which is discussed further in chapter 5, and I also present these comparisons to set up the methodological exercises in section 4.3 below.

Table 4.2: Mean volatility of residual earnings by subgroup and time period

	1975–2005	1975–1981	1983–1989	1991–1997	1999–2005	Change over time	*p*-value of test of increase over time
All workers	**0.057**	0.040	0.057	0.059	0.065	*0.025*	*0.000*
Earnings quintile							
Quintile 1	**0.102**	0.073	0.116	0.108	0.101	*0.028*	*0.002*
Quintile 2	**0.066**	0.055	0.066	0.070	0.071	*0.016*	*0.006*
Quintile 3	**0.053**	0.040	0.053	0.060	0.056	*0.016*	*0.005*
Quintile 4	**0.048**	0.028	0.036	0.054	0.067	*0.039*	*0.000*
Quintile 5	**0.038**	0.022	0.035	0.032	0.050	*0.028*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Job change							
No job change	**0.035**	0.023	0.032	0.036	0.043	*0.021*	*0.000*
Voluntary job change	**0.075**	0.041	0.073	0.089	0.085	*0.044*	*0.000*
Involuntary job change	**0.097**	0.086	0.109	0.096	0.096	*0.010*	*0.146*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Demand shock index							
Low shock index	**0.052**	0.041	0.057	0.056	0.057	*0.017*	*0.000*
High shock index	**0.062**	0.038	0.057	0.061	0.073	*0.035*	*0.000*
p-value of equal means test	***0.000***	*0.614*	*0.908*	*0.405*	*0.003*		
Education							
High school or less	**0.054**	0.041	0.057	0.061	0.057	*0.016*	*0.000*
Some college	**0.059**	0.051	0.057	0.054	0.067	*0.016*	*0.008*
College or more	**0.059**	0.031	0.057	0.061	0.072	*0.041*	*0.000*
p-value of equal means test	***0.115***	*0.000*	*0.988*	*0.396*	*0.027*		
Experience							
6–15 years	**0.064**	0.049	0.062	0.061	0.082	*0.033*	*0.000*
16–25 years	**0.053**	0.037	0.053	0.056	0.058	*0.021*	*0.000*
26–39 years	**0.055**	0.035	0.057	0.064	0.062	*0.027*	*0.000*
p-value of equal means test	***0.000***	*0.001*	*0.254*	*0.457*	*0.000*		
Unemployment rate							
Below 6 percent	**0.060**	0.043	0.051	0.057	0.066	*0.023*	*0.000*
At or above 6 percent	**0.053**	0.040	0.059	0.062	0.057	*0.017*	*0.004*
p-value of equal means test	***0.007***	*0.394*	*0.134*	*0.311*	*0.201*		
Married							
Not married	**0.077**	0.050	0.081	0.072	0.085	*0.035*	*0.000*
Married	**0.053**	0.039	0.053	0.056	0.059	*0.020*	*0.000*
p-value of equal means test	***0.000***	*0.095*	*0.000*	*0.016*	*0.000*		

Table 4.2: Mean volatility of residual earnings by subgroup and time period (continued)

	1975–2005	1975–1981	1983–1989	1991–1997	1999–2005	Change over time	*p*-value of test of increase over time
Union membership							
Non-union	**0.058**	0.036	0.060	0.061	0.065	*0.029*	*0.000*
Union	**0.054**	0.052	0.047	0.052	0.067	*0.015*	*0.098*
p-value of equal means test	***0.360***	*0.000*	*0.003*	*0.247*	*0.835*		
Occupation							
Managers	**0.054**	0.033	0.060	0.061	0.054	*0.021*	*0.000*
Professionals	**0.052**	0.021	0.026	0.051	0.083	*0.062*	*0.000*
Technicians	**0.047**	0.031	0.050	0.050	0.050	*0.019*	*0.034*
Sales	**0.084**	0.055	0.130	0.080	0.078	*0.023*	*0.052*
Office and admin.	**0.060**	0.040	0.057	0.068	0.067	*0.028*	*0.019*
Prod./craft/repair	**0.056**	0.046	0.064	0.053	0.060	*0.014*	*0.005*
Operators/laborers	**0.061**	0.055	0.061	0.066	0.062	*0.008*	*0.104*
Services	**0.061**	0.029	0.057	0.073	0.066	*0.037*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.227*	*0.067*		
Industry							
Ag./extr./constr.	**0.062**	0.031	0.053	0.069	0.083	*0.052*	*0.008*
Manufacturing	**0.068**	0.061	0.073	0.068	0.068	*0.007*	*0.268*
Trans./comm./util.	**0.055**	0.042	0.061	0.058	0.060	*0.018*	*0.000*
Wholesale/retail	**0.051**	0.039	0.046	0.049	0.064	*0.025*	*0.002*
FIRE/prof. svc.	**0.069**	0.049	0.070	0.079	0.071	*0.022*	*0.001*
Public administration	**0.062**	0.023	0.068	0.054	0.083	*0.060*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Employer tenure							
0–5 years	**0.072**	0.052	0.070	0.078	0.082	*0.030*	*0.000*
5–10 years	**0.049**	0.034	0.049	0.045	0.065	*0.032*	*0.000*
>10 years	**0.044**	0.029	0.041	0.051	0.047	*0.018*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Occupation tenure							
0–5 years	**0.075**	0.055	0.072	0.087	0.084	*0.030*	*0.000*
5–10 years	**0.054**	0.040	0.055	0.048	0.073	*0.033*	*0.000*
>10 years	**0.049**	0.029	0.049	0.054	0.053	*0.024*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Industry tenure							
0–5 years	**0.073**	0.054	0.076	0.077	0.081	*0.027*	*0.000*
5–10 years	**0.055**	0.041	0.052	0.048	0.074	*0.033*	*0.000*
>10 years	**0.050**	0.030	0.050	0.057	0.054	*0.024*	*0.000*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.000*	*0.000*		
Observations	**9479**	2324	2372	2334	2449		

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Sample is male heads of household in the PSID; see chapter 3 for details. Column labeled *Change over time* gives change in mean volatility among the specified group from 1975–1981 to 1999–2005. Column labeled *p-value of test of increase over time* gives p-values from t tests in which the alternative hypothesis is that mean volatility for the specified group was higher in 1999–2005 than in 1975–1981. Rows labeled *p-value of equal means test* give p-values from F tests of equality of means among subcategories.

For almost all subgroups, mean volatility rises in each successive sub-period. Gottschalk and Moffitt (1994) also found a broad increase in earnings volatility using a smaller number of subgroups. The rightmost column of Table 4.2

shows the p-value of a t test in which the alternative hypothesis is that earnings volatility was higher in the final sub-period, 1999–2005, than in the first sub-period, 1975–1981. No subgroup experienced a decrease in earnings volatility, and the vast majority of the increases are significant at the 1 percent level. This shows that the increase in volatility of residual earnings is a very broad phenomenon, and is unlikely to be explained by simple compositional effects such as changes in the distribution of education or occupation.

Before reviewing the results in depth, note that comparisons of earnings volatility across workers are less reliable than comparisons across time. The reason is that measurement error in earnings contributes to observed earnings volatility, and the type or degree of measurement error differs across workers (Duncan and Hill, 1985; Bound et al., 1994; Pischke, 1995), which complicates cross-worker comparisons of earnings volatility. Comparisons of earnings volatility over time within a defined population are not affected by heterogeneous measurement error, but would be affected by changes in measurement error over time.

In this section I look closely at earnings volatility across the earnings distribution and at earnings volatility by job change and by sector. There are some other results of interest in Table 4.2, including the absence of significant differences in the volatility of residual earnings across education groups or between union and non-union workers. Many of these results are discussed in more detail when I introduce the explanatory variables of interest in chapter 5.

4.2.1 Earnings volatility across the earnings distribution

One important result in Table 4.2 is that volatility of residual earnings is much higher among the lowest earnings quintile than among any other quintile. This is not a new finding (Gottschalk and Moffitt, 1994), but it is worth emphasizing, because it implies that the difference in welfare between high and low earners is understated if earnings volatility is not taken into account.

Figure 4.6 illustrates this clearly by showing mean volatility of residual earnings by ventile (five-percentile bins) of earnings, in the initial and final years of my sample. It is clear that, in both years, volatility of residual earnings is greater for low earners than for high earners. Another interesting result in this figure is that, while earnings volatility rose between 1975 and 2005 at almost all earnings ventiles, the increase was largest in the bottom fifth of the earnings distribution. That suggests that the welfare gap between low and high earners grew even more than would be suggested by the well-documented increase in cross-sectional earnings inequality (Autor, Katz, and Kearney, 2008).

To quantify how the welfare losses from earnings volatility differ across the earnings distribution and over time, I compare observed earnings and certainty equivalent earnings for workers in each earnings quintile, separately for the initial and final years of my sample. I assume constant relative risk aversion (CRRA) utility functions with parameter $\theta = 2$.[2] To measure certainty equivalent earnings for workers in a given quintile, I first compute expected utility if earnings follow a lognormal distribution, with the mean and variance of log earnings chosen to be typical of the earnings quintile. The most natural mean to choose for this log earnings distribution is median log earnings within the earnings quintile in the specified year. For the variance of the distribution, I use the mean of earnings volatility within the earnings quintile in the specified year (recall that worker-specific earnings volatility, by my preferred measure, is a variance, so the cross-worker mean represents a typical variance of resid-

[2] Estimates of this parameter vary, with plenty of widely-cited evidence for somewhat lower (Chetty, 2006) and higher (Barsky et al., 1997) values.

Figure 4.6: Volatility of residual earnings by earnings ventile

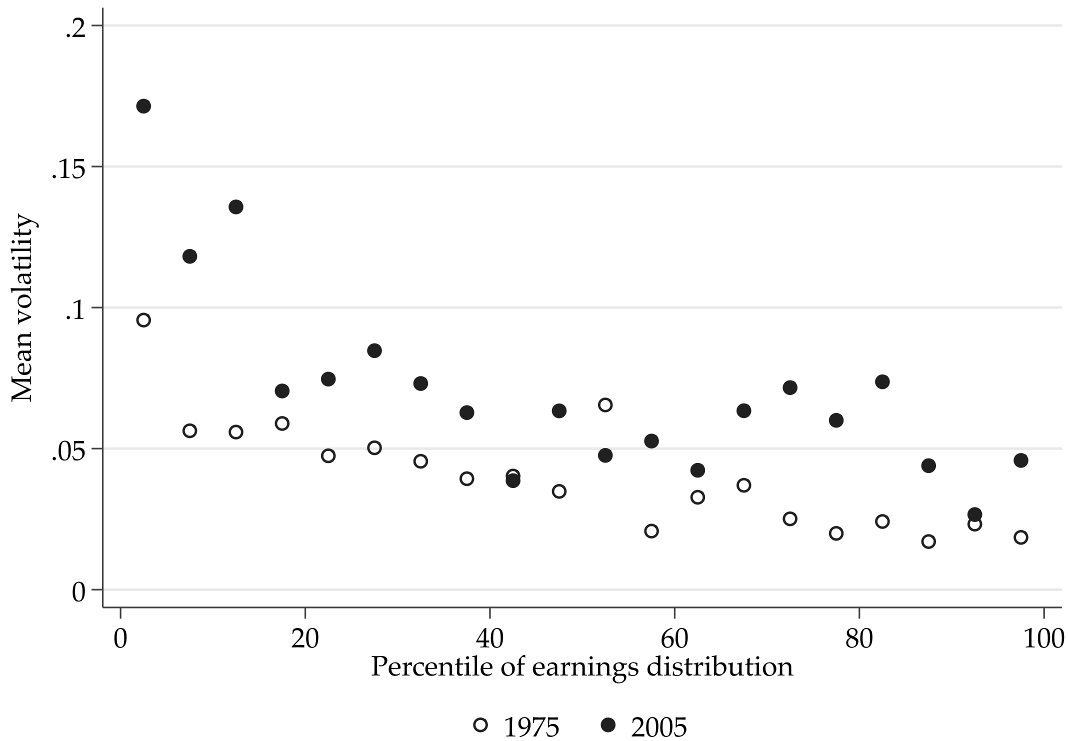

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). The fitted values from these profiles are used to identify earnings ventiles for each year. Sample is male heads of household in the PSID; see chapter 3 for details.

ual log earnings within the earnings quintile). Certainty equivalent earnings is the level of earnings that, with zero earnings volatility, would yield the same amount of utility as the expected utility just described. The difference between observed median earnings and certainty equivalent earnings is the amount that the worker at the median of the earnings quintile would be willing to pay to reduce earnings volatility to zero.

For this exercise, I assume that workers consume exactly what they earn in each period. This assumption is unrealistic in many cases, but it will not affect comparisons of welfare losses over time, and it will give a lower bound on differences in welfare losses across earnings quintiles if, as seems likely, lower earners find it more difficult to smooth consumption in the face of volatile earnings (Blundell, Pistaferri, and Preston, 2008). In reality, it may be more accurate to assume that lower earners have greater risk aversion, because they have more limited or expensive access to consumption smoothing through the financial system. In this case, the welfare losses among low earners, for a given level of earnings volatility, would be higher than the values I report here.

Table 4.3 shows the results of this exercise. In 1975, the median worker in the lowest quintile of earnings would be willing to pay about 3 percent of earnings to eliminate volatility of residual earnings, while the median worker in the highest quintile of earnings would pay about 1 percent. Between 1975 and 2005, willingness to pay increases at each quintile, and the increase is largest among the lowest earnings. By 2005, the median worker in the first earnings quintile would be willing to pay almost 6 percent of earnings to eliminate volatility of residual earnings.

4.2.2 Earnings volatility by job change and by sector

A fundamental question about the increase in economy-wide volatility of residual earnings is to what extent it reflects changes in the incidence and consequences of job change, and to what extent it is a result of changing pay setting within jobs. This subsection discusses some raw descriptive evidence on this topic, and I address the issue more closely in chapter 5.

If job change is the reason for rising economy-wide volatility of residual earnings, the proximate causes can be more frequent job change, more earnings volatility among job changers, or both. After discussing in chapter 3 the

Table 4.3: Welfare losses from earnings volatility by earnings quintile

	1975			2005		
	Median earnings	Certainty equivalent earnings	Willingness to pay as percentage of median earnings	Median earnings	Certainty equivalent earnings	Willingness to pay as percentage of median earnings
Earnings quintile 1	26,921.52	26,108.43	3.02	24,711.44	23,278.82	5.80
Earnings quintile 2	39,036.21	38,170.27	2.22	38,096.80	36,709.91	3.64
Earnings quintile 3	47,112.66	46,165.49	2.01	51,482.17	50,195.01	2.50
Earnings quintile 4	56,871.71	56,026.18	1.49	66,926.82	64,948.56	2.96
Earnings quintile 5	74,034.18	73,260.07	1.05	103,993.98	101,584.72	2.32

Fitted values from life cycle earnings profiles, specified in equation (2.1), are used to identify earnings quintiles for each year. Computation of certainty equivalent earnings assumes that utility is CRRA with risk aversion parameter $\theta = 2$, and that the worker's earnings follow a lognormal distribution where the mean of log earnings is median log earnings within the worker's earnings quintile and where the variance of log earnings is given by mean earnings volatility within the worker's earnings quintile in the specified year. Certainty equivalent earnings is the level of earnings that, with zero earnings volatility, would yield the same amount of utility as the expected utility just described. Sample is male heads of household in the PSID; see chapter 3 for details.

PSID data I use, I presented in Table 3.1 means of the variables used in this dissertation, for the full sample period and four sub-periods. The top of the table shows that at the beginning of the sample period, in nine-year windows centered on 1975 to 1981, just under half of all workers experienced any job change. For the last sub-period, nine-year windows centered on 1999 to 2005, the fraction experiencing no job change over the window actually rose a few percentage points. Incidence of voluntary job change also increased a few percentage points, while incidence of involuntary job change (due to establishment closing or layoff) decreased.

Turning now to the amount of volatility of residual earnings experienced by job stayers and job changers, the results in Table 4.2 show that, not surprisingly, job changers experience higher volatility of residual earnings than job stayers. For job stayers, mean volatility of residual earnings rose by 0.021 between the first and last sub-periods of my sample, which is just under the economy-wide increase of 0.025. Voluntary job changes are associated with less earnings volatility than involuntary job changes, so the shift toward more voluntary changes and fewer involuntary changes would tend to decrease economy-wide volatility of residual earnings. Overall, the incidence and consequences of job change do not appear to be strong proximate explanations for rising earnings volatility.

Another class of explanations for rising economy-wide volatility of residual earnings involves changes in pay setting within jobs. One clue to such changes might be found in whether earnings volatility rose at different rates in different occupation or industry sectors. Table 4.2 shows mean volatility of residual earnings for eight occupation categories and six industry categories, for the full sample period and four sub-periods. There are significant differences across occupations and industries for the full period and for most sub-periods. Mean earnings volatility rose between the first and last sub-periods for all categories, but there is no clear pattern to the increases. For example, managers experienced a below-average increase (0.021 versus 0.025), while professionals experienced the largest increase (0.062).

These results are for all workers, including those who changed jobs during the nine-year window over which earnings volatility is measured. To get cleaner evidence on pay setting within jobs across occupation and industry sectors, in Table 4.4 I restrict the sample to nine-year windows in which the worker did not change jobs. It is still the case that there are significant differences across occupations and industries for the full period and for most sub-periods. Most occupation and industry categories experienced an increase in volatility of residual earnings between the first and last sub-periods but, again, there is no clear pattern to the increases.

4.3 Effects of measurement choices

This section assesses the empirical importance of each of the four methodological issues discussed in chapter 2: when risk is the motivation for studying earnings volatility, volatility of residual earnings is more interesting than volatility of total earnings, life cycle earnings profiles should not absorb aggregate shocks that appear as volatility to workers, the profiles should absorb heterogeneous earnings trends, and volatility measures should be based on variances of residual earnings or squared changes in residual earnings, rather than on standard deviations of residual earnings or absolute changes in residual earnings.

There are two ways in which each of these methodological issues can be empirically meaningful. First, the methodological choices may change the time

Table 4.4: Mean volatility of residual earnings by sector among job stayers

	1975–2005	1975–1981	1983–1989	1991–1997	1999–2005	Change over time	p-value of test of increase over time
Occupation							
Managers	**0.033**	0.017	0.035	0.040	0.031	*0.014*	*0.001*
Professionals	**0.039**	0.012	0.017	0.039	0.069	*0.057*	*0.000*
Technicians	**0.020**	0.016	0.018	0.027	0.014	*-0.002*	*0.761*
Sales	**0.064**	0.036	0.183	0.074	0.034	*-0.002*	*0.545*
Office and admin.	**0.031**	0.014	0.020	0.044	0.040	*0.026*	*0.009*
Prod./craft/repair	**0.033**	0.028	0.029	0.025	0.050	*0.022*	*0.002*
Operators/laborers	**0.035**	0.031	0.040	0.037	0.032	*0.000*	*0.477*
Services	**0.037**	0.020	0.035	0.035	0.047	*0.027*	*0.014*
p-value of equal means test	***0.000***	*0.000*	*0.000*	*0.012*	*0.000*		
Industry							
Ag./extr./constr.	**0.049**	0.012	0.035	0.032	0.084	*0.071*	*0.016*
Manufacturing	**0.034**	0.021	0.022	0.025	0.056	*0.035*	*0.011*
Trans./comm./util.	**0.035**	0.027	0.042	0.036	0.035	*0.007*	*0.105*
Wholesale/retail	**0.028**	0.025	0.027	0.024	0.035	*0.010*	*0.042*
FIRE/prof. svc.	**0.045**	0.025	0.036	0.064	0.041	*0.016*	*0.007*
Public administration	**0.037**	0.020	0.025	0.045	0.050	*0.030*	*0.032*
p-value of equal means test	***0.000***	*0.001*	*0.003*	*0.000*	*0.185*		
Observations	**5057**	1164	1285	1364	1244		

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Job stayers are workers who remained with the same employer over the same nine-year window used to compute volatility of residual earnings. Sample is male heads of household in the PSID; see chapter 3 for details. Column labeled *Change over time* gives change in mean volatility among the specified group from 1975–1981 to 1999–2005. Column labeled *p-value of test of increase over time* gives p-values from t tests in which the alternative hypothesis is that mean volatility for the specified group was higher in 1999–2005 than in 1975–1981. Rows labeled *p-value of equal means test* give p-values from F tests of equality of means among subcategories.

series of economy-wide earnings volatility, affecting the extent to which volatility increased or even whether it increased. Second, these issues may affect comparisons of earnings volatility across workers. This is important both because a number of researchers are interested in such comparisons (Gottschalk and Moffitt, 1994; Stevens, 2001; Drewianka, 2010; Strain, 2011; Ziliak, Hardy, and Bollinger, 2011) and because such comparisons are the basis for the my decomposition approach in chapter 5 to studying why economy-wide earnings volatility increased.

Therefore, I show the effect that each of my methodological recommendations has both on the change in economy-wide earnings volatility over time and on three comparisons of earnings volatility across workers. The cross-worker results use regressions in which the dependent variable is the log of earnings volatility, and I use three correlates that have been discussed in previous research: a worker's age or potential experience, job change, and the unemployment rate.[3] (I use log earnings volatility rather than the level of earnings volatility because some of the changes I consider alter the scale of the volatility measure.)

A final subsection takes up a secondary issue raised by one of my recommendations: if life cycle earnings profiles are to include heterogeneous earnings trends, for how many periods do workers need to be observed in order to obtain sufficiently accurate estimates of these trends?

[3] In chapter 5, I include a quadratic in the unemployment rate in all earnings volatility regressions, to account for the fact that large deviations from typical unemployment rates, whether high or low, may be associated with abnormal earnings changes. However, for simplicity and for comparison with previous literature, in this chapter I do not include the square of the unemployment rate.

4.3.1 Volatility of residual earnings and total earnings

As I argued in chapter 2, when risk is the motivation for studying earnings volatility, volatility of residual earnings is more conceptually sound than volatility of total earnings, because earnings changes explained by a life cycle earnings profile are likely to be anticipated by workers. Here I compare my preferred measure of the volatility of residual earnings to a similar measure of the volatility of total earnings. I measure the volatility of a worker's total earnings in a given year as the variance of the worker's log earnings observations within the nine-year interval centered at that year. This is simply equation (2.2) with log earnings, y, in place of residual log earnings, e.

Figure 4.7 shows how economy-wide volatility of residual and total earnings compare over time. Volatility of total earnings is larger, which is sensible because volatility of total earnings includes life cycle earnings changes, while volatility of residual earnings excludes such changes. However, the basic pattern of each type of volatility over time is similar, and the distance between the two does not vary much over time. For the full sample period, both types of volatility increased roughly the same amount.

The first panel of Table 4.5 shows how some cross-worker comparisons of earnings volatility differ depending on whether volatility of residual earnings or volatility of total earnings is measured. When volatility of total earnings is measured instead of volatility of residual earnings, earnings volatility decreases more with potential experience, because earnings variability due to life cycle dynamics is larger for young workers, who are on a steep portion of the earnings-experience profile, than for old workers, who are on a flatter portion of the profile. The difference in earnings volatility between job changers and job stayers is slightly larger when volatility of total earnings is measured, and

Figure 4.7: Economy-wide volatility of residual earnings and total earnings

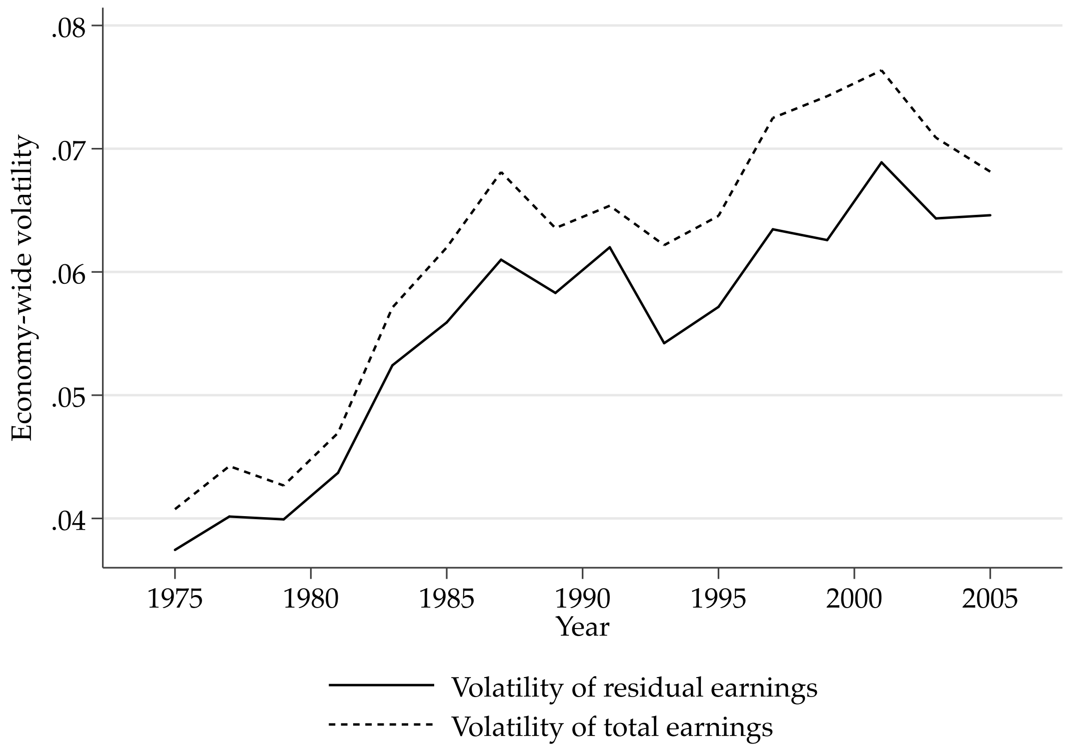

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). A worker's volatility of total earnings is computed using log earnings rather than residual log earnings. Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

this change is statistically significant at the 5 percent level. Finally, the association between earnings volatility and the unemployment rate appears much smaller when volatility of total earnings is measured.

4.3.2 Aggregate shocks

I argued in chapter 2 that life cycle earnings profiles should not absorb aggregate shocks, which are unlikely to be anticipated by workers. This means it

Table 4.5: Effects of some methodological choices on cross-worker comparisons of earnings volatility

	Dependent variable: log of earnings volatility		
	Coefficient on potential experience	Coefficient on indicator of job change	Coefficient on unemployment rate in state
Volatility of residual earnings	−0.023	1.325	0.063
	(0.002)	(0.054)	(0.016)
Volatility of total earnings	−0.035	1.406	0.032
	(0.002)	(0.054)	(0.017)
p-value of test of equal coef.	*0.000*	*0.048*	*0.018*
Profiles: Smoothly varying coef.	−0.023	1.325	0.063
	(0.002)	(0.054)	(0.016)
Profiles: Year-specific coef.	−0.017	1.201	0.033
	(0.002)	(0.055)	(0.017)
p-value of test of equal coef.	*0.001*	*0.000*	*0.013*
Profiles: Worker fixed effects	−0.026	1.392	0.036
	(0.002)	(0.054)	(0.017)
Profiles: Cross-sections	−0.024	1.366	0.031
	(0.002)	(0.053)	(0.016)
p-value of test of equal coef.	*0.101*	*0.181*	*0.480*
Profiles: Worker-specific trends	−0.023	1.325	0.063
	(0.002)	(0.054)	(0.016)
Profiles: Common trends	−0.026	1.392	0.036
	(0.002)	(0.054)	(0.017)
p-value of test of equal coef.	*0.061*	*0.052*	*0.017*
Volatility: Variance of residuals	−0.023	1.325	0.063
	(0.002)	(0.054)	(0.016)
Volatility: Std. dev. of residuals	−0.011	0.662	0.031
	(0.001)	(0.027)	(0.008)
p-value of test of equal coef.	*0.000*	*0.000*	*0.000*

Notes: Within panels of the table, each column represents a separate regression. All regressions include year indicators. Robust standard errors in parentheses. Preferred measure of a worker's earnings volatility in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Each panel modifies one aspect of this measurement process as indicated. Sample is male heads of household in the PSID; see chapter 3 for details.

is better to force the parameters of the profiles to change smoothly over time, rather than estimating them separately for each year.

Figure 4.8 shows the effect of this change on economy-wide volatility of residual earnings over time. The solid line plots economy-wide volatility of residual earnings by year after estimating my preferred life cycle earnings specification with smoothly varying coefficients, and the dotted line plots volatility after estimating profiles with the same regressors but year-specific coefficients. The two series are almost identical, suggesting this issue will not affect conclusions about the extent to which economy-wide earnings volatility has increased over time.

The second panel of Table 4.5 shows that although the trend in economy-wide earnings volatility is not affected by this issue, comparisons of earnings volatility across workers are. When the life cycle profiles are estimated with year-specific coefficients, volatility of residual earnings declines more gradually as a worker's career advances, and the difference in volatility of residual earnings between job changers and job stayers does not appear as large. Both of these differences are statistically significant. Most strikingly, when the life cycle profiles absorb aggregate shocks, the apparent association between volatility of residual earnings and state-specific unemployment rates is only half as large.

4.3.3 Heterogeneous earnings trends

My preferred method of measuring earnings volatility involves a life cycle earnings profile that includes worker fixed effects and worker-specific earnings trends. The natural way to investigate the effects of these heterogeneous trends is to compare volatility of residual earnings based on profiles with heterogeneous trends to volatility based on profiles that include worker fixed effects

Figure 4.8: Effect of absorbing aggregate shocks in life cycle profiles on economy-wide volatility of residual earnings

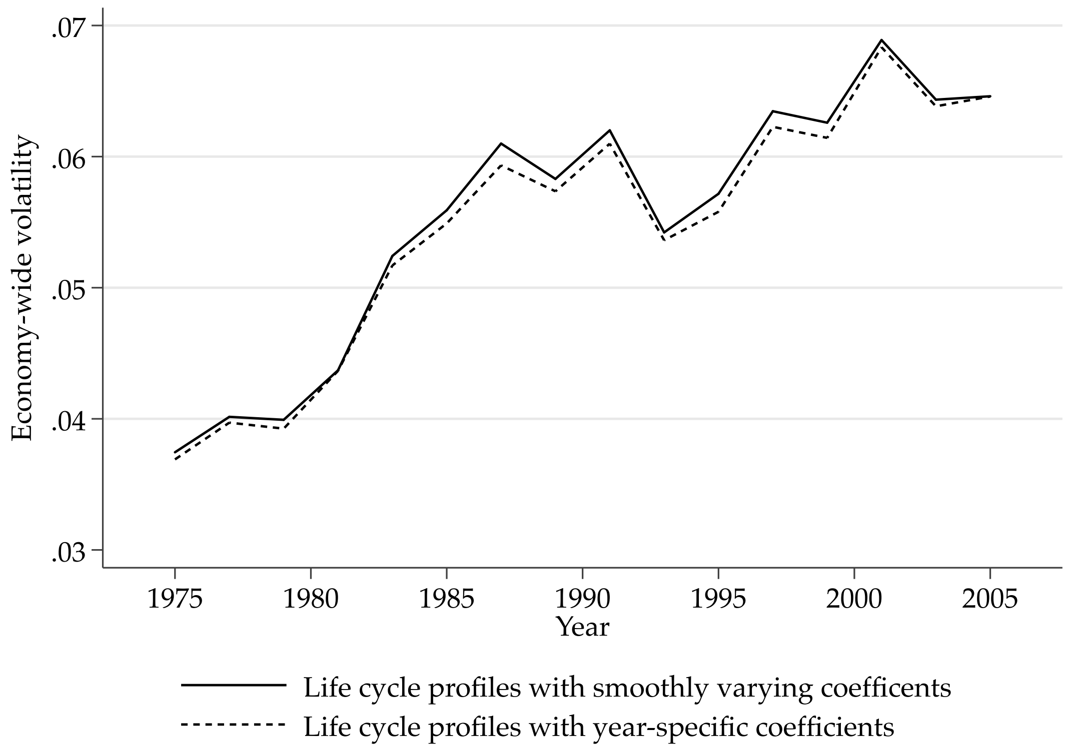

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, which modifies equation (2.1) as specified in the legend. Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

but not worker-specific trends. However, previous studies of volatility of residual earnings do not even include worker fixed effects in the life cycle earnings profiles. So there are two steps from the existing literature to my preferred specification, and each may affect the measurement of earnings volatility. First, adding worker fixed effects may change coefficients in the potential experience polynomial, and therefore change the earnings residuals. Second, further adding heterogeneous earnings trends will change the series of log earnings residuals obtained for each worker.

Figure 4.9: Effect of life cycle earnings profiles estimated on cross-sections on economy-wide volatility of residual earnings

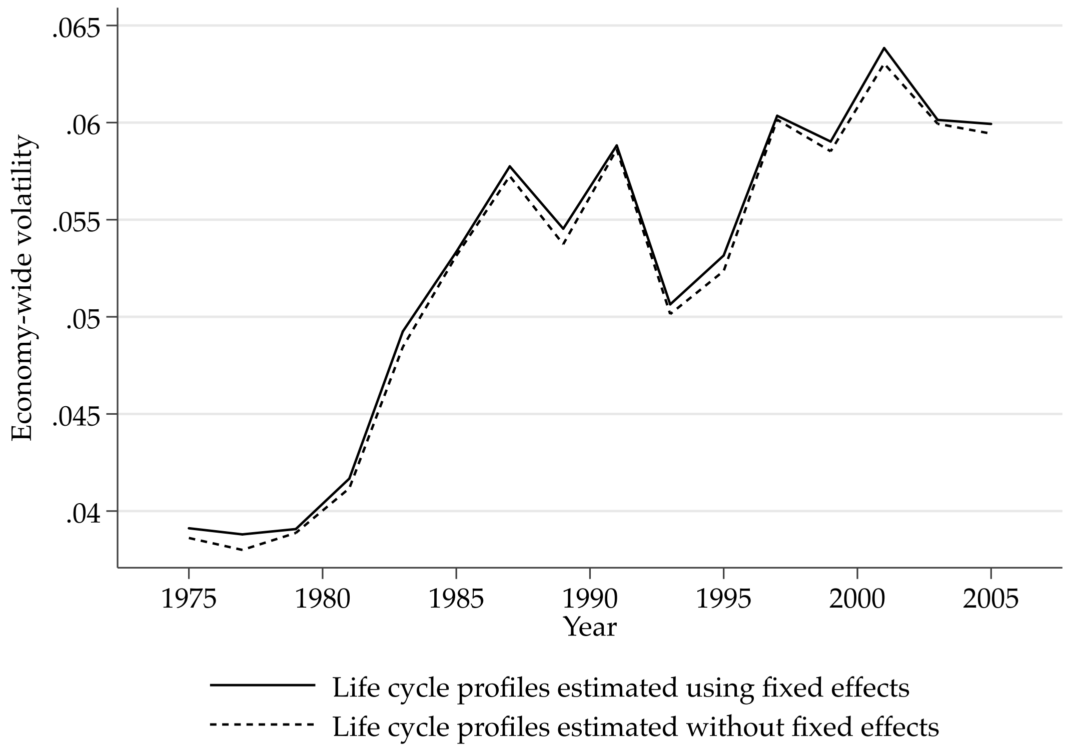

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, which modifies equation (2.1) as specified in the legend. Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

Figure 4.9 shows economy-wide volatility of residual earnings over time depending on whether worker fixed effects are included in the life cycle profiles. The two series are nearly identical. The third panel of Table 4.5 shows that comparisons of volatility of residual earnings across workers are also not affected by this issue: the changes to some regression coefficients of interest are small and statistically insignificant.

The inclusion of heterogeneous earnings trends is empirically more important than the inclusion of worker fixed effects alone. Figure 4.10 shows that economy-wide volatility of residual earnings grew by 73 percent between 1975 and 2005 when worker-specific trends are included in the life cycle profiles, but 53 percent when the heterogeneous trends are omitted. Excluding the heterogeneous trends also causes the associations of volatility of residual earnings with both potential experience and job change to be larger, as shown in the fourth panel of Table 4.5, but the differences are small and not statistically significant at the 5 percent level. Omitting the trends has a larger effect on the relationship between volatility of residual earnings and the unemployment rate, causing the observed association to be cut almost in half.

4.3.4 Functional form choice

I provided a justification in chapter 2 for earnings volatility measures based on variances or squared changes instead of standard deviations or absolute changes. To assess the empirical importance of this issue, I compare my preferred measure, the *variance* of a worker's residual log earnings observations over a nine-year period, to a different measure, the *standard deviation* of the same residuals.

Figure 4.11 shows that the functional form choice is empirically very important. By the variance measure, economy-wide volatility of residual earnings increased 73 percent between 1975 and 2005. By the standard deviation measure, economy-wide volatility of residual earnings increased only 29 percent over the same period.

Table 4.5 shows that the functional form also has a large effect on comparisons of volatility across workers. When the standard deviation is used instead of the variance, the associations of volatility of residual earnings with potential experience, job change, and the unemployment rate all decrease by about half,

Figure 4.10: Effect of omitting heterogeneous earnings trends from life cycle profiles on economy-wide volatility of residual earnings

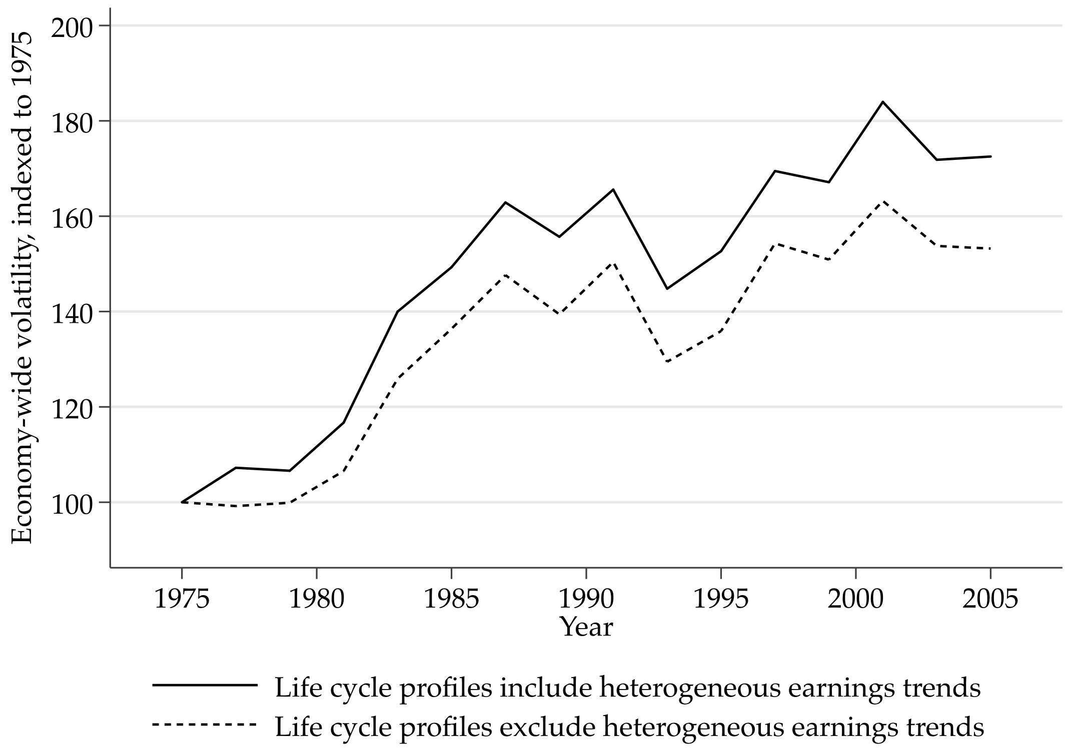

Notes: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, which modifies equation (2.1) as specified in the legend. Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

and each difference is statistically significant at the 1 percent level.

4.3.5 Panel length and heterogeneous earnings trends

The issue of heterogeneous trends in the life cycle earnings profiles raises the question of how many earnings observations are needed in order to obtain a good estimate of a worker's trend. The problem is that processes other than the heterogeneous trends can also create steadily rising or falling earnings over

Figure 4.11: Effect of functional form choice on economy-wide volatility of residual earnings

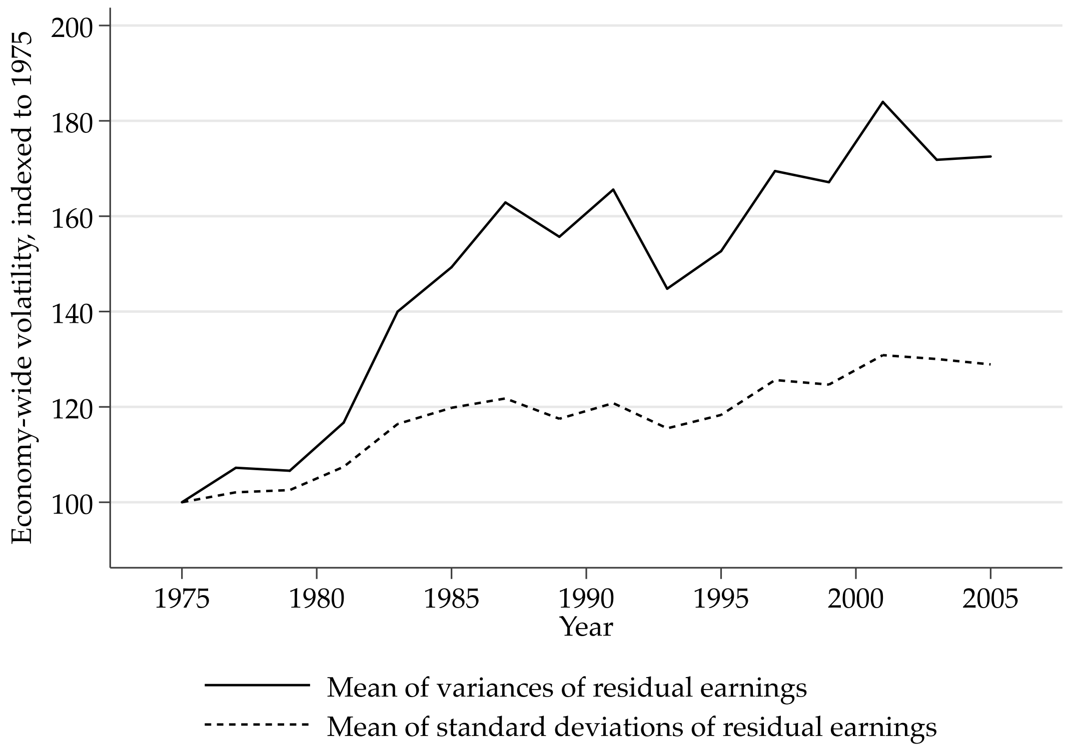

Notes: A worker's volatility of residual earnings in a given year is the variance or standard deviation, as indicated, of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

a short period. For example, the fading of a positive shock in an AR(1) process would appear as a negative earnings trend. If a worker is followed for only a few years, it is impossible to distinguish this from the worker's life cycle earnings trend. The fading of the shock would be absorbed by a bad estimate of the worker's earnings trend, and the worker's estimated volatility would be too low.

My preferred measure of earnings volatility, WA9, is ill-suited to studying this issue, because by definition it requires a worker to be observed at least five times (every two years over a nine-year period). For this subsection, I adopt a shorter measure from the same class: WA3 is the variance of the worker's residual log earnings during a window containing three time periods, instead of nine time periods as in my preferred measure. This shorter window will contain two biannual observations, and I will assign the volatility to the latter of these observations. Using the discussion of this measure in chapter 2, one way to express the WA3 measure is

$$V_{it}^{WA3} = \frac{(e_{it} - e_{i,t-2})^2}{2}. \tag{4.4}$$

To determine how long a worker should be followed for the purpose of measuring earnings volatility, I compute quantities of interest, such as economy-wide earnings volatility, for panels of different lengths. However, as the panel length requirement is shortened, the composition of the sample may change. For example, workers who move geographically may have high earnings volatility and be more likely to drop out of the survey after only a few observations.

To exclude this type of composition shift, I begin with a single sample of workers observed for ten consecutive surveys. (If the worker was observed beyond ten surveys, I keep only the first ten observations.) I compute life cycle earnings profiles, including heterogeneous earnings trends, for this sample, then compute the statistics of interest, such as economy-wide earnings volatility. Then, I drop the final observation for each worker, leaving panels of nine observations. I repeat estimation of the life cycle profiles and the statistics of interest.

I proceed to drop observations—alternating between dropping the last and first remaining observations—and repeat the estimations, stopping after the panels are shortened to three observations. (For panels of just two observations, the life cycle profiles that include heterogeneous trends will perfectly fit the earnings observations, so that measured volatility is zero.) This process simulates the effects of shortening the minimum panel requirement while leaving the sample unchanged. The initial restriction to panels of ten consecutive observations leaves only 706 workers. To check the robustness of my findings to other initial restrictions, and to see whether a larger sample changes the conclusions, I repeat the entire process with progressively shorter initial panel restrictions, from nine to four.

Table 4.6 shows how the heterogeneous earnings trends behave as the panels are shortened. Each row marks an initial panel length restriction, and the columns mark progressively shorter panels from that initial sample. The cells plot the mean difference between the trend computed from the restricted panel and the trend computed from the unrestricted panel. More specifically, let $\hat{\delta}_i^P$ denote the estimated earnings trend for worker i for the initial panel length requirement P, and let $\hat{\delta}_i^R$ be the estimated earnings trend when the panel is shortened to length R. Then the cells plot $\frac{1}{N}\sum_{i=1}^{N}\left|\hat{\delta}_i^R - \hat{\delta}_i^P\right|$.

First consider the initial panel length requirement of ten observations ($P = 10$). When these panels are artificially shortened to nine observations ($R = 9$), the average magnitude of the resulting changes in estimated earnings trends is just 0.006, so that changes in residual log earnings would be misstated by 0.6 percent on average. This error increases steadily through $R = 5$, after which the size of the error accelerates until it is 4.4 percent when the panels are shortened to three observations. The findings are similar if the initial panel length requirement is lowered ($P = 9, 8, \ldots, 4$).

Table 4.7 shows how economy-wide volatility of residual earnings, by my preferred measure, changes as panel length requirements are reduced. At an initial requirement of ten observations, economy-wide earnings volatility is

Table 4.6: Mean absolute change in heterogeneous earnings trends as panels are shortened

	Restricted panel lengths							
Initial panel length	10	9	8	7	6	5	4	3
10	0.000	0.006	0.009	0.014	0.019	0.024	0.033	0.044
9		0.000	0.007	0.012	0.016	0.023	0.031	0.046
8			0.000	0.009	0.015	0.022	0.031	0.051
7				0.000	0.012	0.020	0.032	0.050
6					0.000	0.016	0.030	0.049
5						0.000	0.025	0.046
4							0.000	0.043

Note: Life cycle earnings profiles, including worker-specific earnings trends, are specified in equation (2.1). Table shows mean absolute change in these trends as panels are shortened from the initial length indicated in the first column. Sample is male heads of household in the PSID; see chapter 3 for details.

0.040. As the panels in this sample are artificially shortened, the estimate of economy-wide earnings volatility is surprisingly stable, falling to 0.037 at $R = 5$ and 0.036 at $R = 4$ before dropping to 0.031 at $R = 3$. For lower initial panel length requirements, the findings are similar: the estimate of economy-wide volatility is fairly stable through $R = 5$ and experiences a large drop when moving from $R = 4$ to $R = 3$.

Finally, Table 4.8 shows how one comparison of volatility of residual earnings across workers is affected by the panel length requirement. The statistic of interest in this table is the coefficient from a regression of the log of earnings volatility on an indicator of job change, so the cells in the table are estimates of the proportional difference in volatility of residual earnings between job changers and job stayers. At an initial panel length requirement of ten observations, the difference is 1.202. The difference remains quite stable—between 1.114 and 1.207—as the panels are artificially shortened to a length of $R = 4$, but then drops to 0.964 when the panels are shortened to three observations. The large change between $R = 4$ and $R = 3$ is seen for some, but not all, of the other initial panel length requirements, although the estimates are noisy.

Taken as a whole, these exercises suggest workers should be followed for at least four time periods in order to obtain a "good enough" estimate of their heterogeneous earnings trend. It is important to remember that in my PSID sample, observations are two years apart, because the PSID switched to biannual interviews after 1997. I chose to drop interviews for all previous even-numbered years rather than worry about the issues created by this change. Therefore, workers with four consecutive observations were interviewed six years apart, which is longer than workers can be followed in the SIPP or CPS, two data sources commonly used in the earnings volatility literature.

Table 4.7: Economy-wide volatility of residual earnings as panels are shortened

	Restricted panel lengths							
Initial panel length	10	9	8	7	6	5	4	3
10	0.040	0.040	0.039	0.039	0.038	0.037	0.036	0.031
9		0.040	0.040	0.039	0.039	0.037	0.036	0.031
8			0.047	0.046	0.045	0.043	0.040	0.032
7				0.046	0.045	0.042	0.040	0.032
6					0.054	0.052	0.049	0.041
5						0.052	0.052	0.041
4							0.063	0.053
3								0.054

Note: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details.

Table 4.8: Difference in log volatility of residual earnings between job changers and job stayers as panels are shortened

	Restricted panel lengths							
Initial panel length	10	9	8	7	6	5	4	3
10	1.202	1.114	1.207	1.189	1.182	1.187	1.177	0.964
	(0.096)	(0.099)	(0.108)	(0.115)	(0.129)	(0.138)	(0.164)	(0.203)
9		1.130	1.193	1.235	1.234	1.256	1.312	1.040
		(0.094)	(0.095)	(0.106)	(0.112)	(0.129)	(0.133)	(0.188)
8			1.217	1.193	1.121	1.200	1.224	1.138
			(0.086)	(0.088)	(0.104)	(0.106)	(0.122)	(0.147)
7				1.252	1.173	1.200	1.138	1.168
				(0.079)	(0.084)	(0.097)	(0.109)	(0.142)
6					1.174	1.122	1.107	1.086
					(0.078)	(0.085)	(0.104)	(0.128)
5						1.105	1.094	0.945
						(0.078)	(0.090)	(0.119)
4							1.114	1.211
							(0.082)	(0.096)
3								1.148
								(0.086)

Note: A worker's volatility of residual earnings in a given year is the variance of the worker's earnings residuals during a nine-year window centered on that year; see equation (2.2). The earnings residuals are residuals from a life cycle earnings profile, specified in equation (2.1). Economy-wide volatility in a given year is the mean of workers' volatility in that year. Sample is male heads of household in the PSID; see chapter 3 for details. Robust standard errors in parentheses.

CHAPTER 5

ACCOUNTING FOR INCREASING EARNINGS VOLATILITY

This chapter examines the reasons for rising economy-wide volatility of residual earnings among men in the U.S., using my preferred measure of volatility. Section 5.1 reviews previous research on this topic. In section 5.2, I develop decomposition methods to measure the contributions of various mechanisms to the change in economy-wide earnings volatility. I treat the change over the entire period, 1975 to 2005, as the sum of a series of two-year changes, and perform one decomposition for each two-year change. To find the contribution of a given mechanism to the change in economy-wide earnings volatility over the full period, I total that mechanism's contribution to each of the two-year changes. There are two advantages to this approach. First, if a mechanism contributed to the change in economy-wide earnings volatility, I can see whether this contribution was concentrated within a short time period or spread out over the full sample period. Second, I can relax the assumption that the relationship between earnings volatility and the explanatory variables is constant over time.

In section 5.3, I discuss the mechanisms I consider as candidates to explain rising economy-wide earnings volatility. The most important of these is shocks to labor demand. I construct a demand shock index to capture the exposure of a worker's demographic group to national changes in the occupation-industry distribution of hours worked. Using yearly CPS data, I compute the variance of log hours worked within each occupation-industry sector for each nine-year volatility window. To measure a worker's exposure to these demand shocks, I define the demand shock index for a given window as a weighted average of the sector-specific shocks, where the weights are provided by the sector employment distribution for the worker's sex, education category, and state. Intuitively, a demand shock for managers in the finance industry will raise the demand shock index more for college graduates in New York than for college graduates in Arkansas, and more for college graduates in New York than for low-education workers in New York.

Section 5.4 presents my results. I present both unconditional results, using a separate decomposition for each mechanism, and conditional results, with all mechanisms included in a single decomposition. My primary finding is that much of the increase in economy-wide earnings volatility is due to the fact that workers have become exposed over time to larger or more frequent demand shocks in the labor market. This mechanism explains about half of the increase in economy-wide earnings volatility in both unconditional and conditional decompositions. I show that the association between earnings volatility and the demand shock index is statistically significant across a range of specifications, and remains statistically significant when using only about a quarter of the sample to account for concerns about precision that arise from observations that use overlapping nine-year windows. I also show that the role of demand shocks is robust to a number of alternative measurement strategies, such as removing sector-specific demand trends before constructing the shock index and defining a worker's labor market by sector instead of demographic group.

In light of previous literature that has tied rising economy-wide earnings volatility to job and occupation mobility, I explore the contributions of job and occupation change to rising economy-wide earnings volatility in three ways. First, I measure the unconditional explanatory power of job change indicators and find that job change predicts small decreases in economy-wide earnings volatility over my sample period. Second, I assess the explanatory power of firm, occupation, and industry tenure, both unconditionally and conditional on other variables. If job, occupation, or industry change has become more frequent, the distributions of these tenure variables will shift toward zero. Because low-tenure workers have higher earnings volatility, such a shift could explain part of the increase in economy-wide earnings volatility over time. I find that the explanatory power of the tenure variables is very small in all of the decompositions. Finally, I analyze the extent to which job change mediates the relationship between demand shocks and earnings volatility. For example, demand shocks could affect earnings volatility only through job change, or demand shocks could be associated with more frequent job change even as overall job change rates are roughly constant. I find limited and quantitatively insignificant support for these possibilities.

5.1 Previous research

Previous research has found links between earnings volatility and job loss (Stevens, 2001), the unemployment rate (Gottschalk and Moffitt, 1994; Stevens, 2001), firm revenue volatility (Comin, Groshen, and Rabin, 2009), firm employment level volatility (Strain, 2011), and AFQT score (Strain, 2011). I include these factors in my analysis to the extent my data allows.

A few authors have attempted to explain rising economy-wide earnings volatility, either as its own phenomenon or as part of a broader set of facts. However, no study has used worker-level data to account for increasing economy-wide earnings volatility, and no study has tested alternative explanations for the trend. Comin, Groshen, and Rabin (2009), using COMPUSTAT data, find that volatility of firms' average wages rose by 7.5 percentage points between 1970 and 1999 due to rising firm revenue volatility. Using the most comparable PSID sample, economy-wide earnings volatility rose 2.7 percentage points over the same period. This must be viewed as a very rough accounting. The COMPUSTAT data covers only publicly traded firms, which have more volatile revenue than privately held firms (Davis et al., 2007), and 80 to 85 percent of firms do not report wage data (Pistaferri, 2009). Therefore, the workers in the resulting sample may be very unrepresentative of the PSID sample. Moreover, fluctuations in firms' average wages can come from two sources: wage changes for existing workers, and hires and separations that change the average wage.

Violante (2002) studies theoretically the role of accelerating technical change. Workers learn by doing on the job and endogenously separate from jobs in an attempt to be paired with the newest, most productive machines. As the rate of technical change grows, the technical distance between newer and older machines grows, so there are larger wage changes for those who change jobs. Job stayers experience faster wage growth: assuming that learning profiles on the job are concave, then as technical change accelerates, workers joining a new job carry over less skill from the old job, so they start the new job at a lower, steeper part of the learning profile. Therefore, both job stayers and job changers contribute to rising economy-wide earnings volatility. In a calibration exercise, accelerating technical change can explain most of the observed rise in economy-wide earnings volatility.

One implication of this model is that the rate of job change is weakly increasing in the rate of technical change. Increasing job change is also important in Kambourov and Manovskii (2009a), which builds and calibrates a model to explain rising occupational mobility and increasing wage inequality. The authors note that the productivity shocks that drive the primary results in the paper also predict rising economy-wide earnings volatility.

5.2 Decomposing changes earnings volatility

I use decompositions to quantify the effects of various mechanisms on the change over time in economy-wide volatility of residual earnings. Because volatility of residual earnings is the only type of earnings volatility studied in this chapter, I do not use the "residual" modifier below. Also, earnings volatility is consistently measured using my preferred method, window averaging using nine-year intervals, which was detailed in chapter 2. Therefore, I drop the WA9 superscript when denoting earnings volatility. Worker-specific earnings volatility is denoted V_{it}. Economy-wide earnings volatility is denoted V_t and always refers to the mean of workers' earnings volatilities in the specified year.

The quantity to be decomposed is the change over the full sample period in economy-wide earnings volatility, $V_{2005} - V_{1975}$. Each mechanism $m = 1, \ldots, M$ is associated with explanatory variables x^m. In a traditional Oaxaca decomposition of the change in economy-wide earnings volatility between the initial and final years, the contribution of mechanism m to the change is $(\overline{x}^m_{2005} - \overline{x}^m_{1975})\hat{\beta}_m$, where $\hat{\beta}_m$ is estimated from a regression of earnings volatility on x^m and possibly other variables.

There are two disadvantages to this approach in the current setting. First, while this simple decomposition can identify whether some mechanism contributed to the change in economy-wide earnings volatility over time, it cannot say when that contribution was made—for example, abruptly over a short time period or gradually over the full sample period. Second, it forces the relationship between earnings volatility and the explanatory variables to be constant over time.

To address these issues, I treat the total change in economy-wide earnings volatility, $V_{2005} - V_{1975}$, as the sum of a series of two-year changes, beginning with $V_{1977} - V_{1975}$ and ending with $V_{2005} - V_{2003}$, and perform one decomposition for each two-year change. To find the contribution of a mechanism to the change in economy-wide earnings volatility from 1975 to 2005, I total that mechanism's contribution to each of the two-year changes. By splitting the 30-year change into two-year changes, it is possible to identify the timing of contributions to the change in economy-wide earnings volatility, and the assumption of time-invariant coefficients in the underlying regressions can be relaxed by allowing the coefficients to vary across the two-year decompositions.

I begin with OLS regressions of earnings volatility, which is measured at the worker-year level, on explanatory variables and year indicators:

$$V_{it} = \sum_{m=1}^{M} x_{it}^{m} \beta_m(t) + \alpha_t + \varepsilon_{it}. \tag{5.1}$$

In one set of decompositions, the coefficients in the underlying regressions are constrained to be time-invariant:

$$\beta_m(t) = \beta_m. \tag{5.2}$$

In a second set of decompositions, the coefficients are allowed to vary over time. Year-specific coefficients would be too noisy given the limited sample size available in each year, so I allow the coefficients to change over time along a quadratic path:

$$\beta_m(t) = \beta_{0m} + \beta_{1m} t + \beta_{2m} t^2. \tag{5.3}$$

The decompositions of the change in economy-wide earnings volatility between 1975 and 2005 are performed by summing up the two-year decompositions. The two-year changes in the means of the explanatory variables are multiplied by the regression coefficients from the midpoint year of the two-year decomposition. These midpoint year coefficients exist because the coefficients are either assumed to be time-invariant or to vary as a smooth function of time, so they are defined for any year during the sample. Let $\hat{\delta}^{m}_{t,t+2}$ denote the contribution of mechanism m to the change in economy-wide earnings volatility in one of the two-year decompositions, and let $\hat{\Delta}^{m}$ denote the contribution of mechanism m to the change in economy-wide earnings volatility over the full sample period. Then

$$\begin{aligned}\hat{\Delta}^{m} &= \hat{\delta}^{m}_{1975,1977} + \cdots + \hat{\delta}^{m}_{2003,2005} \\ &= (\overline{x}^{m}_{1977} - \overline{x}^{m}_{1975})\hat{\beta}_{m}(1976) + \cdots + (\overline{x}^{m}_{2005} - \overline{x}^{m}_{2003})\hat{\beta}_{m}(2004).\end{aligned} \tag{5.4}$$

To complete a full description of the decompositions, first consider the decompositions of two-year changes. Denote the portion of the change due to changing coefficients associated with mechanism m by $\hat{\delta}^{\beta_m}_{t,t+2}$, and the portion due to changes in the year indicators by $\hat{\delta}^{\alpha}_{t,t+2}$. Then the decomposition of the first two-year change in economy-wide earnings volatility is

$$\begin{aligned}V_{1977} - V_{1975} = &\sum_{m=1}^{M} (\overline{x}^{m}_{1977} - \overline{x}^{m}_{1975})\hat{\beta}_{m}(1976) \\ &+ \sum_{m=1}^{M} \Big[\overline{x}^{m}_{1977} \left(\hat{\beta}_{m}(1977) - \hat{\beta}_{m}(1976)\right) \\ &\qquad + \overline{x}^{m}_{1975} \left(\hat{\beta}_{m}(1976) - \hat{\beta}_{m}(1975)\right) \Big] \\ &+ (\hat{\alpha}_{1977} - \hat{\alpha}_{1975}) \\ = &\sum_{m=1}^{M} \hat{\delta}^{m}_{1975,1977} + \sum_{m=1}^{M} \hat{\delta}^{\beta_m}_{1975,1977} + \hat{\delta}^{\alpha}_{1975,1977}.\end{aligned} \tag{5.5}$$

For the long decomposition, let $\hat{\Delta}^{\beta_m} = \hat{\delta}^{\beta_m}_{1975,1977} + \cdots + \hat{\delta}^{\beta_m}_{2003,2005}$ be the contribution of changing coefficients associated with mechanism m to the change in economy-wide earnings volatility between 1975 and 2005, and let $\hat{\Delta}^{\alpha} = \hat{\delta}^{\alpha}_{1975,1977} + \cdots + \hat{\delta}^{\alpha}_{2003,2005} = \hat{\alpha}_{2005} - \hat{\alpha}_{1975}$ be the contribution of the year indicators. Then the change in economy-wide earnings volatility between 1975 and

2005 can be decomposed into a portion explained by changes in the explanatory variables, an unexplained portion due to changes in the coefficients, and an unexplained portion due to changes in the year indicators. Using the above definitions and the template provided by equation (5.5), the decomposition of the 30-year change is

$$\begin{aligned} V_{2005} - V_{1975} &= (V_{2005} - V_{2003}) + \cdots + (V_{1977} - V_{1975}) \\ &= \sum_{m=1}^{M} \left[\hat{\delta}^{m}_{2003,2005} + \cdots + \hat{\delta}^{m}_{1975,1977}\right] \\ &\quad + \sum_{m=1}^{M} \left[\hat{\delta}^{\beta_m}_{2003,2005} + \cdots + \hat{\delta}^{\beta_m}_{1975,1977}\right] \\ &\quad + \hat{\delta}^{\alpha}_{2003,2005} + \cdots + \hat{\delta}^{\alpha}_{1975,1977} \\ &= \sum_{m=1}^{M} \hat{\Delta}^{m} + \sum_{m=1}^{M} \hat{\Delta}^{\beta_m} + \hat{\Delta}^{\alpha}. \end{aligned} \tag{5.6}$$

The above discussion applies to the case in which all of the mechanisms, $m = 1, \ldots, M$, are included in the underlying volatility regression (5.1). I call decompositions based on this multiple regression *conditional decompositions*, emphasizing the presence of multiple mechanisms. To assess the unconditional explanatory power of individual mechanisms, I also present *unconditional decompositions*, in which only one mechanism is included in the underlying volatility regression. In this case, the applicable versions of equations (5.1), (5.5), and (5.6) do not include the summation $\sum_{m=1}^{M}$.

The usual concern about inconsistency of the OLS estimates of $\beta_m(t)$ in equation (5.1) applies. As a motivating example in the context of earnings volatility, consider measurement error. Suppose the variance of earnings measurement error differs across workers and that the variance is correlated with worker characteristics. Duncan and Hill (1985), Bound et al. (1994), and Pischke (1995) provide evidence that earnings measurement error is correlated with worker characteristics using the PSID survey instrument. Because earnings measurement error contributes to observed earnings volatility, this produces a worker-specific effect in an earnings volatility regression that appears in the error term in equation (5.1) and is correlated with explanatory variables in that model, making OLS estimates of $\beta_m(t)$ biased and inconsistent.

To address this, I use repeated earnings volatility observations to estimate a modification of equation (5.1) that includes worker fixed effects:

$$V_{it} = \sum_{m=1}^{M} x_{it}^{m} \tilde{\beta}_m(t) + \gamma_i + \tilde{\alpha}_t + \tilde{\varepsilon}_{it}. \tag{5.7}$$

The decomposition (5.6) must be altered because the change in economy-wide earnings volatility, $V_{2005} - V_{1975}$, is now affected by entry to and exit from the sample over time through the worker fixed effects. Denote the average estimated worker effect in year t by $\overline{\gamma}_t$. Then the additional term $\overline{\gamma}_{2005} - \overline{\gamma}_{1975} = \hat{\Delta}^{\gamma}$ belongs on the right hand side of equation (5.6) when the worker fixed effects specification is used as the underlying regression.

The fixed effects approach is conceptually more attractive than the OLS approach, but the fixed effects method has some disadvantages in the context of this chapter. First, measures of earnings volatility are already quite noisy, and worker fixed effects further remove useful variation in earnings volatility. Second, in the empirical results below, I discuss the concern that many of the windows over which earnings volatility is measured overlap. To address a possible consequence—artificially small standard errors when earnings volatility is regressed on covariates—I restrict the sample to earnings volatility observations with non-overlapping windows. This leaves too few earnings volatility observations per worker to implement fixed effects earnings volatility regressions. Because of these issues, I give equal weight to decomposition results based on OLS and fixed effects regressions, and fortunately the major results are not sensitive to the estimation method.

Both because of the number of results generated and because the detailed "unexplained portions" of the decompositions have no natural interpretation, I present in the results only the explained changes in economy-wide earnings volatility, $\hat{\Delta}^m$. These terms are compared to the change in economy-wide earnings volatility, $V_{2005} - V_{1975}$, to assess the explanatory power of mechanism m.

5.3 Mechanisms

In this section I discuss the mechanisms to be assessed. Variables associated with earnings volatility, and for which the mean has changed over time, may have some explanatory power for the change in economy-wide earnings volatility. As far as possible, I include each mechanism that has been linked in previous literature to the levels of workers' earnings volatility or changes in economy-wide earnings volatility. The mechanisms and explanatory variables are listed in Table 5.1, and summary statistics for both the full sample and four sub-periods are presented in Table 3.1. Overall, these variables contain a good deal of information on labor demand and on the characteristics of workers and jobs, and less information on labor supply and labor market institutions.

Demand shocks

The U.S. has experienced major changes in the wage and occupation distributions over the past 40 years, with wages and employment falling in medium-skilled routine work and rising in high-skill abstract work and low-skill service work (Autor, Levy, and Murnane, 2003; Autor and Dorn, 2012). The changes did not occur smoothly, though. During the 1970s and 1980s, the changes were monotonic: higher wage percentiles experienced better wage changes. But during the following two decades, wage changes were most favorable among the

Table 5.1: Variables to explain earnings volatility of worker i in year t

Mechanism	Explanatory variable(s)
Demand shocks	Weighted average of occupation-industry variances of log hours worked during $[t-4, t+4]$ from CPS data, where weights are employment distribution specific to demographic group of worker i; see text for details
Education	3 indicators of education of worker i: high school or less; some college; college or more
Experience	Potential experience of worker i in year t
Unemployment rate	Quadratic in mean unemployment rate during $[t-4, t+4]$ in state worker i lived in as of year $t-4$
Marriage	Indicator of worker i being married in year $t-4$
Union membership	Indicator for union membership of worker i in year $t-4$
Employer tenure	Employer tenure of worker i in year $t-4$
Occupation tenure	Occupation tenure of worker i in year $t-4$
Industry tenure	Industry tenure of worker i in year $t-4$
Occupation	8 indicators of occupation of worker i in year $t-4$: managers; professionals; technicians; sales; clerical, administrative; production, craft, repair; operators, laborers; services
Industry	9 indicators of industry of worker i in year $t-4$: agriculture, extraction; construction; manufacturing; transportation, communication, utilities; wholesale trade, retail trade; finance, insurance, real estate; business services; professional services; public administration

highest percentiles, but less favorable near the median than at the bottom of the wage distribution (Acemoglu and Autor, 2011). Occupation categories also experienced business cycle fluctuations differently: Jaimovich and Siu (2012) show that the long-term employment decline in middle-skill occupations occurred largely during recessions, while abstract and service occupations experienced comparatively slight employment decreases during downturns.

Following this literature, I interpret national changes in the occupation-industry employment distribution as reflecting shifts in the demand for labor, rather than supply shifts that might occur because of changes in workers' job preferences. To capture the influence on earnings volatility of these changes in labor demand, I create an index of demand shocks using CPS data on national changes in the occupation-industry employment distribution.[1] The measure of demand shocks within an occupation-industry sector is the variance of log hours worked in the sector over the same nine-year window that is used to compute earnings volatility.

For each demographic group, I compute the initial distribution of hours worked across occupation-industry sectors, using the first five years of the sample. I then define the demand shock index for the demographic group to be the weighted average of the demand shock index for each occupation-industry sector, where the weights are the sector employment mix specific to the demographic group. Workers are matched to the index based on their demographic group at the beginning of the nine-year window. Intuitively, if there is a large negative demand shock to manufacturing, the demand shock index increases more for low-education workers in Ohio than for low-education workers in Oregon, and more for low-education workers in Ohio than for college graduates in Ohio.

More precisely, let j index demographic groups formed by the intersection of two sex categories, three education categories (high school or less, some college, and college or more), and 51 states (including the District of Columbia). Let k index sectors formed by the intersection of eight occupation categories and six industry categories (listed in Table 5.1). Using annual March CPS data for 1971–2009, let $E_{j\tau}$ be hours worked by demographic group j in year τ, let

[1] This approach modifies previous research that has constructed indices of demand shifts from an initial to a final year, using national sectoral employment changes weighted by employment distributions specific to demographic groups. Examples include Katz and Murphy (1992), Blau, Kahn, and Waldfogel (2000), and Bound and Holzer (2000).

$E_{k\tau}$ be hours worked in sector k in year τ, and let $E_{jk\tau}$ be hours worked by demographic group j in sector k in year τ. Let V_{kt} be the variance of $E_{k\tau}$ over the window $[t-4, t+4]$, and let α_{jk} be the average of $E_{jk\tau}/E_{j\tau}$ over the first five years of the sample. The V_{kt} measure sector demand shifts over the window at the national level, and the α_{jk} describe the initial sector employment mixes of demographic groups, with $\sum_k \alpha_{jk} = 1$. Then the demand shock index is defined to be

$$DV_{jt} = \sum_k \alpha_{jk} V_{kt}. \tag{5.8}$$

The demand shock index could be constructed in other ways, reflecting different assumptions about the labor market to which a worker belongs or different measures of demand. Below, I discuss the relative merits of alternative constructions of the demand shock index and show that my results are robust to using these other constructions.

Education

Rising educational attainment is one of the most salient trends over the past 50 years in the U.S. labor market (Acemoglu and Autor, 2011). Table 4.2 shows that mean earnings volatility is quantitatively similar across three education categories—high school or less, some college, and college or more—for the full sample period, and the null hypothesis of equality across groups cannot be rejected. But the conditional relationship between education and earnings volatility may be stronger. For example, workers of different skill levels may experience different labor market conditions, as emphasized by the literature on within-sector skill upgrading (Autor, Katz, and Krueger, 1998). Gottschalk and Moffitt (1994) found some evidence of earnings volatility differences across education categories, especially among job changers.

Experience

If earnings volatility varies with age or experience, then changes in the demographic composition of the labor force would be expected to change the level of economy-wide earnings volatility. Table 4.2 shows that older workers tend to have lower earnings volatility than younger workers, likely because older workers change jobs less often. That, in turn, could be due to an accumulation of specific human capital or to having had longer to find a good worker-firm match. Gottschalk and Moffitt (1994) also found that earnings volatility declines with age.

Unemployment rate

Business cycle fluctuations might be expected to affect earnings volatility because of especially high or low earnings growth, or through especially high or low job loss, although recessions are associated both with more job losses and fewer job-to-job changes (Davis, Faberman, and Haltiwanger, 2012). Recall that the life cycle profile removes "typical" earnings growth trends before earnings volatility is calculated. When labor markets are especially tight, unemployment is low and earnings will tend to grow faster than usual; earnings volatility may be higher than in typical business cycle conditions. Similarly, earnings volatility could be high in times of especially high unemployment rates, when earnings growth is low. To account for this possibly nonlinear effect, I include a quadratic in the mean unemployment rate in the worker's state as of year $t-4$. Stevens (2001) finds a positive relationship between earnings volatility and unemployment. Gottschalk and Moffitt (1994) find a positive relationship only among involuntary job changers, with a negative relationship among job stayers.

Marriage

Married workers may have higher or lower earnings volatility than unmarried workers. For married workers, the spouse's human capital provides insurance against earnings shocks, which reduces the utility cost of earnings volatility. That leads married workers to be more accepting of earnings volatility. But married workers may also be less willing or able to pursue opportunities that require geographic moves; this will tend to reduce earnings volatility. Also, many workers may wait to marry until after they have settled on a career path, in which case earnings volatility will be lower for married than for non-married workers, but with no causal relationship leading from marriage to earnings volatility. Table 4.2 indicates that married men consistently have lower earnings volatility than unmarried men, and marriage rates have been decreasing in the U.S. (Stevenson and Wolfers, 2007). I include an indicator of marriage as of the beginning of the nine-year earnings volatility window.

Union membership

Unionization rates have fallen steadily in the U.S. for over 50 years; union coverage in my PSID sample of men fell from 25 percent in 1975 to 15 percent in 2005, and from 27 percent to 13 percent among private-sector workers. While union contracts may smooth wage growth, rigid wages may in turn raise the probability of layoffs at unionized firms (Medoff, 1979). Therefore the relationship between earnings volatility and union membership is an empirical question. To account for the effects of union membership, I include an indicator of union membership as of the beginning of the nine-year window that is used to compute earnings volatility. Gottschalk and Moffitt (1994) found earnings volatility was higher among non-union workers than among union workers.

Table 4.2 echoes this finding, although the difference is small, not statistically significant, and inconsistent over time.

Occupation and industry

The industrial and occupational composition of the labor force changed significantly over my sample period (Acemoglu and Autor, 2011). Table 4.2 shows that earnings volatility does differ across industry and occupation categories, so these labor market trends could exert compositional effects on earnings volatility apart from the demand shock effects discussed above. I include indicators for eight occupations and nine industries, which are detailed in Table 5.1. The indicators are assigned as of the beginning of the nine-year window that is used to compute earnings volatility.

Tenure

Specific human capital plays an important role in wage determination. Previous research has documented returns to firm-specific (Altonji and Shakotko, 1987), industry-specific (Parent, 2000), and occupation-specific (Kambourov and Manovskii, 2009b) human capital, with a concave profile to the returns. Table 4.2 confirms that earnings volatility tends to decrease as employer, occupation, or industry tenure increases. To account for the effects of more frequent job change or higher occupation or industry mobility on earnings volatility, I include a measure of each of the three types of tenure—employer, industry, and occupation—as of the beginning of the nine-year earnings volatility window. As discussed in the introduction, this is one way to test existing explanations for rising economy-wide earnings volatility.

Other factors

Performance pay is an element of an increasing fraction of U.S. jobs (Lemieux, MacLeod, and Parent, 2009). It is possible that earnings are more variable in performance pay jobs, because output and thus pay depends partly on shocks the worker cannot control. However, my labor earnings measure is wage and salary earnings, which does not include performance pay.

Measurement error undoubtedly accounts for some of the residuals from equation (2.1) that are used to compute earnings volatility. However, this does not directly present a problem for studying the change in economy-wide earnings volatility over time unless the structure of measurement error changes. Duncan and Hill (1985), Bound et al. (1994), and Pischke (1995) have studied measurement error in the PSID, but to my knowledge no studies have examined the change in measurement error over time, whether in the PSID or any other survey. The PSID switched to computer-assisted interviewing in the mid-1990s, so it is conceivable that the magnitude of measurement error has decreased over time in this survey, in which case the rise in economy-wide earnings volatility is even larger than it appears.

5.4 Results

This section presents the main results on why economy-wide earnings volatility has increased among males in the U.S. I begin with unconditional decompositions, in which the explanatory power of each mechanism for the increase in economy-wide earnings volatility is measured separately without conditioning on any other mechanisms. Labor demand shocks are by far the strongest explanation unconditionally. Next, I discuss regression results and find that the effect of demand shocks on earnings volatility is robust to the inclusion of various sets of control variables and remains statistically significant when using a subset of the sample to remove concerns about overlapping observations. Following this, I present conditional decompositions that quantify the explanatory power of each mechanism when controlling for the others. Demand shocks explain about half of the increase in economy-wide earnings volatility in these decompositions. I then explore the extent to which job change mediates the relationship between demand shocks and earnings volatility. Finally, I show that the decomposition results are robust to alternative constructions of the demand shock index.

5.4.1 Unconditional decompositions

Because my primary question is why economy-wide earnings volatility increased, in lieu of presenting unconditional regression results, I perform unconditional decompositions of the increase in economy-wide earnings volatility for each mechanism discussed above. These decompositions answer the following question: Given the unconditional relationship between earnings volatility and mechanism m, how much of the increase in economy-wide earnings volatility over the sample period can be explained by mechanism m? For each decomposition, the covariates used in the regression on which the decomposition is based are the explanatory variables associated with the mechanism, which are detailed above and summarized in Table 5.1, and a full set of year indicators.

Table 5.2 presents the results for decompositions based on OLS regressions, in which each row represents one decomposition using time-invariant coefficients and one decomposition using coefficients allowed to vary along a quadratic path over time (equation (5.3)). Demand shocks have the greatest explanatory power for the increase in economy-wide earnings volatility, accounting for more than 40 percent of the change between 1975 and 2005 in each specification. Given the relationship between earnings volatility and unemployment, the downward trend in the unemployment rate over the sample period predicts a decrease in economy-wide earnings volatility.[2] Changes in marriage account for about 15 percent of the increase in economy-wide earnings volatility.

The explanatory power of the other mechanisms is uniformly small, confirming that pure composition effects, such as changes in the distribution of education or occupation, explain almost none of the change in economy-wide earnings volatility. Recall that previous theories of rising economy-wide earnings volatility predict shifts toward zero in the distribution of firm or occupation tenure. If low-tenure workers have higher earnings volatility, then these shifts could explain increasing economy-wide volatility. But in the unconditional decompositions, the tenure variables more often predict decreases than increases in economy-wide earnings volatility.

To more directly test whether job change is responsible for rising economy-wide earnings volatility, I also measure the total explanatory power of indicators of voluntary and involuntary job change at any point during the nine-year earnings volatility window.[3] Table 4.2 shows that job changers experience much higher earnings volatility than job stayers, but the results in Table

[2]It may be that the natural rate of unemployment has declined over time, so that the effect on earnings volatility of 7 percent unemployment today is similar to the effect of, say, 9 percent unemployment in the 1970s. In unreported results, I associate the business cycle mechanism with a detrended unemployment rate instead of the level of the unemployment rate. When this is done, the predicted negative effect of the unemployment rate on the change in economy-wide earnings volatility becomes much smaller.

[3]Job changes due to establishment closing or layoff are assumed to be involuntary. All other job changes are labeled voluntary. Workers with both voluntary and involuntary job changes during the window are assigned to the involuntary job change group.

Table 5.2: Unconditional decompositions of change in economy-wide earnings volatility from 1975 to 2005, based on OLS regressions

	Time-invariant coef.		Smoothly varying coef.	
Mechanism	$\hat{\Delta}^m$	Percent explained	$\hat{\Delta}^m$	Percent explained
Demand shock index	0.0115	42.5	0.0124	45.5
Education	0.0006	2.2	−0.0011	−4.0
Experience	0.0003	1.1	−0.0004	−1.6
Unemployment rate	−0.0046	−16.9	−0.0091	−33.5
Married	0.0041	15.0	0.0035	13.0
Union membership	0.0002	0.7	0.0002	0.8
Occupation	−0.0001	−0.5	−0.0013	−4.7
Industry	−0.0000	−0.2	−0.0005	−1.7
Employer tenure	0.0006	2.1	−0.0000	−0.0
Occupation tenure	−0.0013	−4.8	−0.0013	−4.7
Industry tenure	−0.0008	−2.9	−0.0008	−2.9
Job change	−0.0003	−1.0	−0.0003	−1.0

Notes: Each row represents one decomposition. $\hat{\Delta}^m$ is the change in economy-wide earnings volatility attributable to changes in the explanatory variables associated with the mechanism in the first column, computed according to equation (5.4). Decompositions with time-invariant coefficients are based on the regression (5.1) without other mechanisms and using the coefficient specification (5.2). Decompositions with smoothly varying coefficients are based on the regression (5.1) without other mechansisms and using the coefficient specification (5.3). Percent explained is the ratio of $\hat{\Delta}^m$ to the change in economy-wide earnings volatility between 1975 and 2005, which is 0.0272. See text for a description of the explanatory variables associated with each mechanism.

5.2 show that job change cannot explain any of the increase in economy-wide earnings volatility, because the incidence of job change was roughly constant over time in this sample (see the summary statistics in Table 3.1).

Table 5.3 presents the results of unconditional decompositions based on regressions with worker fixed effects. Demand shocks about 60 percent of the increase in economy-wide earnings volatility. The tenure measures do tend to make positive contributions, but never explain more than 15 percent of the change. Changes in the occupation and industry distributions make negative contributions to the change in economy-wide earnings volatility.

Table 5.3: Unconditional decompositions of change in economy-wide earnings volatility from 1975 to 2005, based on fixed effects regressions

	Time-invariant coef.		Smoothly varying coef.	
Mechanism	$\hat{\Delta}^m$	Percent explained	$\hat{\Delta}^m$	Percent explained
Demand shock index	0.0157	57.7	0.0178	65.6
Unemployment rate	0.0010	3.8	−0.0017	−6.4
Married	0.0028	10.5	0.0019	6.9
Union membership	0.0014	5.1	0.0013	4.8
Occupation	−0.0023	−8.3	−0.0033	−12.3
Industry	−0.0027	−9.8	−0.0040	−14.6
Employer tenure	−0.0003	−1.0	0.0041	15.0
Occupation tenure	0.0013	4.9	0.0019	7.1
Industry tenure	0.0018	6.6	0.0015	5.5
Job change	−0.0002	−0.8	−0.0002	−0.9

Notes: Each row represents one decomposition. $\hat{\Delta}^m$ is the change in economy-wide earnings volatility attributable to changes in the explanatory variables associated with the mechanism in the first column, computed according to equation (5.4). Decompositions with time-invariant coefficients are based on the regression (5.7) without other mechanisms and using the coefficient specification (5.2). Decompositions with smoothly varying coefficients are based on the regression (5.7) without other mechansisms and using the coefficient specification (5.3). Percent explained is the ratio of $\hat{\Delta}^m$ to the change in economy-wide earnings volatility between 1975 and 2005, which is 0.0272. See text for a description of the explanatory variables associated with each mechanism.

5.4.2 Regressions

Table 5.4 assesses the robustness of the relationship between earnings volatility and demand shocks as various controls are added. Columns 1 and 2 show the unconditional relationship without and with year indicators, respectively. The specification in column 2 corresponds to the demand shock unconditional decomposition with time-invariant coefficients in Table 5.2. Column 3 adds the worker's education category, potential experience, and a quadratic in the unemployment rate, and column 4 includes all of the mechanisms listed in Table 5.1. In all specifications, the association between earnings volatility and labor demand shocks is statistically significant, although the estimate becomes noisier as controls are added. Table 5.5 presents the same specifications for models with worker fixed effects. The coefficients on the demand shock index are very similar for the OLS and fixed effects results. The economic significance of these results are the subject of the conditional decompositions discussed below.

Because the demand shock index varies at the education-state-year level and because disturbances at the level of education-state demographic groups are likely to be serially correlated, standard errors are clustered on education-state categories. A further concern is serial correlation of worker-specific unobservables, so I also cluster on workers, using the method of Cameron, Gelbach, and Miller (2011) to implement multiway clustering.[4]

One concern about the regression results is the fact that for many observations, the nine-year earnings volatility windows overlap. Because the dependent variable and some of the explanatory variables, such as the demand shock index and the unemployment rate, are measured over the nine-year window, the results in Table 5.4 may overstate the precision of the point estimates, even after clustering on workers. To explore the contribution of this false precision to my results in Table 5.4, I perform the same regressions using only earnings volatility observations for windows centered on 1975, 1985, 1995, and 2005. This prevents any two observations from using overlapping nine-year windows.

This conservative approach underestimates the precision of the point estimates by throwing out too much data. First, it excludes workers whose only earnings volatility observations are in years other than the four selected years, such as a worker in the earnings sample only from 1987 to 1997, with volatility

[4]These two group levels—demographic groups and workers—are not nested because workers may change demographic groups over time by moving to another state.

Table 5.4: Earnings volatility regressions without worker fixed effects

	Dependent variable: earnings volatility			
	(1)	(2)	(3)	(4)
Demand shock index	4.596***	4.723**	7.279***	6.598**
	(1.046)	(1.922)	(2.807)	(3.252)
Educ: Some college			−0.006*	0.001
			(0.004)	(0.005)
Educ: College or more			−0.014**	0.001
			(0.006)	(0.006)
Experience			−0.000*	0.000
			(0.000)	(0.000)
Mean u.e. rate			0.010	0.010
			(0.007)	(0.013)
Mean u.e. rate sq.			−0.001	−0.001
			(0.000)	(0.001)
Married				−0.019***
				(0.006)
Union				0.003
				(0.017)
Employer tenure				−0.001**
				(0.001)
Occ. tenure				−0.001
				(0.000)
Ind. tenure				0.001
				(0.001)
Year indicators		yes	yes	yes
Occ. and ind. indicators				yes
R^2	0.006	0.010	0.013	0.031
Observations	9479	9479	9479	9479

Notes: See text for descriptions of the explanatory variables. Robust standard errors, in parentheses, are clustered on education-state groups and on workers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5.5: Earnings volatility regressions with worker fixed effects

	Dependent variable: earnings volatility			
	(1)	(2)	(3)	(4)
Demand shock index	4.036***	6.417**	6.373**	6.586**
	(1.102)	(2.921)	(2.939)	(2.805)
Mean u.e. rate			0.000	0.000
			(0.009)	(0.009)
Mean u.e. rate sq.			−0.000	−0.000
			(0.001)	(0.001)
Married				−0.015***
				(0.005)
Union				−0.015
				(0.017)
Employer tenure				0.000
				(0.000)
Occ. tenure				−0.000
				(0.001)
Ind. tenure				0.002***
				(0.001)
Year indicators		yes	yes	yes
Occ. and ind. indicators				yes
R^2	0.004	0.007	0.007	0.021
Observations	9210	9210	9210	9210

Notes: See text for descriptions of the explanatory variables. Robust standard errors, in parentheses, are clustered on education-state groups and on workers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

observations in 1991 and 1993. Second, it excludes information from workers whose initial or final earnings volatility observations do not fall in one of the four selected years. For example, consider a worker in the earnings sample from 1977 to 1995, with earnings volatility observations every two years from 1981 to 1991. Only the 1985 earnings volatility observation, which captures the relationship between earnings variation and covariates during 1981–1989, is included in this limited sample. Information on the relationship between earnings variation and covariates for 1977–1979 and 1991–1995 is excluded.

The results for the non-overlapping subsample are shown in Table 5.6. The effect of demand shocks on earnings volatility remains statistically significant in each specification. There are not enough lengthy worker panels to perform a similar exercise with worker fixed effects.

5.4.3 Conditional decompositions

Table 5.7 shows the results of two conditional decompositions of the change in economy-wide earnings volatility over the sample period, using the full set of mechanisms. One decomposition uses the time-invariant coefficients reported in Table 5.4, column 4, and the other decomposition allows coefficients to vary over time along a quadratic path (equation (5.3)). Each of the two decomposition approaches provides an answer to the following question: Given the relationship between earnings volatility and mechanism m after holding the other mechanisms constant, how much of the increase in economy-wide earnings volatility over the sample period can be explained by mechanism m?

Demand shocks remain a strong explanation for the increase in economy-wide earnings volatility after adding the full set of controls. Increasingly variable labor demand accounts for 45 to 60 percent of the rise in economy-wide

Table 5.6: Earnings volatility regressions without worker fixed effects, non-overlapping sample

	Dependent variable: earnings volatility			
	(1)	(2)	(3)	(4)
Demand shock index	4.144***	5.169*	9.612**	8.884**
	(1.279)	(2.842)	(3.967)	(4.105)
Educ: Some college			−0.008	0.000
			(0.006)	(0.006)
Educ: College or more			−0.023***	−0.005
			(0.007)	(0.008)
Experience			−0.000	0.001
			(0.000)	(0.000)
Mean u.e. rate			0.005	0.006
			(0.011)	(0.011)
Mean u.e. rate sq.			−0.000	−0.000
			(0.001)	(0.001)
Married				−0.004
				(0.006)
Union				0.002
				(0.006)
Employer tenure				−0.001**
				(0.000)
Occ. tenure				−0.002**
				(0.001)
Ind. tenure				0.001
				(0.001)
Year indicators		yes	yes	yes
Occ. and ind. indicators				yes
R^2	0.008	0.013	0.019	0.043
Observations	2362	2362	2362	2362

Notes: Only observations from year 1975, 1985, 1995, and 2005 are included, so that the nine-year earnings volatility windows associated with each observation do not overlap for any observation in this sample. See text for descriptions of the explanatory variables. Robust standard errors, in parentheses, are clustered on education-state groups and on workers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5.7: Conditional decompositions of change in economy-wide earnings volatility from 1975 to 2005, based on OLS regressions

	Time-invariant coef.		Smoothly varying coef.	
Mechanism	$\hat{\Delta}^m$	Percent explained	$\hat{\Delta}^m$	Percent explained
Demand shock index	0.0161	59.3	0.0125	45.8
Education	0.0002	0.7	−0.0030	−11.2
Experience	−0.0003	−0.9	0.0000	0.0
Unemployment rate	−0.0044	−16.2	−0.0061	−22.3
Married	0.0035	12.9	0.0031	11.4
Union membership	−0.0004	−1.4	−0.0003	−0.9
Occupation	−0.0014	−5.1	−0.0025	−9.1
Industry	−0.0004	−1.5	−0.0008	−3.0
Employer tenure	0.0005	2.0	−0.0007	−2.4
Occupation tenure	−0.0007	−2.4	−0.0008	−3.0
Industry tenure	0.0004	1.6	0.0004	1.6
Total	0.0133	49.0	0.0019	6.8

Notes: $\hat{\Delta}^m$ is the change in economy-wide earnings volatility attributable to changes in the explanatory variables associated with the mechanism in the first column, computed according to equation (5.4). The decomposition with time-invariant coefficients is based on the regression (5.1) with the coefficient specification (5.2). The decomposition with smoothly varying coefficients is based on the regression (5.1) with the coefficient specification (5.3). Percent explained is the ratio of $\hat{\Delta}^m$ to the change in economy-wide earnings volatility between 1975 and 2005, which is 0.0272. See text for a description of the explanatory variables associated with each mechanism.

earnings volatility, depending on the specification. The decline in the unemployment rate over time predicts a modest decrease in economy-wide earnings volatility.[5] Change in marriage can account for about 12 percent of the increase in economy-wide earnings volatility.

As in the unconditional decompositions, the explanatory power of the other

[5] As discussed in footnote 2 above, it may be appropriate to detrend the unemployment rate to account for an apparent secular decline in the natural rate of unemployment. The key issue is whether earnings volatility should vary with the level of the unemployment rate or with the position of the unemployment rate relative to its natural level. In unreported results, I repeat the regressions and conditional decompositions using the detrended unemployment rate instead of the level of the unemployment rate. The predicted negative effect of the unemployment rate on the change in economy-wide earnings volatility becomes much smaller, and the explanatory power of the demand shock index changes very little.

Table 5.8: Conditional decompositions of change in economy-wide earnings volatility from 1975 to 2005, based on fixed effects regressions

	Time-invariant coef.		Smoothly varying coef.	
Mechanism	$\hat{\Delta}^m$	Percent explained	$\hat{\Delta}^m$	Percent explained
Demand shock index	0.0161	59.2	0.0106	38.8
Unemployment rate	0.0011	4.1	−0.0022	−7.9
Married	0.0028	10.4	0.0014	5.3
Union membership	0.0021	7.6	0.0014	5.2
Occupation	−0.0019	−7.0	−0.0037	−13.6
Industry	−0.0023	−8.3	−0.0034	−12.4
Employer tenure	−0.0000	−0.1	0.0006	2.1
Occupation tenure	−0.0005	−1.7	−0.0002	−0.6
Industry tenure	0.0019	7.0	0.0019	6.9
Total	0.0194	71.2	0.0065	24.0

Notes: $\hat{\Delta}^m$ is the change in economy-wide earnings volatility attributable to changes in the explanatory variables associated with the mechanism in the first column, computed according to equation (5.4). The decomposition with time-invariant coefficients is based on the regression (5.7) with the coefficient specification (5.2). The decomposition with smoothly varying coefficients is based on the regression (5.7) with the coefficient specification (5.3). Percent explained is the ratio of $\hat{\Delta}^m$ to the change in economy-wide earnings volatility between 1975 and 2005, which is 0.0272. See text for a description of the explanatory variables associated with each mechanism.

mechanisms tends to be quite small. The estimated contribution of job change or occupation or industry mobility, as measured by the effect of the tenure variables, is close to zero in both specifications.

Table 5.8 shows the results of conditional decompositions based on worker fixed effect models. The explanatory power of labor demand shocks remains high, explaining 40 to 60 percent of the increase in economy-wide earnings volatility. Changes in the occupation and industry distributions predict decreases in economy-wide earnings volatility, and the explanatory power of the tenure variables, taken together, remains small.

Figure 5.1 provides a way of viewing the contribution of labor demand shocks to the change in economy-wide earnings volatility over the sample period. The solid line is the cumulative change in economy-wide earnings volatility since 1975. The dotted line is the cumulative change in economy-wide earnings volatility that is attributed to demand shocks by the decomposition. That is, denoting the position on the x-axis by τ and using the notation defined in section 5.2 above, the solid line plots $V_\tau - V_{1975}$, and the dotted line shows $\hat{\delta}^m_{1975,1977} + \cdots + \hat{\delta}^m_{\tau-2,\tau}$. The figure uses the conditional regression results in Table 5.4, column (4). Labor demand shocks explained all of the increase in economy-wide earnings volatility as late as the mid-1990s, but since then economy-wide earnings volatility has risen while the demand shock index has fallen.

Figure 5.2 shows the same exercises for each of the other mechanisms. The decomposition results in Table 5.7—which show little explanatory power for most of the mechanisms—do not hide any large, offsetting contributions during the sample period. Most of the negative contribution of the unemployment rate is concentrated in the last decade of the sample.

Figures 5.3 and 5.4 repeat these graphical exercises for decompositions based on fixed effects regressions, specifically the results in Table 5.5, column (4). Again, labor demand shocks explained all of the increase in economy-wide earnings volatility until the last decade of the sample period, when the index of labor demand shocks declined. The other mechanisms have little explanatory power throughout the sample period.

5.4.4 Role of job change

Job change does not appear to explain much of the change in economy-wide earnings volatility, as discussed above in the section on unconditional decompositions. But, it may be that demand shocks affect earnings volatility only

Figure 5.1: Contribution of labor demand shocks to change in economy-wide earnings volatility in decompositions based on OLS regressions

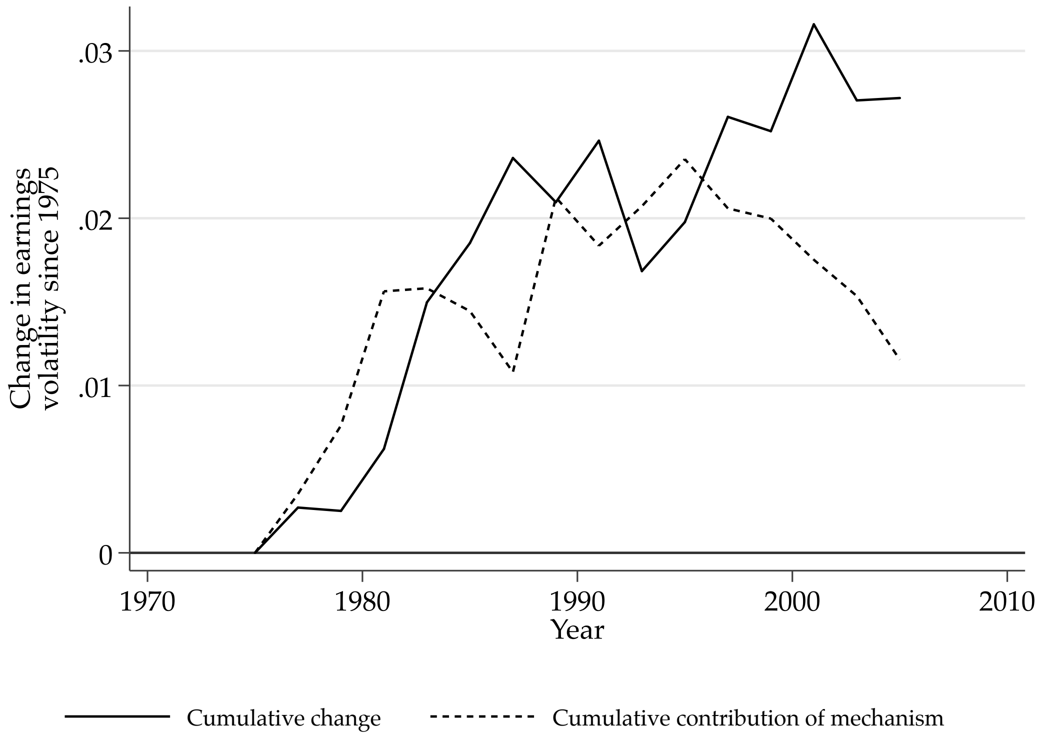

Notes: Solid line is cumulative change in economy-wide earnings volatility since 1975. Dotted line is cumulative contribution of labor demand shocks to this change, calculated by summing contributions of demand shocks in each decomposition of two-year change in economy-wide earnings volatility from 1975 to indicated year; see equation (5.4). Based on results in Table 5.4, column (4).

upon job change, in which case the consequences of job change could explain the rise in economy-wide earnings volatility even if rates of job change did not increase. Also, it may be that demand shocks have tended to increase the rate of job change even though other factors have worked in the opposite direction, keeping total job change rates fairly constant. Either possibility would be support for existing theories of rising economy-wide earnings volatility, as discussed in the introduction. In this section I explore the extent to which job change mediates the relationship between demand shocks and earnings

Figure 5.2: Contribution of various mechanisms to change in economy-wide earnings volatility in decompositions based on OLS regressions

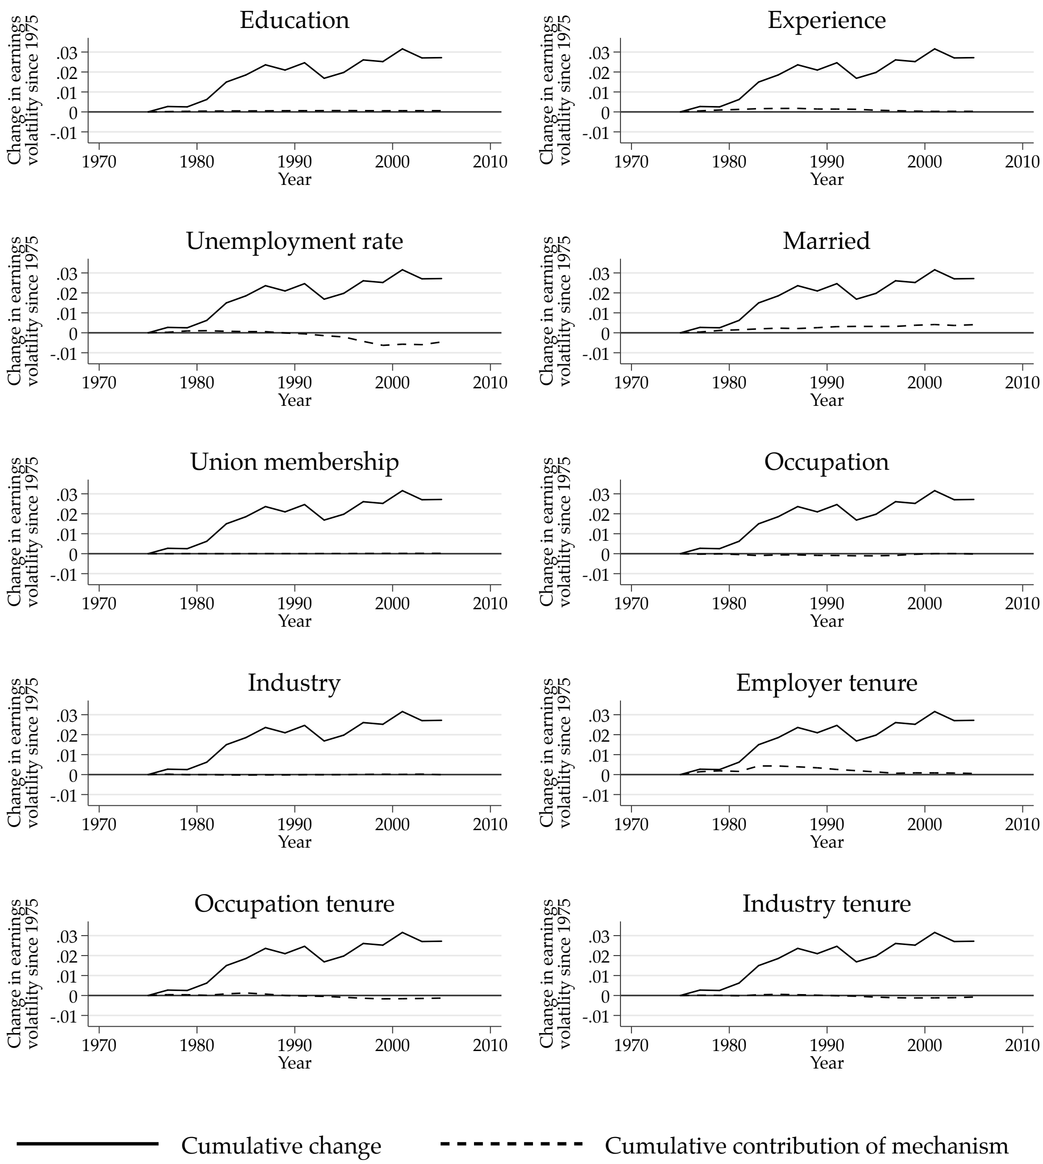

Notes: Solid line is cumulative change in economy-wide earnings volatility since 1975. Dotted line is cumulative contribution of indicated mechanism to this change, calculated by summing contributions of mechanism in each decomposition of two-year change in economy-wide earnings volatility from 1975 to indicated year; see equation (5.4). Based on results in Table 5.4, column (4).

Figure 5.3: Contribution of labor demand shocks to change in economy-wide earnings volatility in decompositions based on fixed effects regressions

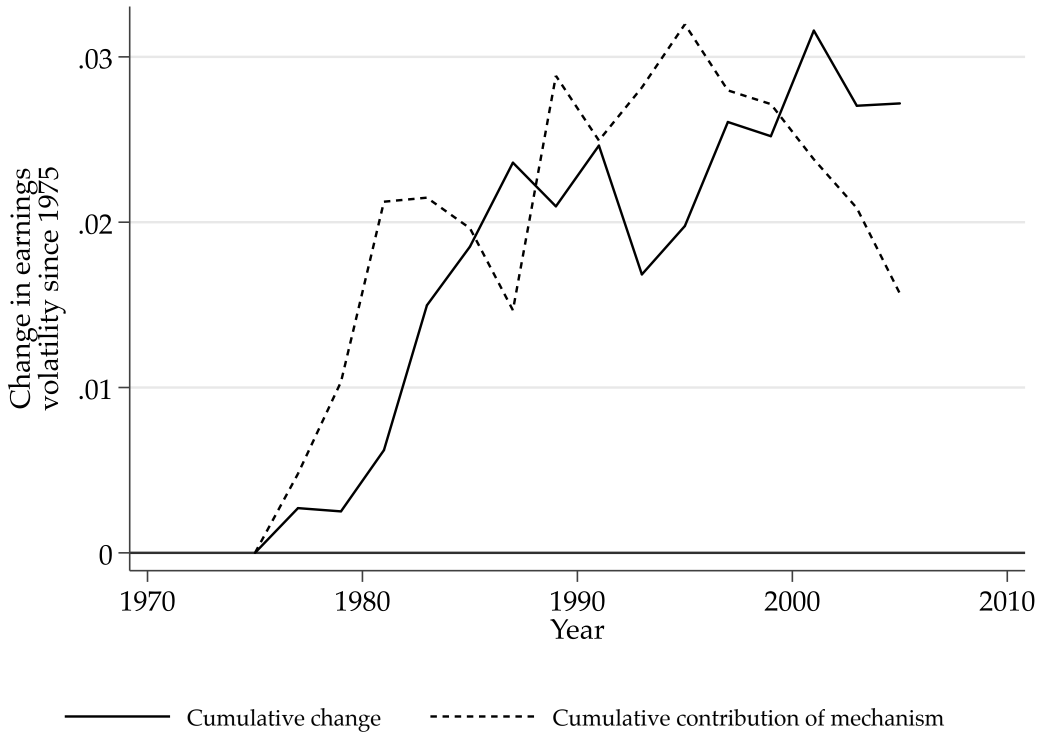

Notes: Solid line is cumulative change in economy-wide earnings volatility since 1975. Dotted line is cumulative contribution of labor demand shocks to this change, calculated by summing contributions of demand shocks in each decomposition of two-year change in economy-wide earnings volatility from 1975 to indicated year; see equation (5.4). Based on results in Table 5.5, column (4).

volatility.

First, I ask whether the effect of demand shocks on earnings volatility is larger among job changers than among job stayers. The first two columns of Tables 5.9 and 5.10 present the results for regressions without and with worker fixed effects, respectively. Though the coefficients are somewhat imprecisely estimated compared to the earlier regression results, the overall story seems to be that demand shocks are positively associated with earnings volatility among job stayers, and the association is about 50 percent larger among job changers.

Figure 5.4: Contribution of various mechanisms to change in economy-wide earnings volatility in decompositions based on fixed effects regressions

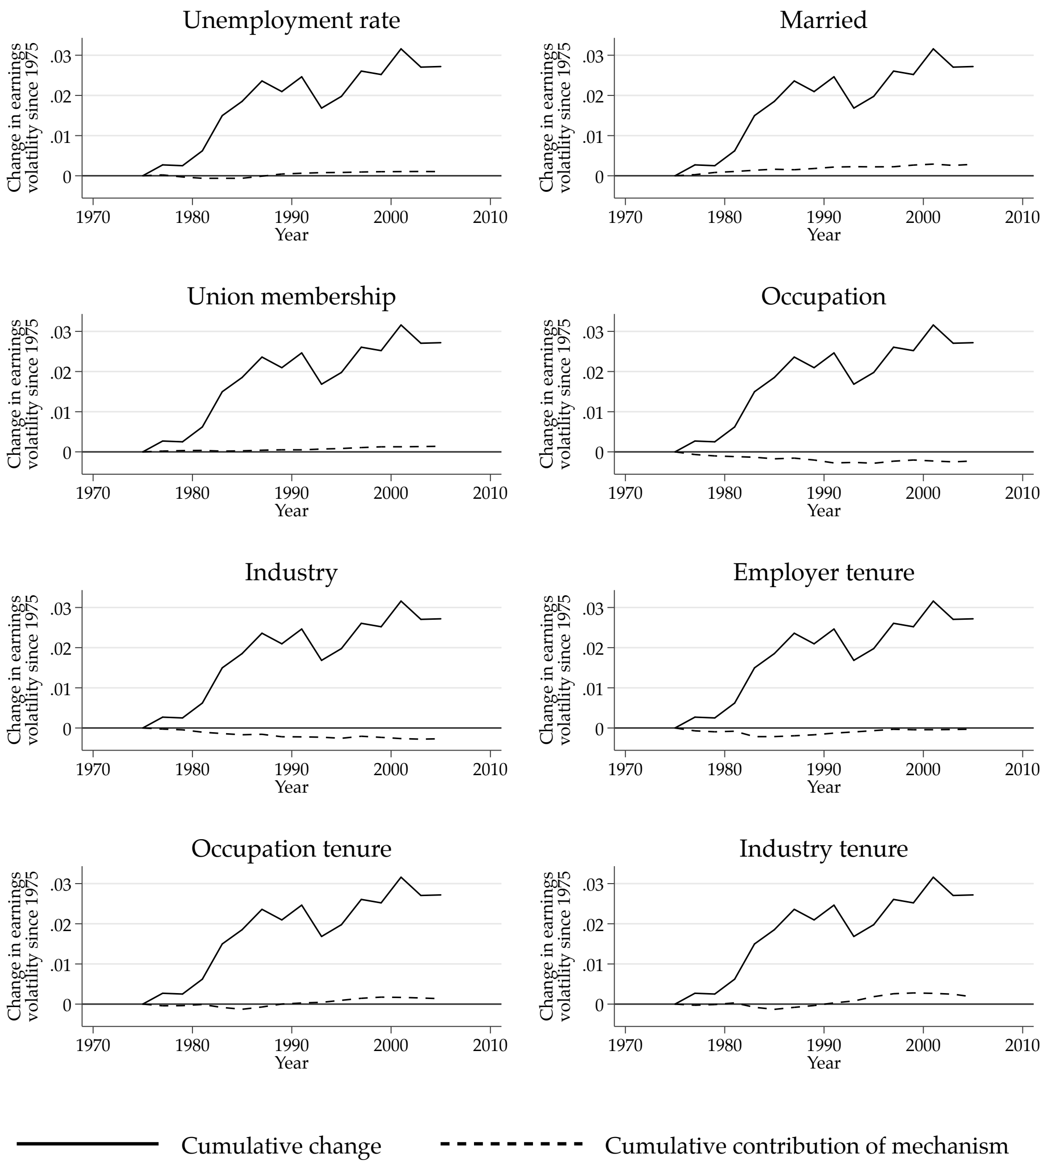

Notes: Solid line is cumulative change in economy-wide earnings volatility since 1975. Dotted line is cumulative contribution of indicated mechanism to this change, calculated by summing contributions of mechanism in each decomposition of two-year change in economy-wide earnings volatility from 1975 to indicated year; see equation (5.4). Based on results in Table 5.5, column (4).

Table 5.9: Role of job change in the relationship between labor demand shocks and earnings volatility, regressions without worker fixed effects

	Dep. var.: earnings volatility		Dep. var.: job change	
	(1)	(2)	(3)	(4)
Demand shock index	2.979	5.663	19.352	3.743
	(2.782)	(4.123)	(94.009)	(64.350)
Demand shocks × job change	1.835	1.598		
	(2.786)	(3.756)		
Job change	0.039***	0.040**		
	(0.009)	(0.018)		
Year indicators	yes	yes	yes	yes
Controls		yes		yes
R^2	0.055	0.067	0.010	0.225
Observations	9479	9479	9479	9479

Notes: Job change variable is an indicator of any job change during the nine-year window used to measure earnings volatility. Controls are education categories, potential experience, unemployment rate, marriage, union membership, employer tenure, occupation tenure, industry tenure, and occupation and industry indicators. See Table 5.1 for details. Robust standard errors, in parentheses, are clustered on education-state groups and on workers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

It is not the case that demand shocks affect earnings volatility entirely through job change.

The final two columns of Tables 5.9 and 5.10 address the extent to which demand shocks affect the probability of job change over the nine-year volatility window, using linear probability models. On average 47 percent of workers change jobs during a nine-year window. In OLS specifications, the estimates indicate that a one standard deviation increase (0.0018) in the demand shock index raises the probability of job change by 0.7 to 3.5 percentage points. In fixed effects specifications, the same increase in the demand shock index decreases the probability of job change by 0.1 to 0.8 percentage points. None of these estimates are statistically significantly different from zero. The demand shock

Table 5.10: Role of job change in the relationship between labor demand shocks and earnings volatility, regressions with worker fixed effects

	Dep. var.: earnings volatility		Dep. var.: job change	
	(1)	(2)	(3)	(4)
Demand shock index	5.392*	5.079	−4.612	−0.674
	(3.198)	(3.098)	(8.702)	(8.604)
Demand shocks × job change	2.411	3.083		
	(2.293)	(2.255)		
Job change	0.029***	0.024**		
	(0.011)	(0.011)		
Year indicators	yes	yes	yes	yes
Controls		yes		yes
R^2	0.036	0.047	0.078	0.117
Observations	9210	9210	9210	9210

Notes: Job change variable is an indicator of any job change during the nine-year window used to measure earnings volatility. Controls are unemployment rate, marriage, union membership, employer tenure, occupation tenure, industry tenure, and occupation and industry indicators. See Table 5.1 for details. Robust standard errors, in parentheses, are clustered on education-state groups and on workers. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

index increased by 0.0024 between 1975 and 2005. Assuming that this increased the probability of job change by 4.6 percentage points ($= 0.0024 \times 19.352$, the largest estimated coefficient), and multiplying by the unconditional difference in mean earnings volatility between job changers and job stayers (0.047), the higher rates of job change caused by demand shocks predict an increase of 0.0022 in economy-wide earnings volatility, less than one tenth of the actual increase of 0.0272.

5.4.5 Robustness checks

In this section I explore the robustness of my primary finding, that labor demand shocks explain about half the increase in economy-wide earnings volatility, to alternative constructions of the demand shock index. I consider three choices made in defining this index, and show results for each of the eight possible combinations of these choices.

First, consider the definition of the worker's labor market. In the results discussed above, labor markets are defined by demographic groups, and these groups are affected differently by national changes in the occupation-industry employment distribution. Among the advantages of this labor market classification is that it recognizes that the sector demand shifts may create local labor market spillovers. For example, service workers may be more affected by a broad negative demand shock to the manufacturing industry in areas where manufacturing employment is relatively high, both because manufacturing workers may shift to lower-skill occupations (Autor and Dorn, 2009) and because the demand of manufacturing workers for services decreases.

Nevertheless, it may be that a better indicator of the worker's labor market is the occupation-industry sector in which the worker is employed, in which case the demand shock index is simply the variability of demand in the worker's sector. A disadvantage of this classification is that workers change occupation and industry somewhat frequently, and presumably at least partially on a forward-looking basis.[6] This creates endogeneity concerns when using occupation-industry sector demand shifts directly to explain earnings volatility, because the sector choice influences both earnings volatility (for example, from searching for a good job within the new sector) and the demand shock index (which would be sector-specific). Another disadvantage of this approach is that occupation and industry classifications likely contain more measurement

[6]For the 1980s, Molloy, Smith, and Wozniak (2011) find that about 3 percent of men move to another state each year. Kambourov and Manovskii (2008) estimate that the annual incidence among males of one-digit occupation change and one-digit industry change is roughly 15 percent and 12 percent, respectively; the incidence of occupation *or* industry change would therefore be greater than 15 percent.

error than education and state categories.

A second choice made in the construction of the demand shock index is to use log hours as the measure of sector labor demands. An alternative is to use the sector's share of all hours worked within the year. This relative demand measure would be unchanged by shifts in the population growth rate, but would also mask any macroeconomic shocks that affect all sectors equally.

The third choice I consider is whether to detrend the measure of sector demand before computing the variability of demand in that sector. In the results above, the sector demands are not detrended before constructing the demand shock index. An alternative is to remove a linear trend in demand specific to each sector, so that long-term shifts in the sector employment mix do not contribute to observed labor demand shocks. For example, there has been a decades-long movement of employment toward skilled occupations dominated by highly educated workers (Acemoglu and Autor, 2011). Arguably, this shift is eventually built into workers' expectations of the future, and should be reflected in life cycle earnings profiles—through worker fixed effects and perhaps heterogeneous trends—rather than in volatility of residual earnings. On the other hand, such trends are rarely permanent (on the recent slowdown in relative demand shifts favoring highly skilled workers, see Beaudry, Green, and Sand (2013)), and detrending the sector demands may overstate the extent to which workers expected the trends. Overall, I do not see a clear case for strongly preferring one of these choices—detrending the sector demands, or not—over the other.

Table 5.11 shows that the association between earnings volatility and labor demand shocks in OLS estimates is quite robust to using other constructions of the demand shock index. For each of the alternative constructions of the index, the table shows the coefficient on the demand shock index in both unconditional and conditional regressions. (The estimates in the first row repeat those in columns (2) and (4) of Table 5.4.) The labor demand shock index has the expected sign in all specifications and is statistically significant in most specifications. Table 5.12 repeats the robustness checks for models with worker fixed effects. Again, the labor demand shock index is statistically significant in most specifications. Table 5.13 shows OLS estimates using the non-overlapping subsample described above. Some of the estimates are much noisier than the full sample results in Table 5.11, but all are of the hypothesized sign and most are statistically significant.

Tables 5.14 and 5.15 show robustness of the decomposition results to the various constructions of the demand shock index. For each construction of the index, the tables show the percent of the increase in economy-wide earnings volatility that is explained by demand shocks in the conditional decompositions which include all explanatory variables in Table 5.1. Table 5.14 shows results in which the underlying regressions are estimated using OLS, and Table 5.15 shows results in which the underlying regressions are estimated using worker fixed effects. The percent change in economy-wide earnings volatility explained by demand shocks varies between 20 percent and 80 percent.

Table 5.11: Coefficient on demand shock index in earnings volatility regressions for various constructions of demand shock index, OLS regressions

Demand shock index construction				
Labor market	Demand measure	Detrended	Regression with year indicators	Regression with year indicators and controls
demographic	log hours	no	4.723** (1.922)	6.598** (3.252)
demographic	log hours	yes	4.531 (3.130)	3.586 (4.575)
demographic	emp. share	no	0.985*** (0.334)	3.582*** (1.142)
demographic	emp. share	yes	7.564*** (2.505)	15.729** (7.966)
sector	log hours	no	4.082** (1.653)	4.529** (2.009)
sector	log hours	yes	4.190* (2.494)	4.532* (2.721)
sector	emp. share	no	1.921 (1.467)	2.748 (2.300)
sector	emp. share	yes	8.071* (4.213)	8.810* (5.335)

Notes: Each result is the estimated coefficient on the demand shock index in an earnings volatility regression, where the construction of the index is described by the first three columns: whether education-state demographic groups or occupation-industry sectors are taken to be the appropriate labor markets, whether log hours or economy-wide employment share is taken to measure a sector's demand, and whether the sector's demand is detrended before the demand shock index is measured. See text for further details. Controls in the rightmost column are education categories, potential experience, unemployment rate, marriage, union membership, employer tenure, occupation tenure, industry tenure, and occupation and industry indicators. See Table 5.1 for details. Robust standard errors, in parentheses, are clustered on workers and on labor markets: education-state groups when demographic groups represent labor markets, and occupation-industry groups when sectors represent labor markets.

Table 5.12: Coefficient on demand shock index in earnings volatility regressions for various constructions of demand shock index, fixed effects regressions

Demand shock index construction				Regression with year indicators and controls
Labor market	Demand measure	Detrended	Regression with year indicators	
demographic	log hours	no	6.417**	6.586**
			(2.921)	(2.805)
demographic	log hours	yes	6.247	6.659
			(4.244)	(4.131)
demographic	emp. share	no	1.465	1.643*
			(0.957)	(0.927)
demographic	emp. share	yes	6.413*	6.618*
			(3.801)	(3.567)
sector	log hours	no	4.883**	4.853**
			(1.958)	(1.957)
sector	log hours	yes	5.469**	5.422**
			(2.636)	(2.629)
sector	emp. share	no	2.293	2.206
			(1.457)	(1.395)
sector	emp. share	yes	6.406**	6.183**
			(3.103)	(2.786)

Notes: Each result is the estimated coefficient on the demand shock index in an earnings volatility regression, where the construction of the index is described by the first three columns: whether education-state demographic groups or occupation-industry sectors are taken to be the appropriate labor markets, whether log hours or economy-wide employment share is taken to measure a sector's demand, and whether the sector's demand is detrended before the demand shock index is measured. See text for further details. Controls in the rightmost column are unemployment rate, marriage, union membership, employer tenure, occupation tenure, industry tenure, and occupation and industry indicators. See Table 5.1 for details. Robust standard errors, in parentheses, are clustered on workers and on labor markets: education-state groups when demographic groups represent labor markets, and occupation-industry groups when sectors represent labor markets.

Table 5.13: Coefficient on demand shock index in earnings volatility regressions for various constructions of demand shock index, OLS regressions, non-overlapping sample

Demand shock index construction				Regression with
Labor market	Demand measure	Detrended	Regression with year indicators	year indicators and controls
demographic	log hours	no	5.169*	8.884**
			(2.842)	(4.105)
demographic	log hours	yes	7.492	7.047
			(6.377)	(7.105)
demographic	emp. share	no	3.192***	4.670***
			(1.230)	(1.786)
demographic	emp. share	yes	8.453	15.432
			(14.019)	(18.574)
sector	log hours	no	5.006***	5.679***
			(1.700)	(1.838)
sector	log hours	yes	5.022**	5.525**
			(2.281)	(2.332)
sector	emp. share	no	3.130*	3.728**
			(1.815)	(1.855)
sector	emp. share	yes	8.386***	8.285***
			(2.442)	(2.272)

Notes: Only observations from year 1975, 1985, 1995, and 2005 are included, so that the nine-year earnings volatility windows associated with each observation do not overlap for any observation in this sample. Each result is the estimated coefficient on the demand shock index in an earnings volatility regression, where the construction of the index is described by the first three columns: whether education-state demographic groups or occupation-industry sectors are taken to be the appropriate labor markets, whether log hours or economy-wide employment share is taken to measure a sector's demand, and whether the sector's demand is detrended before the demand shock index is measured. See text for further details. Controls in the rightmost column are education categories, potential experience, unemployment rate, marriage, union membership, employer tenure, occupation tenure, industry tenure, and occupation and industry indicators. See Table 5.1 for details. Robust standard errors, in parentheses, are clustered on workers and on labor markets: education-state groups when demographic groups represent labor markets, and occupation-industry groups when sectors represent labor markets.

Table 5.14: Percent of change in economy-wide earnings volatility explained by demand shocks, for various constructions of demand shock index, in decompositions based on OLS regressions

Demand shock index construction			Time-invariant coefficients	Smoothly varying coefficients
Labor market	Demand measure	Detrended		
demographic	log hours	no	59.3	45.8
demographic	log hours	yes	38.8	19.6
demographic	emp. share	no	44.3	60.8
demographic	emp. share	yes	81.4	66.8
sector	log hours	no	55.3	53.2
sector	log hours	yes	61.0	68.3
sector	emp. share	no	36.2	44.6
sector	emp. share	yes	43.1	39.4

Notes: Each result is the percentage of the change in economy-wide earnings volatility between 1975 and 2005 that is explained by an index of labor demand shocks, where the construction of the index is described by the first three columns: whether education-state demographic groups or occupation-industry sectors are taken to be the appropriate labor markets, whether log hours or economy-wide employment share is taken to measure a sector's demand, and whether the sector's demand is detrended before the demand shock index is measured. See text for further details. The results come from conditional decompositions similar to those in Table 5.7, and include all the controls listed in that table.

Table 5.15: Percent of change in economy-wide earnings volatility explained by demand shocks, for various constructions of demand shock index, in decompositions based on fixed effects regressions

Demand shock index construction			Time-invariant coefficients	Smoothly varying coefficients
Labor market	Demand measure	Detrended		
demographic	log hours	no	59.2	38.8
demographic	log hours	yes	72.0	81.9
demographic	emp. share	no	20.3	18.3
demographic	emp. share	yes	34.3	36.8
sector	log hours	no	59.3	55.3
sector	log hours	yes	73.0	73.9
sector	emp. share	no	29.1	41.6
sector	emp. share	yes	30.3	29.2

Notes: Each result is the percentage of the change in economy-wide earnings volatility between 1975 and 2005 that is explained by an index of labor demand shocks, where the construction of the index is described by the first three columns: whether education-state demographic groups or occupation-industry sectors are taken to be the appropriate labor markets, whether log hours or economy-wide employment share is taken to measure a sector's demand, and whether the sector's demand is detrended before the demand shock index is measured. See text for further details. The results come from conditional decompositions similar to those in Table 5.8, and include all the controls listed in that table.

CHAPTER 6
CONCLUSION

The literature on earnings volatility contains an array of terminology, measures, and conclusions. The first part of this dissertation provides some clarity. I identify conceptual differences among previous measures of earnings volatility, then make some methodological recommendations to guide measurement in future studies, based on the fact that potential welfare losses are the primary motivation for studying earnings volatility.

I show that, by my preferred measure, economy-wide earnings volatility has increased among men in the U.S. since the 1970s. This phenomenon has presented a puzzle for labor economists since it was first documented by Gottschalk and Moffitt (1994), and the increase has continued since then. The empirical part of this dissertation is the first to attempt to solve that puzzle using worker-level data and the first to analyze multiple potential reasons for increasing economy-wide earnings volatility. I develop decomposition approaches to quantify the extent to which various mechanisms can account for the increase in economy-wide earnings volatility over time.

I construct a demand shock index to capture the labor market demand shifts to which workers are exposed through national changes in the occupation-industry mix of hours worked. The relationship between earnings volatility and demand shocks is strong both unconditionally and after adding numerous controls, and is robust to many alternative constructions of the demand shock index. I find that labor demand shocks explain about half of the increase in economy-wide earnings volatility between 1975 and 2005.

Worker and job characteristics explain close to none of the change in economy-wide earnings volatility. This shows that rising economy-wide earnings volatility is not attributable to salient labor market trends, such as the decline of unions and the increase in the proportion of workers with a college degree. I also test previous explanations for rising economy-wide earnings volatility, which focus more frequent job, occupation, or industry switches, and find that these things cannot explain increasing economy-wide earnings volatility.

The literature on labor market risk often focuses on job change. Hallock (2009), for example, identifies a decrease in job duration among men, and Kambourov and Manovskii (2008) show that occupation and industry mobility have increased. In this paper, I find that larger or more frequent demand shocks have caused higher earnings volatility upon job change, but I also find that these shocks pass through to job stayers. This suggests a substantial change in the implicit contract governing employment relationships, an additional dimension of increasing labor market risk that deserves further study.

REFERENCES

Abowd, John M. and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57 (2):411–445.

Acemoglu, Daron and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics, Volume 4b*. Elsevier, 1043–1171.

Altonji, Joseph G. and Robert A. Shakotko. 1987. "Do Wages Rise with Job Seniority?" *The Review of Economic Studies* 54 (3):437–459.

Autor, David and David Dorn. 2009. "This Job is "Getting Old": Measuring Changes in Job Opportunities using Occupational Age Structure." *American Economic Review: Papers & Proceedings* 99 (2):45–51.

Autor, David H. and David Dorn. 2012. "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market." Working paper.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *The Review of Economics and Statistics* 90 (2):300–323.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics* 113 (4):1169–1213.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4):1279–1333.

Baker, Michael and Gary Solon. 2003. "Earnings Dynamics and Inequality among Canadian Men, 19761992: Evidence from Longitudinal Income Tax Records." *Journal of Labor Economics* 21 (2):289–321.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *The Quarterly Journal of Economics* 112 (2):537–579.

Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2013. "The Great Reversal in the Demand for Skill and Cognitive Tasks." NBER Working Paper No. 18901.

Blau, Francine D., Lawrence M. Kahn, and Jane Waldfogel. 2000. "Understanding Young Women's Marriage Decisions: The Role of Labor and Marriage Market Conditions." *Industrial and Labor Relations Review* 53 (4):624–647.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5):1887–1921.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers. 1994. "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics* 12 (3):345–368.

Bound, John and Harry J. Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics* 18 (1):20–54.

Brown, James N. and Audrey Light. 1992. "Interpreting Panel Data on Job Tenure." *Journal of Labor Economics* 10 (3):219–257.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2011. "Robust Inference With Multiway Clustering." *Journal of Business & Economic Statistics* 29 (2):238–249.

Celik, Sule, Chinhui Juhn, Kristin McCue, and Jesse Thompson. 2012. "Recent Trends in Earnings Volatility: Evidence from Survey and Administrative Data." Working paper.

Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *American Economic Review* 96 (5):1821–1834.

Comin, Diego, Erica L. Groshen, and Bess Rabin. 2009. "Turbulent firms, turbulent wages?" *Journal of Monetary Economics* 56 (1):109–133.

Dahl, Molly, Thomas DeLeire, and Jonathan A. Schwabish. 2011. "Estimates of Year-to-Year Volatility in Earnings and in Household Incomes from Administrative, Survey, and Matched Data." *Journal of Human Resources* 46 (4):750–774.

Davis, Steven J., R. Jason Faberman, and John Haltiwanger. 2012. "Labor market flows in the cross section and over time." *Journal of Monetary Economics* 59 (1):1–18.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." In *NBER Macroeconomics Annual 2006*. MIT Press, 107–179.

Deaton, Angus and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3):437–467.

DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. 2013. "Rising Inequality: Transitory or Permanent? New Evidence from a Panel of U.S. Tax Returns." Working paper.

Drewianka, Scott. 2010. "Cross-sectional Variation in Individuals' Earnings Instability." *Review of Income and Wealth* 56 (2):291–326.

Duncan, Greg J. and Daniel H. Hill. 1985. "An Investigation of the Extent and Consequences of Measurement Error in Labor-Economic Survey Data." *Journal of Labor Economics* 3 (4):508–532.

Dynan, Karen, Douglas Elmendorf, and Daniel Sichel. 2012. "The Evolution of Household Income Volatility." Working paper.

Dynarski, Susan and Jonathan Gruber. 1997. "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 1997 (1):229–284.

Fields, Gary S. and Efe A. Ok. 1996. "The Meaning and Measurement of Income Mobility." *Journal of Economic Theory* 71:349–377.

———. 1999. "Measuring Movement of Incomes." *Economica* 66:455–471.

Gorbachev, Olga. 2011. "Did Household Consumption Become More Volatile?" *American Economic Review* 101 (5):2248–2270.

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

———. 2009. "The Rising Instability of U.S. Earnings." *Journal of Economic Perspectives* 23 (4):3–24.

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1):192–205.

Guvenen, Fatih. 2009. "An empirical investigation of labor income processes." *Review of Economic Dynamics* 12 (1):58–79.

Guvenen, Fatih and Anthony Smith. 2010. "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach." NBER Working Paper No. 16327.

Haider, Steven J. 2001. "Earnings Instability and Earnings Inequality of Males in the United States: 1967–1991." *Journal of Labor Economics* 19 (4):799–836.

Hallock, Kevin F. 2009. "Job Loss and the Fraying of the Implicit Employment Contract." *Journal of Economic Perspectives* 23 (4):69–93.

Hryshko, Dmytro. 2012. "Labor income profiles are not heterogeneous: Evidence from income growth rates." *Quantitative Economics* 3 (2):177–209.

Jaimovich, Nir and Henry E. Siu. 2012. "The Trend is the Cycle: Job Polarization and Jobless Recoveries." NBER Working Paper No. 18334.

Juhn, Chinhui and Kristin McCue. 2010. "Comparing Measures of Earnings Instability Based on Survey and Administrative Reports." Working paper.

Kambourov, Gueorgui and Iourii Manovskii. 2008. "Rising Occupational and Industry Mobility in the United States: 1968–97." *International Economic Review* 49 (1):41–79.

———. 2009a. "Occupational Mobility and Wage Inequality." *Review of Economic Studies* 76 (2):731–759.

———. 2009b. "Occupational Specificity of Human Capital." *International Economic Review* 50 (1):63–115.

Katz, Lawrence F. and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107 (1):35–78.

King, Miriam, Steven Ruggles, J. Trent Alexander, Sarah Flood, Katie Genadek, Matthew B. Schroeder, Brandon Trampe, and Rebecca Vick. 2010. "Integrated Public Use Microdata Series, Current Population Survey: Version 3.0." Machine-readable database. Minneapolis: University of Minnesota.

Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent. 2009. "Performance Pay and Wage Inequality." *The Quarterly Journal of Economics* 124 (1):1–49.

MaCurdy, Thomas E. 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.

Medoff, James L. 1979. "Layoffs and Alternatives under Trade Unions in U.S. Manufacturing." *American Economic Review* 69 (3):380–395.

Moffitt, Robert A. and Peter Gottschalk. 1995. "Trends in the Covariance Structure of Earnings in the U.S.: 1969–1987." Working paper.

———. 2002. "Trends in the Transitory Variance of Earnings in the United States." *The Economic Journal* 112 (478):C68–C73.

———. 2012. "Trends in the Transitory Variance of Male Earnings: Methods and Evidence." *Journal of Human Resources* 47 (1):204–236.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25 (3):173–196.

Parent, Daniel. 2000. "Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics." *Journal of Labor Economics* 18 (2):306–323.

Pischke, Jorn-Steffen. 1995. "Measurement Error and Earnings Dynamics: Some Estimates from the PSID Validation Study." *Journal of Business & Economic Statistics* 13 (3):305–314.

Pistaferri, Luigi. 2009. "Comment on: 'Turbulent firms, turbulent wages?' by Comin, Groshen, and Rabin." *Journal of Monetary Economics* 56 (1):134–136.

Shin, Donggyun. 2012. "Recent Trends in Mens Earnings Volatility: Evidence from the National Longitudinal Survey of Youth, 1985-2009." *The B.E. Journal of Economic Analysis & Policy (Topics)* 12 (2).

Shin, Donggyun and Gary Solon. 2011. "Trends in men's earnings volatility: What does the Panel Study of Income Dynamics show?" *Journal of Public Economics* 95 (7–8):973–982.

Stevens, Ann Huff. 2001. "Changes in Earnings Instability and Job Loss." *Industrial and Labor Relations Review* 55 (1):60–78.

Stevenson, Betsey and Justin Wolfers. 2007. "Marriage and Divorce: Changes and their Driving Forces." *Journal of Economic Perspectives* 21 (2):27–52.

Strain, Michael R. 2011. "Understanding Earnings Instability: Skills and Employers." Working paper.

Sullivan, James X. 2008. "Borrowing During Unemployment: Unsecured Debt as a Safety Net." *Journal of Human Resources* 43 (2):383–412.

Violante, Giovanni L. 2002. "Technological Acceleration, Skill Transferability, and the Rise in Residual Inequality." *The Quarterly Journal of Economics* 117 (1):297–338.

Ziliak, James P., Bradley Hardy, and Christopher Bollinger. 2011. "Earnings volatility in America: Evidence from matched CPS." *Labour Economics* 18 (6):742–754.