Hedging Labor Income Risk*

Sebastien Betermier Thomas Jansson Christine A. Parlour
and Johan Walden †

July 28, 2011

Abstract

We use a detailed panel data set of Swedish households to investigate the relation between their labor income risk and financial investment decisions. In particular, we relate changes in wage volatility to changes in the portfolio holdings for households that switched industries between 1999 and 2002. We find that households do adjust their portfolio holdings when switching jobs, which is consistent with the idea that households hedge their human capital risk in the stock market. The results are statistically and economically significant. A household going from an industry with low wage volatility to one with high volatility will *ceteris paribus* decrease its portfolio share of risky assets by up to 35%, or USD 15,575.

*We have benefited from helpful comments by Jonathan Berk, Ralph Koijen, Richard Stanton, seminar participants at U.C. Berkeley and the LSE, and conference participants at the HEC Mathematical Finance Days and the WFA Meetings. Any errors are our own.

†Betermier is at the Desautels Faculty of Management, McGill Email: `sebastien.betermier@mcgill.ca`, Jansson is at the Swedish Central Bank, E-mail: `thomas.jansson@riksbank.se`, Parlour and Walden are at the Haas School of Business, U.C. Berkeley E-mail: `parlour@haas.berkeley.edu`, `walden@haas.berkeley.edu`

1 Introduction

Labor income accounts for about two thirds of national income in the U.S. and, since the seminal work of Mayers (1973), it has been assumed to play an important role in theoretical asset pricing. In studies such as Bodie, Merton, and Samuelson (1992), Danthine and Donaldson (2002), Qin (2002), Santos and Veronesi (2006) and Parlour and Walden (2011), risky labor income—or more generally, human capital risk—affects investors' portfolio decisions, which in turn has general equilibrium asset pricing implications. Broadly, the theory suggests that the behavior of capital markets can only be understood together with labor markets. More specifically, the theory suggests that an important function of capital markets is to allow investors to hedge their labor income risk.

Are investors' portfolio decisions affected by their labor income risk? Studies that use aggregate labor income data find mixed evidence. Fama and Schwert (1977) find that adding a labor factor does not improve the performance of the unconditional CAPM. By contrast, Jagannathan and Wang (1996) find that an aggregate labor factor significantly improves the performance of a conditional CAPM in explaining the cross section of expected returns. Lustig and Van Nieuwerburgh (2008) argue that in a standard representative agent model the observed aggregate consumption dynamics are inconsistent with a positive relation between returns on human capital and financial returns. On the other hand, using co-integration analysis Benzoni, Collin-Dufresne, and Goldstein (2007) argue that returns to human capital and financial returns should be highly correlated, which may explain the hump-shape life-cycle portfolio holdings of households. Given highly aggregated data, noisy measurements, and incomplete real-world markets, it seems unlikely that an approach based data at the aggregate level can lead to a conclusive answer.

In this paper, we use data at the individual household level. We study panel data on the employment and portfolio holdings of a large subset of the Swedish population between 1999 and 2002, and examine whether there is a relation between the workers' wage structure (measured by wage level and volatility) and their portfolio holdings of risky assets. More specifically, we focus on households in which some of the members switch industries over time and examine how they adjust their portfolios in response to their job changes. This approach allows us to control for a variety of household unobserved "taste" characteristics that are invariant to the switch itself, which is one of the main challenges for empirical work on this topic.

We find that households do adjust their portfolio holdings of risky assets in when switching jobs, which is consistent with the idea that human capital risk affects portfolio decisions. This effect, which is highly statistically significant, is especially strong for job changes that lead to large changes in wage volatility: a household that experiences an increase in wage volatility by 20% decreases its portfolio share of risky assets by 20%. This means that a household going from the industry with the least variable wage in the sample (recycling metal waste) to the industry with the most variable wage (fund management) *ceteris paribus* decreases its share of risky assets by up to 35%, or 15,575 USD. If wages are on average positively correlated with the stock market, then this effect corresponds to the workers' hedging demand for aggregate human capital risk.

Our main contribution is thus to document hedging behavior in stock markets, in line with the theoretical literature, by following individual households over time and thereby controlling for cross sectional "taste" differences, e.g., in risk-preferences, familiarity bias, or heterogeneous information among households. In particular, our approach allows us to control for any source of heterogeneity that is reflected in portfolio holdings. Our data is also of better quality than that used in most previous studies. We use the Longitudinal Individual Data for Sweden (LINDA) database from 1999 to 2002, which provides detailed income and wealth information for a large representative sample of about 3% of the Swedish population at the end of each year.

Although we establish a strong link between *changes* in human capital risk and *changes* in portfolio holdings, the results are weaker when we examine *levels*. We take this as evidence of cross-sectional "taste" differences. If any of these taste factors vary with the business cycle, then our results are consistent with a world in which a human capital factor is of little help in an unconditional CAPM (as argued in Fama and Schwert, 1977), but significantly improves the performance of a conditional CAPM (as argued in Jagannathan and Wang, 1996).[1] Heterogeneity in these "taste" preferences may explain the mixed evidence for the importance of labor income risk in the aggregate.

The weaker results that we obtain when we examine levels are also consistent with the lack of hard evidence from previous studies that have relied on household level data. Heaton and Lucas (2000) use the Panel of Individual Tax Returns, which provides information on income and assets for a large panel with annual frequency. They compute, for each

[1] See also Campbell (1996), Lettau and Ludvigson (2001), Palacios-Huerta (2003), and Santos and Veronesi (2006).

individual, an estimate of wage volatility and then study the effect on their average portfolio share of risky assets. They find that, while levels of entrepreneurial risk have a significant influence on portfolio holdings, the effects of wage income risk is not significant. Guiso, Jappelli, and Terlizzese (1996) use a cross-sectional dataset of Italian households in 1989 which asks them to attribute probability weights to intervals of nominal income increases one-year ahead. They find evidence that households that expect high future wage volatility hold relatively low shares of risky assets. Gakidis (1998) and Vissing-Jorgensen (2002) use panel data from the Panel Study of Income Dynamics and also find that high levels of future wage volatility have a negative effect on both the probability of being a stockholder and the share invested in risky assets conditional on owning stocks. On the other hand, Massa and Simonov (2006) look at individual stock holdings using panel data from Sweden and find that households tend to hold stocks that are closely related to their labor income, which goes against the hypothesis of hedging of labor income risk. They argue that this is because of a preference for familiar stocks due to heterogeneous information, which would fall within our definition of individual taste differences. Our main result—that we find a significant hedging demand for human capital risk when following individual households over time—is consistent with Massa and Simonov's results, since they find that the familiarity bias is considerably smaller for households that switch professions or locations, or who experience an unemployment shock.

A limitation of our approach is that job switches may not be exogenous events. If job switching decisions are driven by the same taste preferences that affect portfolio rebalancing decisions and these preferences change, or are are not fully reflected in the initial portfolio holdings of the switcher households, then our estimates may still be prone to an omitted variable bias. We address this issue with several robustness tests. In particular, we use information on the households' behavior in the years before our tests begin (except for portfolio holdings, we also have information about the households from 1996-99). If the decision to switch to a riskier or safer industry during the bear market years of 1999-02 depends on the type of households, then their type should also affect their decision to switch industries during the previous bull market years of 1996-98. In this case, we should observe a relation between households' changes in wage volatility in 1996-98 and their portfolio rebalancing decisions in 1999-02. We find no such relation, which suggests that our main findings are driven by hedging motives.

The rest of this paper is organized as follows. In Section 2, we lay out the theoretical predictions along with our main empirical strategy. We describe the data in Section 3 and the methodology in Section 4. In Section 5, we provide the empirical results, and in Section 6 we offer some concluding remarks. Further information about the theoretical background, the construction of variables, and the robustness tests is provided in an online Appendix.

2 Theoretical background, predictions and empirical strategy

2.1 Theoretical background and predictions

Recently, a literature has studied the general equilibrium asset pricing implications of human capital risk, see Dreze (1979), Danthine and Donaldson (2002), Qin (2002), Santos and Veronesi (2006), Lustig and Van Nieuwerbugh (2008), Parlour and Walden (2011), Palacios (2010), and Berk and Walden (2010). These studies examine the interplay between labor income risk and stock market risk in agents' portfolio problems. Documenting that agents treat labor income and capital market investments jointly, by hedging labor income risk, is necessary for the theoretical literature on human capital risk, portfolio choice and asset pricing to have any practical implications.

In the appendix we introduce a stylized GE model to motivate the predicted relation between workers' wage volatility and their investment in the stock market. Briefly, the static model, which is a simplified version of Parlour and Walden (2011), introduces a framework where risk averse agents can choose how much they work for a firm and also how to invest their wealth in capital markets. Firms rely on labor to produce a consumption good, which they sell in the market, using the proceeds to pay wages and dividends. Wages are perfectly correlated with stock returns, implying that wage volatility determines the covariance between human capital returns and financial returns, and hence the magnitude of hedging demand in the stock market. The model yields two predictions:

H1. *Levels:* The higher a worker's wage volatility, the lower his/her exposure to the market through financial assets.

H2. *Changes:* A worker who switches to a sector with higher wage volatility decreases his/her exposure to the market through financial assets.

The model in the appendix is very stylized, but these predictions are valid under more general conditions. They extend to a dynamic setting with a constant investment opportunity set, which leads to identical results at each point in time. They also extend to the introduction of idiosyncratic labor income risk, in which case wages are no longer perfectly correlated with stock returns. If the correlation between human capital and financial returns is positive and constant across all industries, then an increase in the volatility of human capital returns (proxied by wage volatility) will still translate one-for-one into an increase in the covariance between human capital and financial returns. More generally, if the correlation between human capital and financial returns is positive in expectation and the cross-sectional distribution of idiosyncratic human capital risk across industries is i.i.d., then wage volatility will provide an unbiased noisy measure of the covariance and the predictions will hold.[2]

2.2 Empirical strategy

The main challenge for empirical studies on this topic is that there may be other sources of heterogeneity that are correlated with labor income and also affect portfolio investment decisions. For example, as Massa and Simonov (2006) point out, workers may want to invest more in the industry they work in because they are more familiar with this industry. Or, it may be that the less risk averse agents choose to work in riskier industries and invest more in the stock market. Indeed, we show in the stylized model in the appendix that if enough risk tolerant agents choose to work in high wage-risk firms, then a statistical test of the relation between wage risk and investment portfolios may yield an outcome of "anti-hedging." In other words, the endogeneity introduced by heterogeneous tastes makes such a test inconclusive. Studies that rely on cross-sectional data are especially prone to this omitted variable bias because these taste differences among households are unobservable and hence very difficult to control for.

The previous argument implies that hypothesis H1, on levels, is difficult to test. In this paper, we therefore focus on testing H2, i.e., we focus on portfolio changes when households switch jobs. By conditioning on households' portfolio holdings before their switch, we are able to control for any taste differences that are reflected in their initial portfolio holdings. For example, differences in risk aversion between households would typically be reflected in

[2] A positive covariance between shocks to wages and stock returns has been reported in Heaton and Lucas (1996), Campbell, Cocco, Gomes, and Maenhout (2001) and Cocco (2005), using household-level data.

different initial portfolio holdings.

Furthermore, rebalancing due to changing market conditions will also typically be controlled for. The analysis of Merton (1969) suggests that investors should invest a fraction $\frac{\mu-r}{\gamma_i\sigma^2}$ of their wealth in the risky asset, where μ and σ^2 are the expected return and the variance of the asset respectively, and γ_i the relative risk aversion of CRRA agent i. If households revise down their views on μ during bear market years, they decrease the share of wealth invested in risky assets and the extent to which they do so depends on their level of risk aversion. If the highly risk averse agents are also the ones who switch into the lower risk jobs, this introduces a link between job switching and portfolio rebalancing. Now, since there is a direct link between risk-aversion and the initial portfolio holdings in this case, this effect would also be controlled for in our tests.

Nevertheless, a limitation of our approach is that a job switch may not be an exogenous event. First, a job switch may be part of a major life change, which also affects a household's attitude toward savings, risk, and other determinants of portfolio holdings, i.e., the switch and portfolio rebalancing may be due to a "taste shock." Second, when frictions that lead to infrequent portfolio adjustments are present, the initial holdings of risky assets may not control for all sources of heterogeneity. For example, two agents with different degrees of risk aversion may have the same portfolio holdings if they rebalanced at different points in time. If, in addition, there is a correlation between job switching and risk aversion, e.g., in that more risk averse households switch to safer jobs in bad times, then our estimates of a hedging effects may be biased. We address these issues in Sections 4.4 and 5.3, respectively, by introducing further controls. Although we cannot rule out endogeneity, our results suggest that hedging is indeed present.

3 Description of the datasets

3.1 Overview

LINDA (Longitudinal INdividual DAta for Sweden) is an annual cross-sectional sample of around 300,000 individuals, or approximately 3% of the entire Swedish population.[3] Select individuals and their family members are tracked over the years. The sampling procedure ensures that the panel is representative of the population as a whole, and each annual cohort is cross-sectionally representative. The values of all the variables in year t correspond to

[3]The data set is a joint project between Uppsala University, The National Social Insurance Board ("Försäkringskassan") Statistics Sweden, and the Swedish Ministry of Finance.

the values on December 31 of that year.

The data are primarily based on filed tax reports (available on an annual basis from 1968) and include various measures of income, government transfers and taxes in addition to individual characteristics such as gender, marital status, education, municipality of residence, and country of birth. We do not have information on the identity of a worker's employer but we do know the industry he or she works in. In LINDA, any working individual is assigned a five-digit SNI code – the Swedish equivalent to the NAICS/SIC codes in the USA – for the industry in which he or she made most income during the year. Unless specified otherwise, we work with SNI codes at the three-digit level because they provide sufficient granularity: in total there are 223 3-digit codes.

From 1999 onwards, the market values of financial and real assets (e.g. stocks, bonds, mutual funds, and owner-occupied homes) are included in LINDA. The values for the financial assets are actual values and not estimates, because in Sweden banks and financial institutions are required by law to report the market values of individual holdings – except for the very small bank accounts for which the interest rate earned is below 100 SEK a year. The values of real estate holdings are estimated from Statistics Sweden, which uses tax-assessed values and actual transaction prices in the surrounding areas.

To control for agent heterogeneity, we also use a Statistics Sweden demographic data set which provides information on the population density of the various Swedish regions. Since the region where individuals live is available in LINDA, we can merge these two datasets and use population density as a control in our regressions on portfolio holdings. This data set groups regions into six different categories, based on the population composition at the end of year 2002.

3.2 Excluded data

We have access to the LINDA dataset from 1993 to 2003. While we use the entire data in a couple of instances, our primary period of focus is 1999-2002. There are three reasons for this. First, we need information on the portfolio holdings, which is only available from 1999. Second, the 2000-2002 period corresponds to the Bear market in Sweden. Since our measure of changes in portfolio holdings involves a three-year horizon and is sensitive to market returns, the 1999-2002 period provides a homogeneous environment for our tests. Finally, this period allows us to conduct robustness checks against Calvet, Campbell, and Sodini (2009), who have access to all individual stock holdings for the entire Swedish population during the same period. We have information on the market value of broad asset categories such as directly-held stocks and mutual funds and we show that our measure of changes in households' holdings of risky assets over time approximates the changes reported in Calvet, Campbell, and Sodini (2009) quite well. Overall, there are 230,000 households that exist in the data for the entire 1999-2002 period and that do not undergo any major change in their civil status (see below).

We also run several additional filters to eliminate unusual data (e.g. households with very low or negative wealth, no industry code, outliers). We end up with a sample of 73,346 households. Unless specified otherwise, our tests are based on this sample. More information on our filters is provided in the appendix.

4 Construction of variables

Portfolio decisions are typically made at the household level so we track households (h) over the years (t). Our approach requires that we keep track of the industries where household members work. We also need measures of portfolio holdings and wage volatility. While aggregating household financial holdings is straightforward, imputing wage volatility to a household is less so.

4.1 Household characteristics and industries

In LINDA, two adult individuals belong to the same household in a given year if they are either married, legal partners, or if they live together and have children in common. We study the households that existed for the entire 1999-2002 period and where the head couple (or the single head member) remained the same. To identify the head of the household, we select the two adults who generate the greatest levels of income in 2001. We sort these two individuals by income, and adopt the convention that Individual #1 (Ind1) generates the highest income and Individual #2 (Ind2) is the other adult. In the case in which only one adult exists or generates income we treat Ind2 as missing.

We define a "switcher" as a household in which at least Ind1 changed SNI codes between 2000 and 2001. In other words, our switcher worked in the old industry in 1999 and 2000, switched to a new industry in 2001, and stayed in the same new industry in 2002. This also includes individuals who entered or quit the workforce in 2001. We choose 2000–2001 as the switch year to take into account the fact that investors may not adjust their portfolios immediately before or after a job change, as documented in Calvet, Campbell, and Sodini (2009). Households where individuals switch to industries with higher (lower) wage volatility are referred to as the "up-switchers" ("down-switchers"). For comparison, we also define a "non-switcher" as a household where neither Ind1 nor Ind2 changed industries between 1999 and 2002.

Summary statistics for the overall population as well as for the 3,815 switchers are displayed in Table 1 for 1999. The ex ante characteristics of switchers are broadly similar to the overall population. However, switchers are slightly more likely to live in one of Sweden's big three metropolitan areas, to have a college degree, and to have studied business.

[Table 1 about here.]

4.2 Portfolios

4.2.1 The share of risky assets

For each household, we examine its non-retirement portfolio of directly-held stocks and risky mutual funds. We refer to this portfolio as the portfolio of risky assets. Unfortunately, retirement portfolios are not available in LINDA, but we note that in 1998, Sweden switched from a defined benefit plan ("Allmän Tjänste Pension," ATP) to a defined contribution plan (see Sunden, 2006). Since no changes were made retroactively, pension capital accumulated up to our time period was low-risk. Risky mutual funds include pure-equity funds as well as funds that invest only a positive fraction of their assets in stocks. Ideally we would like to separate these two types of mutual funds but unfortunately this information is not available after 1999. From the 1999 data, however, it seems that the vast majority of these funds are pure-equity (about 85%).

At the end of each year t, we define the "risky share," denoted by $w_{h,t}$. This is the share of household h's holdings of risky assets over its financial wealth, which is the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds, and capital insurance and other products. So, $w_{12,02}$ refers to household #12's share of risky assets in its financial wealth at the end of the year 2002.

Summary statistics on portfolio shares of the overall population as well as those of switchers in 1999 appear in Panel A of Table 2. All the moments are equal-weighted by household. Although the switchers are broadly representative of the population, they are slightly more likely to invest in stocks than the other households.

[Table 2 about here.]

Compared to US investors, Swedes in our data hold more risky assets and are more likely to invest in mutual funds. To see this, consider statistics from the US 2001 Survey of Consumer Finances (SCF). In the first set of columns in Table 3 we report the (equal-weighted) moments of the 2001 portfolio shares for the overall Swedish population. In the second set of columns (SCF I), we report the moments of the equivalent portfolio shares for the US population from the SCF. Note that to make the comparison relevant, these US statistics are not the ones that are usually reported from the SCF. In the standard definition of the risky share from the SCF, the amount of mixed mutual funds is halved and retirement assets are included. To see how these modifications affect our statistics from the SCF, we also report the standard statistics in the third set of columns (SCF II).

[Table 3 about here.]

Comparing the first two sets of columns of Table 3, it is evident that the participation rate in risky assets is much higher in Sweden than in the USA. High Swedish stock-market participation rates have been documented elsewhere (Georgarakos and Pasini, 2009), and suggest that the selection bias in stock market participation is not as important as it is in the USA. Swedish households also tend to invest much more of their risky assets in mutual funds than American households. This may be due to the introduction in the late 1970's of highly accessible mutual funds (so-called "Allemansfonder"), which offered high tax-incentives. The tendency towards well–diversified investments is consistent with our empirical analysis because our measure of hedging is the share of financial assets invested in risky assets. As we cannot observe Swedish households' detailed portfolio of stock holdings, observing a high portfolio share in mutual funds indicates that these households are likely to be mostly invested in the overall stock market. As a result, if these households hedge their labor income risk, they are likely to do so by levering up or down their holdings of mutual funds.

4.2.2 Active portfolio rebalancing

In Panel B of Table 2, we also report statistics on portfolio shares in 2002. The equal-weighted average of the risky share dropped by about 9% (in levels) between 1999 and 2002. This drop is consistent with the significant decrease in the value of the Swedish stock market from 2000 to 2002. The total return on the Morningstar index for stock mutual funds[4] was 0.596 (i.e., the return rate was -41%). In comparison, the total return on the 12-month Swedish government bills (SSVX) during the same time period was 1.135 (Source: Thomson Reuters).

To distinguish changes that simply come from changes in the returns on risky assets from changes that come from portfolio rebalancing decisions, we follow Calvet, Campbell, and Sodini (2009) and decompose the total change in the risky share $\Delta w_{h,02}$ of any household into a passive change, $\Delta^p w_{h,02}$, and an active change, $\Delta^a w_{h,02}$,

$$\Delta^p w_{h,02} = w_{h,99} \left(\frac{R_{02}}{w_{h,99} \cdot R_{02} + (1 - w_{h,99}) \cdot Rf_{02}} - 1 \right), \tag{1}$$

$$\Delta^a w_{h,02} = \Delta w_{h,02} - \Delta^p w_{h,02}, \tag{2}$$

where R_{02} and Rf_{02} correspond to the cumulative total returns on the risky and risk-free portfolios from 1999 to 2002. Since we do not observe the exact composition of these portfolios, we assume that $R_{02} = 0.596$ and $Rf_{02} = 1.135$ based on the indices described above. As we note below, our results approximate well those of Calvet, Campbell, and Sodini (2009) who have information on the households' exact portfolio holdings.

The passive change $\Delta^p w_{h,02}$ corresponds to the change in the risky share if household h did not trade any financial assets between 1999 and 2002. The active change $\Delta^a w_{h,02}$ is defined as the difference between the total change and the passive change. It represents portfolio rebalancing decisions. A positive (negative) active change means that household h bought (sold) risky assets between 1999 and 2002.

In Fig. 1, we show this decomposition of the total change into a passive change and an active change, as a function of initial share, $w_{h,99}$. To filter out noise and get a smooth approximation of total change as a function of $w_{h,99}$, household changes have been projected (regressed), using three cubic splines in the figure. Several insights follow from this decomposition. First, the average active change in the risky share across all households is close to zero, which is consistent with the general equilibrium restriction on portfolio rebalancing. Second, not all households experienced the same passive decrease in their risky share. The reason is purely mechanical. The passive change in the risky share is always negative because of the Bear market during these years and it follows a U-curve. By definition, if

[4] Available on www.morningstar.se. Morningstar mutual fund index for stock mutual funds are available, both for investments in Sweden and abroad.

a household invested only in risk-free assets ($w_{h,99} = 0$) or in risky assets ($w_{h,99} = 1$) in 1999, changes in the value of the stock market do not affect the composition of the one-asset portfolio, so the passive change in the risky share is zero. For very unbalanced portfolios ($w_{h,99}$ close to 0 or 1), the passive change is small because, even with a highly negative stock return, the portfolio remains very unbalanced. For example, if a household owned \$99 of stocks and \$1 of bonds in 1999 ($w_{h,99} = 0.99$), a 40% decrease in the value of the stock market would decrease its risky share by only 0.6% (in levels). However, for balanced portfolios, the passive change in the risky share is much greater. If the same household owned \$50 of stocks and \$50 of bonds in 1999 ($w_{h,99} = 0.5$), then a 40% decrease in the value of the stock market would decrease its risky share by 12.5% (in levels). Finally, we note that our computation of active and passive changes based on the indices of risky and risk-free assets provides a close approximation to the results in Calvet, Campbell, and Sodini (2009). They have access to the exact stock holdings of the entire Swedish population and compute active and passive changes based on all individual stock returns between 1999 and 2002. The predicted values of the active and passive changes in Fig. 1 are very similar to those in Fig. III.A in Calvet, Campbell, and Sodini (2009).

[Figure 1 about here.]

4.3 Wage volatility

Computing a measure of annual wage volatility for switcher households is difficult because we only have data for at most two years after a 2001 switch. So we compute industry-averages of wage volatility (which we describe in detail below) and then attribute these values to all individuals based on the industry in which they worked that year, and aggregate by household each year.

Even though industry-averages of wage volatility are crude proxies for individual agents, if agents are unaware of how their particular careers will evolve, then industry averages may well reflect an agent's ex ante information about the true values. Therefore, these variables should be informative. Furthermore, for the switcher households, these measures should do a good job in identifying the *change* in wage volatility or productivity that is associated with changing industries.

In the large LINDA sample from 1993 to 2003, we select all the individuals who work in the same industry for at least five consecutive years. Then, we compute the volatility of the annual growth rate of each individual's real disposable income during these years, and average this volatility across all the households within the same sector. This measure takes into account unemployment risk. If a worker is let go during a year, he will still be assigned his former SNI code as long as he was employed during part of the year.

Table 4 reports the top and bottom ten industries ranked by wage volatility. It is not surprising to find that industries such as "fund management," "legal representation activities," and "motion picture and video production" have high wage volatility whereas industries such as "recycling of metal waste and scrap" and "mining of iron and ores" have low wage volatility.

[Table 4 about here.]

Once we have computed these measures of the volatility and level of wages for each three-digit industry, we assign them to each individual-year given their SNI code. Finally, we aggregate these measures by household, weighting each individual by the amount of disposable income he or she earned during that year. In other words, if the household is composed of two working individuals, then the household labor income volatility measure is a weighted average of the individuals' volatility. In reality, the household labor volatility should also include the covariance between both individuals' labor income. However, given that we are working with industry-level estimates for their labor income, estimating this covariance precisely is difficult. In our regression we try to correct for this by creating a dummy to catch whether both individuals work in the same three-digit SNI code.

Another simple measure of wage volatility is whether an individual works in the public or the private sector. We have this information available in LINDA. It is well-known in Sweden that jobs in the public sector are less risky than in the private sector, in terms of unemployment risk and wage volatility. It is therefore not surprising to find in LINDA that the average wage volatility for employees in the public sector (12.9% per year) is lower than that in the private sector (14.9%). We use this measure as a robustness check. Note that while we keep the same sample of households, with this alternative measure we need to re-define which households are considered switchers and non-switchers. For this measure, the up-switchers (down-switchers) are households where at least Ind1 switches from the public sector (private) to the private one (public). Non-switchers are households where both individuals don't switch between the public and the private sectors between 1999 and 2002.

4.4 Endogeneity

As we discussed in Section 2.2, a potential source of concern is that a job switch may be part of a major life change — a "taste shock" — which jointly affects a household's attitude toward risk and portfolio decisions, without having anything to do with hedging. For example, if households when reaching a certain age and starting families reevaluate their attitudes toward risk, this may lead to simultaneous job switches and portfolio changes.

While we do not observe the reason for job switches, we can compare the characteristics of switchers and other households before and after the change, to rule out any observable differences between switchers and non switchers. The summary statistics from Tables 1 and 3 indicate that in 1999, the sample of switchers is fairly representative of the entire population. The equivalent summary statistics for 2002 are identical, which indicates that any major life change is likely to be idiosyncratic.

We then run five additional tests to compare the characteristics of switchers and other households. First, we compare statistics on wage volatility for three categories of households: the up-switchers, the down-switchers, and the non-switchers. Second, we study the distribution of industries in 1999 for the switchers and check whether they worked in different types of industries compared with non switchers. Third, we examine whether individuals who have already switched jobs are more likely to switch jobs again in the future. In all these three tests, which we report and discuss in greater detail in the appendix, we find no systematic differences between switchers and other households other than the fact that switchers are more likely to come from industries with higher job turnover.

Fourth, we look at the transition matrix of SNI codes for switchers between 1999 and 2002 and exclude the cases in which an unusually high number of individuals switch from a particular SNI code in 2000 to another particular SNI code in 2001. The empirical results remain the same. Finally, in the next section we compare the portfolio rebalancing decisions of the up-switchers to those of down-switchers and non-switchers. As we shall see, the active change in the risky share between 1999 and 2002 for the non-switchers is lower than for the down-switchers, but higher than for the up-switchers. This last result is consistent with switchers being of the same "type" and responding to shocks to their employment.

Altogether, we find little evidence of job switching being associated with major life changes that also affect investment decisions. As we mentioned earlier, we further discuss potential endogeneity issues in the light of our results in Section 5.3.

5 Empirical tests and results

5.1 Cross-section analysis of H1

What is the relation between a household's wage volatility and its financial portfolio? We begin with a cross-sectional analysis and test hypothesis H1.

H1: The higher a worker's wage volatility, the lower his/her exposure to the market through financial assets.

If agents only differ in the industries in which they work, we would expect a cross-sectional comparison of agents' wage volatility and investments in risky assets to have a negative relation.

In our data, we do find some evidence of hedging but the results go the wrong way in some cases, in line with the results in Massa and Simonov (2006). Thus, our results are consistent with the mixed findings from the previous literature. It could be that investors do not hedge labor income risk, but it could also be that there are cross-sectional taste differences between agents that drive wage volatility and portfolio decisions jointly, so that individual agents hedge but it does not show cross-sectionally. Our tests that control for such fixed effects in the next section support the latter view.

As in Vissing-Jorgensen (2002) and Massa and Simonov (2006), we assume that the investment decision takes place in two steps: first, the investor decides whether to enter the stock market, and then he selects his portfolio holdings. To account for the first stage participation decision, we use a two-step estimation procedure following Heckman (1979). We model the decision to enter the stock market by estimating $\mathbf{1}\{w_{h,02} > 0\}$, the observed probability of participation in the portfolio of risky assets in 2002, with the probit regression,

$$\mathbf{1}\{w_{h,02} > 0\} = \alpha_1 + \beta_1 \cdot LABOR_{h,02} + {\gamma_1}' \cdot X_{h,02} + \epsilon_{1,h,02}, \tag{3}$$

where $X_{h,t}$ is a vector of explanatory variables for household h in year t, and $LABOR_{h,t}$ includes wage volatility along with an interaction variable for households where both individuals work in the same industry. We report results for year 2002 because it allows us to include 1999 values for some potentially endogenous regressors such as wealth and income. If we choose $t = 2000$ or $t = 2001$ the results are similar.

In this and the subsequent regressions, the choice of control variables in the vector $X_{h,02}$ is critical because of the potential endogeneity issues. We control for each household's composition, where it is located, the sources and composition of household wealth and financial sophistication.

The various measures of household composition, location, real estate, and education (e.g. age, population density, college degree) are standard in the literature so we refer the reader directly to Tables 5 and Tables 6 and the appendix for further details. Measures of labor income and employment include the logarithm of family disposable income, a dummy on whether at least one of the adults is receiving unemployment insurance, a dummy on whether at least one of the adults is receiving a retirement pension, and the ratio of debts to family income. In addition to our measures of labor income risk $LABOR_{h,t}$, we add two dummies on whether both adults work in the private sector or the public sector. To avoid any endogeneity issues, both net worth and the ratio of house value to net worth are from year 1999. We avoid controlling for portfolio shares in previous years, because portfolio shares are extremely predictable over time, which means that including them would capture most of the information from the other variables, including $LABOR_{h,02}$.

Then, in the second stage, we regress the portfolio shares $w_{h,02}$ on $LABOR_{h,02}$, our proxy for wage volatility. Our main focus is on the portfolio share of risky assets (the risky share) but we also repeat the exercise for the portfolio shares of stocks and mutual funds. We also include the vector $X_{h,02}$ of control variables and Heckman's lambda variable ($\lambda_{h,02}$), which controls for possible selection at the first stage. The equation is as follows,

$$w_{h,02} = \alpha_2 + \beta_2 \cdot LABOR_{h,02} + \gamma_2{}' \cdot X_{h,02} + \theta_2 \cdot \lambda_{h,02} + \epsilon_{2,h,02}, \tag{4}$$

where h only includes the households that participate in the stock market in 2002. Households hedge their labor income risk if $\beta_2 < 0$.

The results of the second stage regressions are reported in Table 5. We run three specifications of Eq. (4). In the first column, we take a look at what the results look like if we do not control for selectivity. In the second column, we include $\lambda_{h,02}$ but only study the effect of wage volatility. In the third column, we include both $\lambda_{h,02}$ and the public-private sector dummies to see how much of the industry-wide differences in wage volatility comes from the differences between the private and the public sectors.

[Table 5 about here.]

Most of the control variables are strong predictors of the risky share. This is not surprising, and it is consistent with the literature. The coefficient on $\lambda_{h,02}$ also confirms the selectivity among market participants, despite the high overall participation rate in risky assets. We report the t-stats for the bootstrapped standard errors of the estimates and find that θ_2 is significantly different from zero. When we control for selectivity, the effect of wage volatility becomes more significant.

The results from Table 5 are consistent with H1. An increase in wage volatility does lead to a decrease in the risky share that is significant at the 1% level. This decrease is also fairly significant from an economic perspective. From the second column, a 5% increase in wage volatility (in levels) leads to a 1% decrease in share of risky assets (in levels). The magnitude of this effect is lower in the third column but that is because some of it is being picked up by the public-private sector dummies. A household where both individuals work in the public sector has a risky share almost 2% higher than a household where both individuals work in the private sector. These results are in line with those of Guiso, Jappelli, and Terlizzese (1996), Gakidis (1998), and Vissing-Jorgensen (2002).

However, once we decompose the risky share into the share of directly held stocks and the share of mutual funds, we get mixed results. In Table 6 we repeat the estimations of column 3 in Table 5 but this time with the shares of stocks and mutual funds as dependent variables. While a more formal analysis should involve estimating a system of simultaneous equations, we find that this heuristic analysis already provides interesting information. The key result is the opposite effect that $LABOR_{h,02}$ has on the shares of stocks and mutual funds. An increase in wage volatility leads to a significant *increase* in the share of stocks and a significant *decrease* in the share of mutual funds.

[Table 6 about here.]

The positive effect of $LABOR_{h,02}$ on the shares of direct stock-holdings reinforces the idea that our cross-sectional analysis is prone to an omitted-variable bias. This is consistent with what is found in Massa and Simonov (2006), who look at the levels of individual stock holdings and find that households' investments in stocks also come from factors other than hedging, such as a preference toward stocks they are more familiar with, for information reasons. Indeed, they argue that less-informed agents choose to invest more in stocks closely related to their labor income because they are more familiar with these stocks, via either location or professional proximity.

5.2 Analysis of job switches, H2

As we discussed in Section 2.2, the main weakness of the cross-sectional analysis above is that one can conjecture other sources of heterogeneity that are correlated with labor income and affect portfolio selection. Since our cross-section analysis cannot control for these unobserved taste differences, we turn to our main estimation strategy and look instead at changes in the portfolio shares of households over time, with a particular focus on those households where individuals change industries, i.e., we test hypothesis H2.

H2: A worker who switches to a sector with higher wage volatility decreases his/her exposure to the market through financial assets.

Our focus on changes in portfolio holdings over time is similar to adding fixed effects to Eq. (4) in that it allows us to control for any unobserved heterogeneity that is constant over time and correlated with the independent variables. It is important to point out, however, that a standard panel estimation of Eq. (4) with fixed-effects is hardly applicable in our setting. As mentioned earlier, since our time-series is short and not all households adjust their financial portfolios frequently, it is difficult to measure changes in the levels of wage volatility of households over time as well as their effect on the households' risky share. Consequently, a standard panel estimation would have very little power. We overcome this issue by modifying the standard panel model in three major ways.

The first unique feature is that we focus specifically on the households that switched industries between 2000 and 2001 and their portfolio re-balancing decisions between 1999 and 2002. This feature provides us with a pool of observations where the variation in our measures of changes in wage volatility over time is relatively high. The three-year horizon also provides a relatively large window of time to capture portfolio re-balancing decisions.

The second unique feature has to do with the way we control for the initial portfolio holdings. Instead of adding lagged values of the risky share to the right-hand side of Eq. (4), we study the variation in the active change in the risky share $\Delta^a w_{h,02}$ that is *orthogonal* to the initial level of the risky share $w_{h,99}$. This allows us to fully control for past portfolio choices and compare households that had the same initial risky share in 1999. Among these households, we can ask whether the ones that switch to riskier industries between 1999 and 2002 reduce their risky share relative to those that do not switch industries and to those that switch to safer industries.

Finally, the third unique feature is that even though our focus is on the switchers, we also use the group of non-switchers as a benchmark in the first stage where we back out the variation in $\Delta^a w_{h,02}$ that is orthogonal to $w_{h,99}$. Instead of running a first-stage regression of $\Delta^a w_{h,02}$ on $w_{h,99}$ over the pool of switchers and then using the residuals as our dependent variable for our second-stage regression on changes in wage volatility, we compare the switchers to the non-switchers in the first stage. That is, we begin with the pool of non-switchers and model their active change in the risky share, $\Delta^a w_{h,02}$, on their initial risky share, $w_{h,99}$. We keep the predicted values from this estimation. We then turn to the switchers and compute the difference between their active change in the risky share, $\Delta^a w_{h,02}$, and the *predicted value* of the active change for the non-switchers given the same level of $w_{h,99}$. This difference term becomes our dependent variable, which we can then regress on changes in wage volatility for the switchers between 1999 and 2002. Fig. 2 provides a visual representation of this construction, which allows us to test whether households that switch to sectors with the same level of wage volatility are equivalent (observationally) to the non-switchers.

[Figure 2 about here.]

This approach complements the one taken in Massa and Simonov (2006), who also use panel data from LINDA but focus more on the cross-sectional differences between households' labor income risk and their portfolio holdings. While their approach provides the opportunity to estimate the effect of any "taste" variable that does not vary much over time (if at all) and that can be measured like their indices of familiarity, it comes at the cost of not being able to include fixed effects and control for other sources of unobserved heterogeneity. In our approach, we only look at *changes* in household characteristics and portfolio holdings between 1999 and 2002. In doing so we are not able to estimate the effects of any of these "taste" variables, but we can fully control for all of them, whether they are observed or unobserved. This approach allows us to focus purely on the effects of the time variation in the wage volatility of households. We will see below that we find strong support for hedging along the time dimension. Their study and ours thus together suggest that both tastes (broadly defined) and hedging are present in the data.

From Fig. 1, it is clear that a household's active change in risky share depends on its initial risky share.[5] We control for this dependence on the initial risky share, using the same approach as in Fig. 1, i.e. by regressing the changes on three cubic splines. In the first stage we carry out this estimation for the population of non-switchers. The fitted values are depicted in the two left quadrants of Fig. 3. In the top left quadrant, we use the baseline sample of non-switchers that we defined in Section 4.1, which is tailored to the main wage volatility measure. In the bottom left quadrant, we use a slightly modified sample of non-switchers that is tailored to our second measure of wage volatility (whether individuals work in the public or the private sector, see Section 4.3). The results for both samples are very similar.

[Figure 3 about here.]

As a first test of whether switching jobs affects portfolio holdings, we also generate splines for the populations of households that switch to industries with higher wage volatility (the up-switchers) and those that switch to industries with lower wage volatility (the down-switchers) and we plot the *additional* $\Delta^a w_{h,02}$ (i.e. relative to the non-switchers) in the top right quadrant of Fig. 3. In the bottom right quadrant, we generate the same splines for households that switch between the private and public sectors. The top line (red) in each quadrant is the locus of predicted values for the down-switchers, and the bottom line (blue) is the equivalent line for the up-switchers. For clarity, we only select the switchers whose wage volatility changes by more than 1% (in levels). This involves about two-thirds of the switchers.

The results from Fig. 3 provide strong evidence in favor of hedging. The first key result is that the active change in the risky share $\Delta^a w_{h,02}$ is *always* greater for the down-switchers than for the up-switchers, which is consistent with the predictions. The difference between the two groups is economically important as well. If we take the average difference between the predicted values of the up- and down-switchers (weighted equally by $w_{h,99}$), we find in the top left quadrant that switchers who experience an increase in wage volatility tend to decrease their risky share by 1.57% relative to those that experience a decrease in wage volatility. From the bottom left quadrant, we see that households that switch to the private sector tend to decrease their risky share by 2.6% relative to those that switch to the public

[5] Such a dependence even arises for purely mechanical reasons. For example, the active change can only be positive if the initial share is zero, whereas it can only be negative if the initial share is one.

sector. These results are very robust to the types of basis functions used (see the online Appendix).

The second result from Fig. 3 is that the average differences between the active changes of the risky share $\Delta^a w_{h,02}$ of switchers and non-switchers are *negative* for the up-switchers and *positive* for the down-switchers. In other words, the up-switchers tend to decrease their risky share relative to the non-switchers, and the down-switchers tend to increase their risky share relative to the non-switchers. This result, although not as strong as the previous result, is still quite significant. We verify the result statistically, using a simple but very robust non-parametric sign test. The results are reported in Table 7. The hypotheses that the fitted curves for the up- and down-switchers are respectively above and below the fitted curve for the non-switchers are both strongly rejected at the 1% level. It is thus clear that changes in labor income risk affect the portfolio decisions of households, in line with our theoretical predictions.

[Table 7 about here.]

We next analyze the magnitude of these effects, to understand how big the hedging demand for labor income risk is. Let $\widehat{\Delta^a w_{s,02}}$ be the difference between the *observed* active change in the risky share $\Delta^a w_{s,02}$ of switcher household $h = s$ and the *predicted* active change in the risky share of non-switcher household $h = ns$ given the same initial share $w_{s,99}$. In Fig. 2, $\widehat{\Delta^a w_{s,02}}$ corresponds to the double-arrow vertical vector. We test the effect of a change in labor income risk on $\widehat{\Delta^a w_{s,02}}$ by estimating the following equation,

$$\widehat{\Delta^a w_{s,02}} = \alpha_3 + \beta_3 \cdot \Delta LABOR_{s,02} + \gamma_3 \cdot (\Delta Z_{s,02} - \overline{\Delta Z_{s,02}}) + \epsilon_{3,s,02}, \tag{5}$$

where s represents the switcher population, $\Delta LABOR_{s,02}$ represents the change in our measure of labor income risk between 1999 and 2002, and $(\Delta Z_{s,02} - \overline{\Delta Z_{s,02}})$ is a set of demeaned independent regressors. Note that we restrict the switchers to participate in the stock market in 1999. We do not include Heckman's lambda variable ($\lambda_{s,99}$), which controls for possible selection in 1999. Since our measure of $\widehat{\Delta^a w_{s,02}}$ is orthogonal to levels of the risky share in 1999, the selection bias is no longer an issue. As a test, we tried a version where we include $\lambda_{s,99}$. It comes up as insignificant and does not affect the other results.

We test the parameters α_3 and β_3. The first test is whether $\beta_3 < 0$. The theory predicts that switchers who experience an increase in labor income risk should decrease their risky share relative to the other switchers. The second test is whether $\alpha_3 = 0$. Since we demeaned the ΔZ variables, α_3 corresponds to the value of $\widehat{\Delta^a w_{s,02}}$ if $\Delta LABOR_{s,02} = 0$. The theory predicts that switchers who do not experience any change in their level of labor income risk should not invest differently than non-switchers. Their active change in the risky share should, on average, equal the predicted value of the active change of the non-switchers.

In addition to employment, other household characteristics may have changed during 1999-2002. $\Delta Z_{h,02}$ is defined as the vector of these changes. These variables include a dummy on whether the household moved from a low density area to a high density area, a dummy on whether at least one member of the household has emigrated, and a variable that captures the change in the number of children. We also look at the change in the logarithm of family disposable income, the change in the debt-to-income ratio and we include dummies on whether at least one of the individuals found a job, lost a job, or retired from the job market during the time period. In terms of real estate, we include two dummies on whether households started or stopped owning real estate as well as a variable that captures the change in the ratio of house value to net worth. In terms of education, we include a dummy on whether at least one of the individuals has graduated. In terms of changes in wealth, one has to be careful because of the potential endogeneity issues. We try two specifications: one with the change in net wealth between 1999 and 2002, and one without it. In both cases, all the other coefficients are approximately the same, which confirms that we can include net worth.

The results of our estimation are reported in Table 8. We run six specifications of Eq. (5). In the first column, we include all the variables in the vector $\Delta Z_{h,02}$. Unlike our regressions on the levels, only a select few of the control variables predict our measure of change in the risky share. So, to improve the precision of the estimation, we only retain in the second column the variables whose coefficient was statistically significant in the first column. In the third column, we exclude the change in net worth, to check whether it affects the other coefficients. In the fourth column, we interact $\Delta LABOR_{s,02}$ with dummies on whether the switchers are up-switchers or down-switchers. This is to check whether the effect of $\Delta LABOR_{s,02}$ is symmetric across both types of switchers. In the fifth column, we test whether the effect of the absolute value of $\Delta LABOR_{s,02}$ is quadratic rather than linear. Finally, in the sixth column, we focus on our sample of switchers with respect to the public-private measure. $\Delta LABOR_{s,02}$ becomes a dummy variable, so we include dummies for the up- and down- switchers and test that these dummies are negative and positive respectively.

[Table 8 about here.]

The results provide further evidence in favor of hedging, i.e., they support hypothesis H2. For the linear model (columns 1 to 3), an increase in wage volatility by 3% (in levels) leads to an active decrease in the share of risky assets by 1% (in levels). This means that a household going from the industry with lowest wage volatility to the industry with highest wage volatility would decrease its risky share by almost 10%. The one-tailed test that $\beta_3 < 0$ is statistically significant at the 1% level. The magnitude of this hedging effect is even stronger in the quadratic model in column 5. Because of the quadratic nature of the model, the effect on portfolio shares is quite small for small changes in wage volatility. But for large changes in wage volatility, the effect on the risky share increases considerably. For example, an increase in wage volatility of 20% leads to a decrease in the share of risky assets of almost 20%. The same household going from the industry with lowest wage volatility to the industry with highest wage volatility would decrease its risky share by 35%. Finally, we can check in column 4 that this hedging effect is fairly symmetric across the up- and down-switchers. Neither β_3 coefficient is as statistically significant as in the first three columns, but both coefficients are about the same size economically (although slightly greater for the down-switchers).

As for the second test on the value of α_3, we focus on the first five columns of Table 8.[6] Across all the estimations, we cannot reject the null hypothesis that $\alpha_3 = 0$. This is again consistent with the theory, i.e., switchers who do not experience any change in their level of labor income risk should have the same active change in the risky share as non-switchers. While this test is not as statistically powerful as the test on β_3, we see that the estimated value of α_3 is minimal from an economic perspective. The difference between the active changes in the risky share of switchers with no change in wage volatility and non-switchers is about 0.5%.

In terms of the estimation with the public-private sector dummies in column 6, the effects of the dummies are strong as well and consistent with the theory. Households where the high-income individual switches to the private sector decrease their risky share by

[6]Recall that the estimation with the public-private sector dummies in column 6 is run without an intercept.

1.6% relative to non-switcher households. Households where the high-income individual switches to the public sector increase their risky share by 0.8% relative to the non-switcher households. The one-tailed tests that the dummies for the up- and down-switchers are negative and positive respectively are statistically significant at the 1% and 10% level, respectively.

An alternative potential explanation for the fact that the coefficients of the changes in wage volatility are negative is if wage volatility is correlated with wealth. A change in wage volatility could be associated with a change in wealth, which could be the real driving force behind portfolio changes. As mentioned earlier, we control for this potential factor by looking at the change in net worth between 1999 and 2002. The addition of this variable acts not only as a control but it also indicates the effect of an increase in wealth on the risky share. If we compare columns 2 and 3, we find that the addition of net worth does not influence the effects of wage volatility and labor productivity. Moreover, we find that an increase in net worth leads to a significant decrease in the risky share. Note that we also control for changes in family income. Supposedly, households that switch to an industry where they obtain a wage increase have become wealthier. If we estimate Eq. (5) excluding labor income, we also find that the effects of wage volatility and labor productivity remain the same. And the coefficient on the labor income in all the columns is also negative. These results suggest that this other potential explanation goes the other way, hence strengthening our results.

It could also be the case that this hedging effect comes from a change in the switchers' housing situation, if this change is correlated with their change in labor income risk. We control for these housing effects by including the change in the households' ratio of housing wealth to net worth between 1999 and 2002 as well as dummies on whether they bought or sold their home and moved from a high density region to a low density region. While most of these variables have a significant effect on the households' change in the risky share, they do not affect the negative coefficients of the changes in wage volatility. These coefficients remain the same if we exclude all the housing variables. We conclude that the labor income hedging effect we observe does not come from housing.

Finally, it is important to point out that once we decompose the risky share into the share of directly-held stocks and the share of mutual funds, we no longer obtain the mixed results on hedging that our cross-sectional analysis was subject to. In Table 9 we repeat the estimation of column 2 in Table 8 but this time with the stocks and the mutual funds as the dependent variables. For example, for the stocks, our dependent variable becomes the difference in the observed active change of the share of directly-held stocks between switchers and predicted value of the active change for the comparable non-switchers.

[Table 9 about here.]

There are two main observations from Table 9. First, if we compare it to Table 6, we find that while the negative effect of $\Delta LABOR_{s,02}$ on the share of mutual funds remains, the positive effect of $\Delta LABOR_{s,02}$ on the share of stocks is no longer significant, neither statistically nor economically. In other words, the "anti-hedging" effect on directly-held stocks we found in the cross-section is no longer present in the time-series, which suggests that it really captures *time-invariant* differences in households' tastes. This result is consistent once again with the findings in Massa and Simonov (2006). The second observation from Table 9 is that the significantly negative effect of $\Delta LABOR_{s,02}$ on the shares of mutual funds is almost identical in size to the one on risky assets (from Table 8). This result confirms our intuition from Section 4.2.1 that households are most likely to hedge their labor income risk by levering up or down their holdings of mutual funds. Altogether, these two related observations provide additional support for hypothesis H2.

5.3 Additional controls for endogeneity

In Section 4.4 we addressed the issue of "taste shocks" in the form of major life changes as a source of endogeneity, finding very limited evidence for such effects. Another source of endogeneity is that households' tastes may not be fully reflected in their initial holdings of risky assets, in which case our analysis of job switchers in Section 5.2 may still be prone to an omitted variable bias.

For example, if there are transaction costs so that households only rebalance their portfolios infrequently, and there is also a systematic relation between job switching and risk-aversion so that households with high risk-aversion tend to down-switch in down-turns, this introduces a source of endogeneity that is not controlled for in our tests.[7] Specifically, with infrequent rebalancing, two households may in 1999 have the same portfolio share in risky assets but have different levels of risk aversion: a household with low risk-aversion may have just rebalanced its risky share downward after the market run-up (along the lines of

[7]We thank the referee for suggesting this example.

Merton (1969)), whereas a household with high risk-aversion may have a higher share in risky assets than what is optimal in the long-term because it has not yet rebalanced. In the market downturn between 1999-2002, the household with high risk-aversion then became a down-switcher, and both households rebalanced their portfolios. Systematic differences in rebalancing may then have occurred, not because of hedging motives but rather because of differences in risk-aversion, which lead to both heterogeneous rebalancing and switching decisions. We stress that this effect is driven by a friction that leads to similar initial portfolio holdings for households with different tastes, together with a correlation between job switching and taste.

To address the type of effects discussed above, we first note that wage volatility has the same effect on portfolio holdings of risky assets in both the cross-section and the time-series (i.e., with and without fixed effects), which suggests that hedging is indeed present. By definition, any source of endogeneity that is not reflected in the agents' initial portfolio holdings does not contaminate our first cross-sectional estimation. The results for the 1999 cross-section are nearly identical to the ones for 2002 that we reported in Table 5. Wage volatility also has a significantly negative effect on the risky share in the initial cross-section, so an omitted variable would therefore have to drive this effect of wage volatility both with and without fixed effect adjustments, raising the bar for such an alternative explanation.

Furthermore, we can verify that previous behavior of switchers and non-switchers in the years leading up to our test does not differ. We do not have information on the households' portfolio holdings prior to 1999 but we observe whether they also switched industries between 1996 and 1998, a period during which the market conditions were quite different from the recession of the early 2000s, notably with a large market run-up. Thus, with this additional control, to fail to detect differences between heterogenous households, not only would their behavior have to be similar during the switching period, but also in the years before, during radically different market conditions.

Presumably, if the decision to switch to a riskier or safer industry during a recession depends on the type of an individual, then her type should also affect her decision to switch industries in a good economy. We measure this effect by computing, for each household, the change in their wage volatility from 1996 to 1998 using the same method as for their 1999-2002 volatility change. Our analysis is twofold. First, we study whether we can infer anything from the job switching behavior of households in the years 1996-1998 (after having controlled for their portfolio holdings in 1999). Then, we test whether their change in wage volatility during these early years has any effect on their portfolio rebalancing decisions between 1999 and 2002.

[Table 10 about here.]

We find that while these 1996-98 changes in wage volatility do seem to pick up some additional unobserved heterogeneity in preferences, controlling for them does not affect our main results. When comparing the job switching behavior of households between the 1996-98 and 1999-2002 periods, we find some evidence that there may be some unobserved heterogeneity behind the job switching decision that we are not fully capturing by conditioning on the households' portfolio holdings in 1999. In Table 10 we report the likelihood of "up" and "down" switches between 1996 and 1998 for our three types of households (i.e. our up-, down-, and non-switchers between 1999 and 2002), which we also split into three terciles to control for their portfolio share of risky assets held in 1999. Across all three terciles, the households that switched to safer industries between 1999 and 2002 (i.e. the down-switchers) were the most likely to switch to the riskier industries in the previous "boom" period. Likewise, the households that switched to the riskier industries between 1999 and 2002 were the most likely to switch to the safer industries in the previous period. This evidence suggests that if this switching behavior depends on the households' type, then observing the households' change in wage volatility between 1996 and 1998 will tell us something about their type that is unrelated to hedging during the 1999-2002 period.

[Table 11 about here.]

We test whether adding the change in the households' wage volatility between 1996 and 1998 as another control variable in Eq. (5) affects our main results. In Table 11 we report the results of two additional regressions. In the first estimation, we simply add this new variable as another control in Eq. (5). In the second and more conservative estimation, we begin by regressing the same dependent variable $\Delta \widehat{^{a}w_{s,02}}$ on this variable to pick up anything that has to do with it. Then, we take the residuals from this regression and regress them on the change in wage volatility between 1999 and 2002 and all the other control variables.

In both estimations, the effect of the change in wage volatility between 1996 and 1998 is not statistically significant. Moreover, the effects of all the other variables including the change in wage volatility between 1999 and 2002 are nearly identical to those in Table 8.

These robustness checks suggest that any potential endogeneity that is not reflected in the households' initial portfolio shares in 1999 is unlikely to bias our results.

6 Conclusion

The literature on labor income risk and the levels of portfolio holdings has led to mixed results. On the one hand, there is evidence that agents hedge human capital risk (Guiso, Jappelli, and Terlizzese, 1996; Vissing-Jorgensen, 2002). On the other hand, at the individual stock holdings level, households tend to own stocks that are closely related to their labor income (Massa and Simonov, 2006).

In this paper we take advantage of a unique Swedish panel dataset and provide a new approach to this issue by focusing on the households that switched industries between 1999 and 2002. We study the effect of their industry change — in particular the effect of changes in their wage volatility — on their portfolio holdings of risky assets. We find that households do hedge labor income risk and that the effect is economically significant. A household that moves from the lowest to the highest wage volatility industry decreases its exposure to risky assets by risky by 35%.

Our results are therefore in line with the findings of Guiso, Jappelli, and Terlizzese (1996) and Vissing-Jorgensen (2002). Our results are also, however, consistent with those of Massa and Simonov (2006), since we do not find consistent cross-sectional evidence of hedging. Our overall conclusion is therefore that individual agents hedge labor income risk, but that this hedging effect is more difficult to observe in the cross-section because of the presence of "taste" heterogeneity among agents. This result also has asset pricing implications. If the strength of these two offsetting effects vary with the business cycle, then it is not surprising that the unconditional CAPM with human capital fails (as documented by Fama and Schwert, 1977) whereas the conditional CAPM with human capital is successful in explaining the cross section of expected returns (as documented by Jagannathan and Wang, 1996).

References

Benzoni, L., P. Collin-Dufresne, and R. Goldstein (2007): "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets are Cointegrated," *The Journal of Finance*, 62, 2123–2167.

Berk, J., and J. Walden (2010): "Limited Capital Market Participation and Human Capital Risk," NBER Working Paper 15709.

Bodie, Z., R. Merton, and W. Samuelson (1992): "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model," *Journal of Economic Dynamics and Control*, 16, 427–449.

Calvet, L., J. Campbell, and P. Sodini (2009): "Fight or Flight? Portfolio Rebalancing by Individual Investors," *The Quarterly Journal of Economics*, 124(1), 301–348.

Campbell, J. (1996): "Understanding Risk and Return," *Journal of Political Economy*, 104(2), 298–345.

Campbell, J., J. Cocco, F. Gomes, and P. Maenhout (2001): "Investing Retirement Wealth: A Life-Cycle Model," in *Risk Aspects of Investment-Based Social Security Reform*, ed. by J. Campbell, and M. Feldstein, pp. 439–482. University of Chicago Press.

Cocco, J. (2005): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*, 18(2), 535–567.

Danthine, J.-P., and J. Donaldson (2002): "Labour Relations and Asset Returns," *The Review of Economic Studies*, 69(1), 41–64.

Fama, E., and W. Schwert (1977): "Human Capital and Capital Market Equilibrium," *Journal of Financial Economics*, 4(1), 95–125.

Gakidis, H. E. (1998): "Stocks for the Old? Earnings Uncertainty and the Life-Cycle Portfolio Choice," Ph.D. thesis, M.I.T.

Georgarakos, D., and G. Pasini (2009): "Trust, Sociability and Stock Market Participation," Working paper NETSPAR No.04/2009-015.

Guiso, L., T. Jappelli, and D. Terlizzese (1996): "Income Risk, Borrowing Constraints, and Portfolio Choice," *The American Economic Review*, 86(1), 158–172.

Heaton, J., and D. Lucas (1996): "Evaluating the Effects of incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104(3), 443–487.

——— (2000): "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *The Journal of Finance*, 55(3), 1163–1198.

Heckman, J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47, 153–161.

Jagannathan, R., and Z. Wang (1996): "The Conditional CAPM and the Cross-Section of Expected Returns," *The Journal of Finance*, 51, 3–53.

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth and Expected Stock Returns," *The Journal of Finance*, 55(3), 815–849.

Lustig, H., and S. Van Nieuwerbugh (2008): "The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street," *Review of Financial Studies*, 21(5), 2097–2137.

Massa, M., and A. Simonov (2006): "Hedging, Familiarity, and Portfolio Choice," *The Review of Financial Studies*, 19(2), 633–685.

Mayers, D. (1973): "Nonmarketable assets and the determination of capital asset prices in the absence of a riskless asset," *The Journal of Business*, 46(2), 258–267.

Merton, R. (1969): "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 247–257.

Palacios, M. (2010): "Human Capital as an Asset Class: Implications from a General Equilibrium Model," Working Paper.

Palacios-Huerta, I. (2003): "The Robustness of the Conditional CAPM with Human Capital," *The Journal of Financial Econometrics*, 1(2), 272–289.

Parlour, C., and J. Walden (2011): "General Equilibrium Returns to Human and Investment Capital under Moral Hazard," *Review of Economic Studies*, 78, 394–428.

Qin, J. (2002): "Human-capital-adjusted capital asset pricing model," *The Japanese Economic Review*, 53, 182–198.

Santos, T., and P. Veronesi (2006): "Labor Income and Predictable Stock Returns," *The Review of Financial Studies*, 19, 1–44.

Sunden, A. (2006): "The Swedish Experience with Pension Reform," *Oxford Review of Economic Policy*, 22(1), 133–148.

Vissing-Jorgensen, A. (2002): "Toward an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," Working Paper.

Figures

Figure 1: Total, active, and passive changes between 1999 and 2002

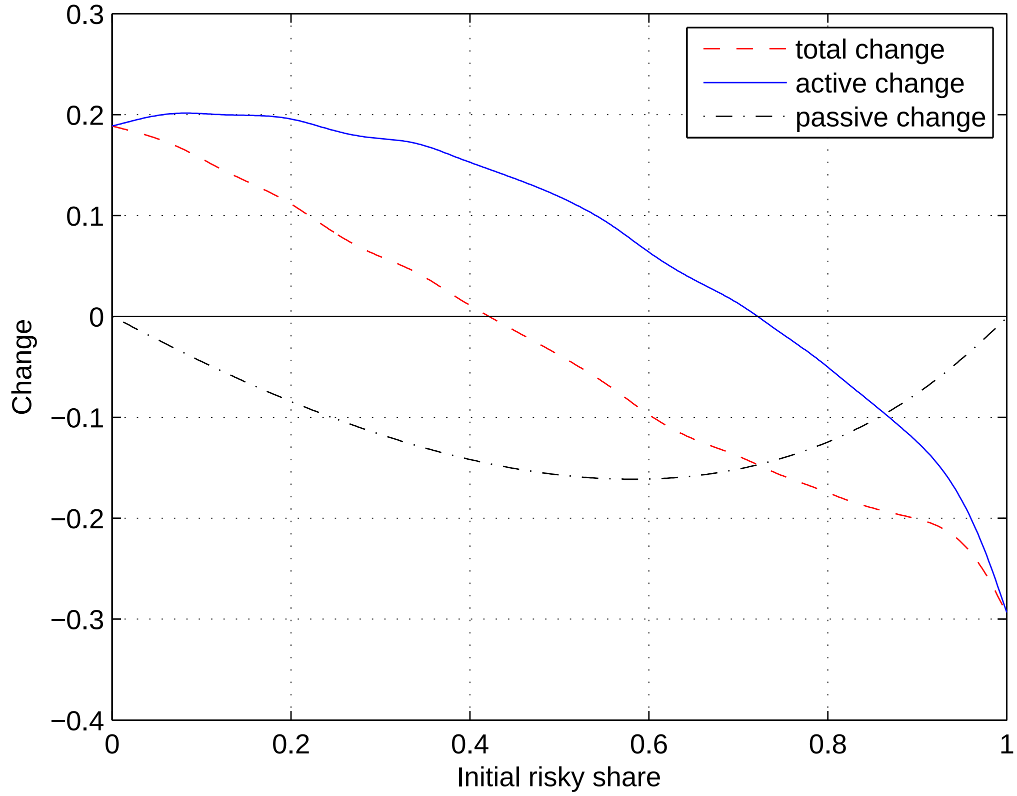

Decomposition of the predicted values of households' change in the risky share between 1999 and 2002, $\Delta w_{h,02}$, as a function of their initial risky share, $w_{h,99}$, into a passive change, $\Delta^p w_{h,02}$, and an active change $\Delta^a w_{h,02}$. The risky share is defined as the percentage of financial assets held in stocks and risky mutual funds (financial assets are the sum of checking and savings accounts, money-market funds, bond-only mutual funds, stocks, and risky mutual funds.). To filter out noise and get a smooth approximation of total changes $\Delta w_{h,02}$ as a function of $w_{h,99}$, the total changes are projected (regressed), using three cubic splines. We then calculate the passive change as the change in the risky share conditional on no portfolio rebalancing between 1999 and 2002 (using stock and government bond index returns as proxies). Finally, the active change is defined as the difference between the projected total change and the passive change.

Figure 2: Construction of our dependent variable in the analysis of switchers

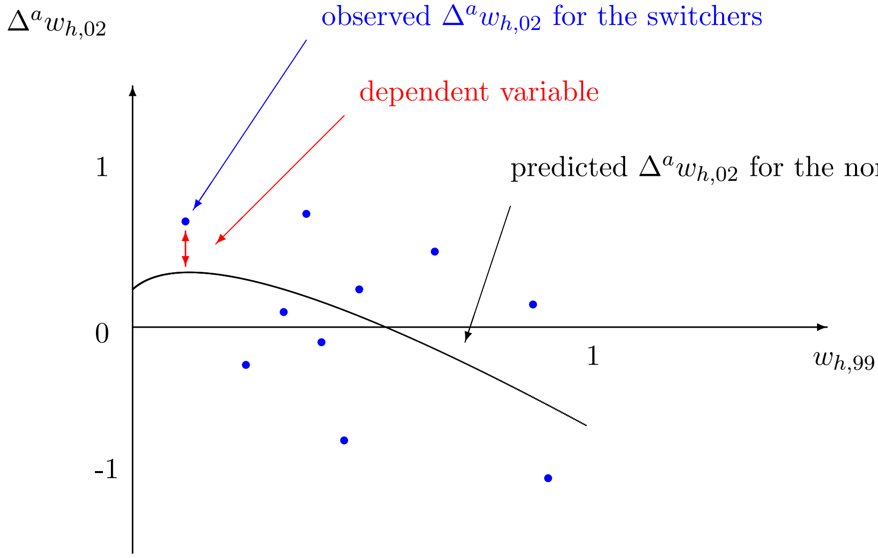

In this graph we explain how we derive our dependent variable for our analysis of switchers. The solid line (black) represents the predicted values of the active change in the risky share for non-switcher households. The active change corresponds to a household's change in its portfolio share of risky assets between 1999 and 2002 that comes from portfolio rebalancing decisions. These values come from a cubic-spline estimation with three degrees of freedom. They are plotted against the initial risky share in 1999. The dots (blue) represent the observed active changes in the risky shares for the switcher households (i.e. households where at least one adult member switched industries between 1999 and 2002). Our dependent variable $\Delta^a \widehat{w_{s,02}}$ is defined as the double-arrow vertical vector (red).

Figure 3: Fitted active changes for up-, down-, and non-switchers

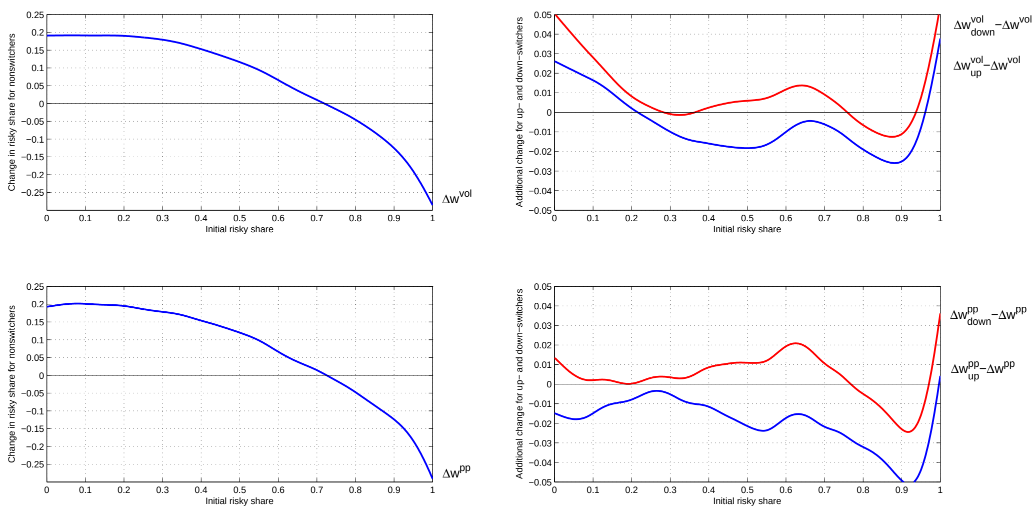

In the two left quadrants we report the fitted active changes in the risky share between 1999 and 2002 for non-switcher households. They are plotted against the initial risky share in 1999. The active change corresponds to a household's change in its portfolio share of risky assets between 1999 and 2002 that comes from portfolio rebalancing decisions. In the top-left quadrant, non-switchers are defined as households where both individuals do not switch industries between 1999 and 2002. In the bottom left quadrant, non-switchers are defined as households where both individuals do not switch between the public and the private sectors. In the two right quadrants, we report the fitted values of the *additional* active changes in the risky share for up-switchers and down-switchers (that is, relative to the predicted change of the non-switchers given the same initial risky share). In the top right quadrant, up-switchers (down-switchers) are defined as switchers that experience an increase (decrease) in wage volatility. In the bottom right quadrant, up-switchers (down-switchers) are defined as switchers that switch from the public (private) to the private (public) sector. In each of the right quadrants, the top (red) line corresponds to the down-switchers, and the bottom (blue) line corresponds to the up-switchers.

Tables

Table 1: Household characteristics in 1999

	All Households				Switchers	
Variable	Mean	Std Dev	Min	Max	Mean	Std Dev
Demographics						
age	43.82	9.43	18	62	41.19	9.41
nordic	.98	.14	0	1	.98	.12
number of children	1.2	1.15	0	11	1.3	1.14
Civil Status						
married	.62	.48	0	1	.6	.49
single	.16	.37	0	1	.18	.38
Education						
student	.07	.26	0	1	.1	.29
college degree	.48	.5	0	1	.52	.5
business degree	.15	.36	0	1	.2	.4
Population Density						
high	.34	.47	0	1	.41	.49
medium	.55	.50	0	1	.51	.50
low	.11	.32	0	1	.09	.28
Labor income						
family income	326.73	170.58	1.8	3209.81	325.93	172.80
is unemployed	.16	.36	0	1	.17	.38
is retired	.08	.28	0	1	.07	.25
Housing and Wealth						
homeowner	.9	.33	0	1	.89	.31
net worth	1,098.42	2,037.08	1.02	157,096.07	1,018.90	1,482.36
fin wealth	445.01	1156.74	3	77,619.77	400.34	860.92

We report the summary statistics in 1999 for our population of 73,456 households, which includes 3,815 switchers. Switcher households are by definition households where at least one adult member switched industries between 1999 and 2002. All monetary values are defined in thousands of Swedish kronor (SEK). The SEK/USD exchange rate on December 30th, 1999 was 8.52. Exact definitions of the reported variables are described in the Appendix.

Table 2: Participation rates and portfolio shares in 1999 and 2002

	All Households			Switchers		
Variable	Mean	Std Dev	Part.	Mean	Std Dev	Part.
Panel A: 1999						
risky assets	.62	.32	.90	.62	.32	.90
stocks	.28	.29	.48	.3	.29	.49
mutual funds	.52	.32	.81	.51	.33	.80
Panel B: 2002						
risky assets	.52	.31	.94	.52	.31	.95
stocks	.19	.23	.62	.2	.23	.63
mutual funds	.42	.29	.88	.42	.29	.87

Portfolio shares are conditional on participation. The category "stocks" consists of all non-retirement directly-held stocks. The category "mutual funds" consists of all mutual funds that are fully or partially invested in stocks. The category "risky assets" is the sum of stocks and mutual funds. These portfolio shares are out of financial wealth, which consists of all checking and savings accounts, money-market funds, bond-only mutual funds, stocks, and risky mutual funds. The data set has 73,456 observations, which includes 3,815 switchers. Switcher households are by definition households where at least one adult member switched industries between 1999 and 2002.

Table 3: Participation rates and portfolio shares in 2001: Sweden vs. USA

	LINDA			SCF I			SCF II		
Variable	Mean	Std Dev	Part.	Mean	Std Dev	Part.	Mean	Std Dev	Part.
risky assets	.57	.3	.94						
stocks	.22	.24	.59	.40	.31	.41	.29	.26	.41
mutual funds	.46	.29	.88	.30	.26	.30	.19	.19	.30

The first column (LINDA) refers to observations from the LINDA dataset in 2001. The data set has 73,456 observations. The other two columns refer to observations from the 2001 Survey of Consumer Finances (SCF). In the second column (SCF I), we adjust the SCF portfolios so that they are comparable to the ones computed in LINDA. In particular, we exclude retirement assets and we sum up the holdings of pure-equity and mixed mutual funds. The third column (SCF II) reflects more closely the true risky portfolio shares in the USA. The holdings of mixed mutual funds are halved to reflect the fact that they are not fully invested in stocks, and the retirement assets are included.

Table 4: Rankings of industries by their levels of wage volatility

SNI	Description	Est.
	Bottom 10	
371	Recycling of metal waste and scrap	.07
271	Manufacturing of iron and steel	.08
131	Mining of iron and ores	.08
173	Finishing of textile	.09
272	Manufacturing and casting of iron tubes	.09
172	Weaving of cotton	.09
365	Manufacturing of games and toys	.09
274	Production of precious metals, copper	.10
403	Steam and hot water supply	.10
175	Manufacturing of ribbons, curtains	.10
	Top 10	
21	Renting of household goods	.21
13	Mixed farming	.21
722	Publishing of software	.22
741	Legal representation activities	.23
672	Other finance activities	.24
744	Advertising	.24
924	Other Entertainment	.25
553	Restaurants	.26
921	Motion picture and video production	.26
671	Finance administration, fund management	.30

Wage volatility is defined as the average volatility of annual returns to real disposable income across all individuals within a 3-digit SNI code who have stayed in the same 5-digit SNI code for at least 5 consecutive years between 1993 and 2003. The rankings are based on 191 observations.

Table 5: Effects of wage volatility on the risky share in 2002.

	(1)		(2)		(3)	
Variable	Est.	t-stat	Est.	t-stat	Est.	t-stat
wage vol.	-.12	-2.51	-.22	-5.69	-.14	-2.86
wage vol. same ind.	.001	.02	.004	.16	-.004	-.15
public	.014	4.26			.011	2.69
private	-.004	-1.62			-.007	-2.12
Intercept	1.102	23.41	.707	8.70	.685	7.6
age	-.002	-1.95	-.005	-3.15	-.005	-3.14
$(\text{age})^2$	.015	1.05	.03	1.77	.031	1.89
nordic	.018	1.82	.047	4.79	.046	4.57
has emigrated	-.014	-.99	-.028	-1.56	-.029	-2.12
no. children	.024	17.2	.027	17.37	.027	16.79
single parent	.014	2.78	.028	4.69	.026	4.38
married	-.007	-2.31	-.006	-2.01	-.006	-1.76
student	.01	1.88	.017	2.81	.017	2.48
college degree	.012	4.37	.025	7.1	.021	5.99
business major	.012	3.53	.01	2.77	.012	3.68
high pop. density	.001	.03	-.017	-3.77	-.016	-3.02
medium pop. density	.033	8.77	.03	.003	.031	7.78
family income	-.04	-11.16	-.017	-3.17	-.015	-2.44
is unemployed	-.007	-.19	-.004	-1.16	-.003	-.087
is retired	-.017	-4.86	-.017	-4.5	-.017	-4.76
debt / income 99	.003	4.14	.003	3.82	.003	4.7
homeowner	.013	2.57	.018	3.27	.019	3.11
house / networth 99	-.016	-8.5	-.019	-9.18	-.019	-8.8
net worth 99	-.004	-3.34	.004	2.37	.004	2.41
lambda			.292	4.31	.292	5.49
No. Obs	69,097		69,097		69,097	
F	67					
R-sq	.022					
Chi-sq			1,782		2,086	

We report second-stage estimates of the portfolio holdings of risky assets (stocks and risky mutual funds) as a percentage of financial assets (e.g. the "risky share") in 2002 (Eq. (4) in the text). Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. The sample is restricted to households with positive holdings. Four separate OLS regressions are run. In columns 2 to 4, lambda is the inverse mills ratio from the first stage estimation of the decision to participate in the stock market (Eq. (3) in the text). We report the bootstrapped t-stats. In column 1, where we do not control for lambda, we report the heteroskedasticity-consistent t-stats. All the goodness-of-fit F and Chi-sq tests are statistically significant at the 1% level. Other explanatory variables are described in the Appendix.

Table 6: Effects of wage volatility on the portfolio shares of stocks and mutual funds in 2002

	stocks		mutual funds	
Variable	Est.	t-stat	Est.	t-stat
wage vol.	.238	4.89	-.37	-8.08
wage vol. same ind.	-.002	-.001	-.07	.16
public	-.009	-2.73	.02	5.05
private	.012	4.41	-.019	-6
X variables	yes	yes	yes	yes
lambda	.607	11.59	-.315	-5.14
No. Obs	69,097		69,097	
Chi-sq	3,048		7,885	

We report second-stage estimates of portfolio holdings of directly-held stocks and mutual funds as a percentage of financial assets in 2002. The sample is restricted to households with positive holdings. Two separate OLS regressions are run. The dependent variables are the share of directly-held stocks over financial wealth and the share of risky mutual funds (equity and mixed) over financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. Lambda is the inverse mills ratio from the first stage estimation of from the first stage estimation of the decision to participate in the stock market. We report the t-statistics for the bootstrapped standard errors. All the goodness-of-fit Chi-sq tests are statistically significant at the 1% level. Other explanatory variables in the vector $X_{h,02}$ (see Tablẽreftab:Levels1 are included but we do not report the results.

Table 7: Sign tests on the active changes in the risky share between 1999 and 2002 for switchers

	wage vol.		public-private	
Variable	up	down	up	down
Sign	-	+	-	+
Est.	-8,770***	15,227***	-20,035***	17,077***

We test that the predicted values from the splines for the up- and down- switchers that are shown in Fig. 3 are different from the predicted values from the splines for the non-switchers. There are 59,025 observations for the wage volatility measure and 59,047 observations for the public-private measure. *** indicates statistical significance at the 1% level.

Table 8: Effect of changes in wage volatility for switchers on the active change in their risky share between 1999 and 2002

	(1)		(2)		(3)		(4)		(5)		(6)	
Variable	Est.	t-stat	Est.	t-stat	Est.	t-stat	Est.	t-stat	Est.	t-stat	Est.	t-stat
Intercept	.005	1.09	.005	1.09	.005	1.08	.004	.6	.003	.47		
Δ wage vol	-.312	-2.25	-.323	-2.34	-.31	-2.24						
Δ wage vol - up							-.29	-1.23				
Δ wage vol - down							-.36	-1.4				
Δ sign(wage vol) $\cdot$ (wage vol)^2									-.482	-2.28		
to private											-.016	-2.84
to public											.008	1.39
has emigrated	-.058	-.87										
Δ no. children	-.022	-2.66	.023	2.48	.019	2.09	.023	2.24	.022	2.2	.022	2.74
has graduated	-.014	-.71										
low to high density	-.013	-1.21										
has retired	-.037	-1.92	-.035	-1.75	-.036	-1.91	-.035	-1.85	-.035	-1.83	-.036	-2.42
Δ family income	-.074	-4	-.074	-4.03	-.08	-4.35	-.074	-4.02	-.073	-3.99	-.045	-3.55
found a job	.024	1.33										
lost a job	-.016	.85										
Δ debt / income	-.007	-1.68	-.007	-1.83	-.01	-2.21	-.008	-1.83	-.007	-2	-.001	-1.33
bought house	.016	.47										
sold house	-.14	-3.49	-.147	-3.66	-.12	-3.01	-.147	-3.67	-.147	-3.67	-.16	-5.01
Δ house / networth	.021	2.08	.022	2.27	.038	4.54	.022	2.26	.022	2.3	.012	1.49
Δ net worth	-.022	-2.66	-.02	-3			-.02	-2.6	-.02	-2.58	-.032	-5.4
No. Obs	2,565		2,565		2,565		2,565		2,565		3,890	
F	4.67		7.95		7.78		7.07		7.94		12.53	
Adj R-sq	.021		.021		.018		.021		.021		.026	

We report estimates of the change in the portfolio share of risky assets between 1999 and 2002 (Eq. (5) in the text). The sample is restricted to households with positive holdings in 1999. Six separate OLS regressions are run. The dependent variable is the difference between the observed active change in the risky share for switchers and the predicted active change in the risky share for non-switchers (between 1999 and 2002) given the same initial risky share in 1999. See Fig. 2 for a visualization of the construction. We report the t-statistics for the heteroskedasticity-robust standard errors. All the goodness-of-fit F tests are statistically significant at the 1% level. Other explanatory variables are described in the Appendix.

Table 9: Effects of changes in wage volatility on changes in the the portfolio shares of stocks and mutual funds between 1999 and 2002

	stocks		mutual funds	
Variable	Est.	t-stat	Est.	t-stat
Intercept	.004	.82	0.000	-.12
Δ wage vol.	.007	.05	-.297	-2.12
Δ no. children	.001	.13	.015	1.54
has retired	.007	.39	.030	-1.68
Δ family income	-.009	-.57	-.082	-4.3
Δ debt / income	-.002	-.8	-.009	-2.35
sold house	-.097	-2.22	-.130	-3.25
Δ house / networth	.013	1.2	.021	2.07
Δ net worth	-.006	-.82	-.019	-2.81
No. Obs	1,346		2,294	
F	1.24		7.1	
Adj R-sq	.001		.021	

We report estimates of the change in the portfolio shares of directly-held stocks and mutual funds between 1999 and 2002. The sample is restricted to households with positive holdings in 1999. Two separate OLS regressions are run. In column 1 (2), the dependent variable is the difference between the observed active change in the share of directly-held stocks (risky mutual funds) for switchers and the predicted active change in the share of directly held stocks (risky mutual funds) for non-switchers (between 1999 and 2002) given the same initial share of directly held stocks (risky mutual funds) in 1999. See Fig. 2 for a visualization of the construction. We report the t-statistics for the heteroskedasticity-robust standard errors. Only the F statistic in the second estimation is statistically significant at the 1% level. Other explanatory variables are described in the Appendix.

Table 10: Probability of up- and down-switches between 1996-98 for the various types of households.

	$w_{h,99}$ low tercile			$w_{h,99}$ medium tercile			$w_{h,99}$ high tercile		
	up	non	down	up	non	down	up	non	down
prob up switch 96-98	.051	.064	.255	.038	.061	.0222	.058	.06	.277
prob down switch 96-98	.15	.061	.064	.147	.057	.033	.165	.0579	.0477

We report estimates of the probability of up- and down- switches between 1996 and 1998 for various groups of households. By definition, an up- (down-) switch corresponds to an industry switch that leads to an increase (decrease) in wage volatility for the household. First, we split all households into three terciles based on their risky share in 1999 ($w_{h,99}$). Then, within each tercile, we classify households as up-, non-, or down-switchers based on their decision to switch between 1999 and 2002. The probability of an up switch in 1996-98 is computed as the fraction of households within each group that had at least one up-switch between 1996 and 1998 and more up-switches than down-switches in the event of multiple job switches.

Table 11: Effects of changes in wage volatility between 96-98 and 99-02 on changes in the portfolio shares of risky assets between 1999 and 2002

	Regular		Two-stage	
Variable	Est.	t-stat	Est.	t-stat
Intercept	.003	.69	0.000	-.12
Δ wage vol.	-.32	-2.23	-.3	-2.11
Δ no. children	.02	1.91	.02	1.9
has retired	-.032	-1.72	-.03	-1.71
Δ family income	-.07	-3.66	-.068	-3.63
Δ debt / income	-.006	-1.52	-.006	-1.51
sold house	-.15	-3.46	-.15	-3.44
Δ house / networth	.02	2.12	.021	2.12
Δ net worth	-.021	-2.92	-.02	-2.51
			first-stage	
Δ wage vol. 96-98	.23	.79	.44	1.48
No. Obs	2,456		2,456	
F	6.57		7.02	
Adj R-sq	.02		.019	

We report estimates of the change in the portfolio share of risky assets between 1999 and 2002. The sample is restricted to households with positive holdings in 1999. Two separate OLS regressions are run. In column 1, we conduct the same regression as in column 2 of Table 8 but adding an additional control variable: the households' change in wage volatility between 1996 and 1998 (which is computed the same way as the one between 1999 and 2002). The dependent variable is the difference between the observed active change in the risky share for switchers and the predicted active change in the risky share for non-switchers (between 1999 and 2002) given the same initial risky share in 1999. In column 2, we perform a two-stage analysis, where in the first stage we regress the same dependent variable on the change in volatility between 1996 and 1997. Then, in the second stage, we regress the residuals from the first-stage regression on the same variables as in Table 8. We report the t-statistics for the heteroskedasticity-robust standard errors. All the goodness-of-fit F tests are statistically significant at the 1% level. Other explanatory variables are described in the Appendix.