

For Better, For Worse: Intra-Household Risk-Sharing over the Business Cycle

Stephen H. Shore¹²³

First Draft: March 6, 2006

This Draft: February 20, 2008

¹Johns Hopkins University, Department of Economics, shore@jhu.edu, 458 Mergenthaler Hall, Baltimore, MD 21218, 410-516-5564.

²JEL Classification: E32 - Business Fluctuations, Cycles; D13 - Household Production and Intrahousehold Allocation; D31 - Personal Income, Wealth, and Their Distributions.

³I thank Michael Pan and Seth Zeleznik for excellent research assistance. I thank Amalia Miller, Todd Sinai, Amir Yaron, and seminar participants at the Wharton School, the University of Washington, and the 2007 Society of Labor Economists (SOLE) Annual Meeting for helpful comments. Two anonymous referees and editor Daron Acemoglu also provided helpful comments. All errors remain my own. I gratefully acknowledge the financial support of the National Institutes of Health - National Institute on Aging, Grant number P30 AG12836, the Boettner Center for Pensions and Retirement Security at the University of Pennsylvania, and National Institutes of Health - National Institute of Child Health and Development Population Research Infrastructure Program R24 HD-044964, all at the University of Pennsylvania.

Abstract

Marriage allows couples to diversify labor income risks and dynamically coordinate labor supply decisions in response to shocks. For instance, when one spouse loses a job, the other can work more. This paper argues that these risk-sharing benefits of marriage are countercyclical; husbands' and wives' income changes are more positively correlated when the economy is growing rapidly. As a result, while individuals face substantially more idiosyncratic labor income risk in bad times than in good, households do not. Improved intra-household risk-sharing undoes for couples the increased riskiness that comes to individuals in bad times.

Since panel datasets are typically too short to capture more than a few business cycles, I exploit variation in the cross-sectional covariance of husbands' and wives' incomes to infer the covariance of couples' past income changes. Couples who have been married through periods of greater economic expansion have more positively correlated labor incomes in the cross-section, after controlling for a time trend in assortative mating. This implies that the correlation of couples' labor income changes falls by roughly 20 percent (e.g., from 10 percent to -10 percent) from booms to busts.

1 Introduction

Marriage helps couples manage labor income risk. First, pairing allows couples to diversify idiosyncratic shocks. As long as a wife's exogenous income shocks are not perfectly correlated with her husband's shocks, pooling can reduce risk for both. Second, a couple can coordinate work choices dynamically in response to exogenous shocks to comparative advantage; if one spouse's market wage falls, the other can choose to work more. By adjusting relative labor supply, a couple reduces the impact of shocks. Both types of risk reduction will be reflected in the covariance of couples' income changes.

Previous research has found that there is more idiosyncratic risk in recessions than booms (Storesletten, Telmer, and Yaron, 2004). This paper seeks to understand the extent to which this increased risk is crowded out (or crowded in) by improved (or reduced) intra-household risk-sharing. How does the covariance of couples' income changes vary over the business cycle? The answer is not clear *a priori*. The benefits of diversification and dynamic coordination are most valuable in recessions. Not only are shocks larger, but households are more risk-averse with respect to a given shock.¹ In recessions, this might lead couples to avoid jobs with correlated shocks (e.g., sharing an employer), or for spouses to respond more strongly *ex post* to their partners' shocks. On the other hand, dynamic coordination is likely to be less effective in recessions. Adjusting labor supply in response to a shock will be more difficult in recessions, when many other unemployed workers are looking for work. Since the risk-sharing benefits of marriage are likely to be more needed but less effective in recessions, the net prediction is unclear.

Understanding how these patterns vary over the business cycle is important for two reasons. From a policy perspective, the generosity of policy interventions aimed at reducing risk (e.g., unemployment insurance) is generally countercyclical. Business-cycle variation in policy is informed by the belief that there is more need for social insurance programs in recessions. This paper argues that this may be true for individuals but not for married

¹With consumption commitments or habit preferences, relative risk aversion will be higher in bad states.

couples. From the perspective of economic theory, this paper will document that intra-household decisions respond to the macroeconomic environment; improved household risk-sharing undoes the increased riskiness that comes to individuals in bad times.

The most direct test for business-cycle variation in the covariance of couples' income changes would use panel data to measure that covariance at various points in time, comparing estimates from booms to those from busts. This approach is used by Parker and Skoufias (2004) and by Juhn and Potter (2007), who estimate the added worker effect – the increase in a wife's labor supply following her husband's unemployment (Long, 1953; Lundberg, 1985; Cullen and Gruber, 2000) – for panels of Mexican and U.S. households, respectively. This approach is complicated by the short length of most panel data sets, which span only a few business cycles, and by the possibility that intra-household coordination may lead or lag the business cycle. As a result, it is difficult to differentiate business-cycle effects from time-trends or variation in other factors.

To overcome this problem, I develop a technique to infer the covariance of couples' past income changes from the cross-sectional covariance of couples' subsequent income levels. This paper builds on Deaton and Paxson (1994), who use *life-cycle variation in the cross-sectional variance* of consumption to measure the magnitude of shocks to permanent income. For example, they estimate the magnitude of permanent shocks between the ages of 25 and 35 by comparing the cross-sectional variance of consumption at age 25 with the cross-sectional variance for the same cohort at age 35. The larger the variance of permanent shocks, the more the cross-sectional variance will rise with age. In this paper, I use *marriage-cycle variation in the cross-sectional covariance* of couples' incomes to measure the covariance of husbands' and wives' permanent income changes. The larger the covariance of couples' permanent shocks, the more the cross-sectional covariance will rise with the number of years of marriage. I find that the cross-sectional covariance of couples' incomes falls with number of years of marriage for roughly the first 10 years of marriage, and then rises with years of marriage after roughly 15 years of marriage. This suggests that couples' permanent shocks

are negatively correlated early in marriage but positively correlated later in marriage.

Storesletten et al. build on Deaton and Paxson by exploiting differences in the cross-sectional variance of cohorts with different macroeconomic histories. They measure cyclical variation in the magnitude of idiosyncratic income shocks by *comparing the cross-sectional variance of income for cohorts of individuals born in different years*, and therefore who had worked through more or fewer recessions. In this paper, I *compare the cross-sectional covariance of husbands' and wives' incomes for cohorts of couples wed in different years*, and therefore whose marriages spanned periods of greater or lesser economic expansion. If the correlation between husbands' and wives' permanent income changes is procyclical, then a cohort composed of couples whose marriages spanned a period of greater economic expansion will have a more positive cross-sectional covariance of husbands' and wives' incomes. Since this method infers the covariance of past shocks from the current cross-sectional covariance, it exploits the large variation in the year of marriage, even in data sets that span few years. The major drawback of this approach is that it can infer only the covariance of past *permanent* shocks from the subsequent cross-sectional covariance; any variation in past transitory shocks will not appear in the subsequent cross-section.²

I construct a panel of cohorts – cohort-year observations include data from a given year about a cohort of couples that all wed in the same year – from the Panel Study of Income Dynamics (PSID). As is standard in the literature on income processes, I remove the predictable (to the econometrician) component of log labor income to calculate “excess” log income. I then calculate the cross-sectional covariance of husbands' and wives' excess log incomes for each cohort in each year. I find that cohorts of couples whose marriages spanned periods of greater economic expansion have more positive cross-sectional covariances. In cohorts of couples who had been married through more good times, high-earning husbands tend to be married to high-earning wives, and low-earning husbands tend to be married to low-earning wives; in cohorts of couples who had been married through relatively bad times,

²Stephens (2002) documents that the added worker effect is persistent. This paper documents cyclical variation in the degree of this persistent dynamic coordination.

high-earning husbands tend to be married to low-earning wives and *vice versa*. The effect is both quite large and statistically significant. Given a simple income process for husbands and wives, the results imply that increasing annual GDP growth by 1 percent increases the correlation of couples' permanent income changes by roughly 4 percent (e.g., from 1 percent to 5 percent). If GDP growth rises by 5 percent from a bust to a boom, this implies business-cycle variation of 20 percent in the correlation of couples' shocks to permanent income (e.g., from -10 percent to 10 percent).

I find that the variance of innovations to individual income is countercyclical. After controlling for the time trend in assortative mating, the covariance of couples innovations to income is procyclical. At the household level, increased individual risk and improved intra-household risk-sharing roughly net out, so that the variance of shocks to the household's (husband's plus wife's) labor income is acyclical. In other words, marriage undoes most if not all of the increased riskiness that comes to individuals in recessions.

The remainder of the paper is organized as follows: Section 2 details the data used; Section 3 develops techniques to infer the covariance of couples' income shocks from cross-sectional data and applies these techniques to estimate business-cycle variation in the covariance of couples' incomes; Section 4 shows that this covariance moves procyclically; Section 5 considers selection problems and interprets the results; and Section 6 concludes.

2 Data

Data are drawn from the Panel Study of Income Dynamics (PSID), a nationally representative panel of U.S. households that has tracked families since 1968. It includes data on households, including the education, income, employment status, and age of husbands and wives.

For Storesletten et al. as well as Deaton and Paxson, data are aggregated to cohort-year moments. Cohorts in those papers are age-based, where an age cohort refers to households

with household heads all born in a given year. Since this paper examines couples' income dynamics, I use marriage-based cohorts. A marriage cohort refers to couples who all wed in a given year. Throughout this paper, "age cohort" and "marriage cohort" will be used to differentiate these two. Marriage data are obtained from the Marriage History File of the PSID, which identifies the year of marriage, the order of the marriage (e.g., first marriage, second marriage), and the year and reason for any marriage dissolution (e.g., death or divorce). This information is available retrospectively, so that a marriage is included in the data even if it took place before data were collected.

The smallest unit of analysis in Storesletten et al. and Deaton and Paxson is the household and the variables of interest are the household's income or consumption. Since this paper aims to examine intra-household dynamics, the relevant unit for marriage cohorts is the couple and the variables of interest are the wife's income and the husband's income. I restrict the sample to married couples who have been married for between zero and 35 years, and to marriages that are the husband's first. To reduce (but perhaps not eliminate) potential divorce-related sample-selection problems, I exclude any marriages that are subsequently shown to end in divorce. Section 5.1 explores selection problems in more detail. I follow Storesletten et al. in limiting the sample to individuals with ages between 22 and 60.

Data used in this paper differ from previous research in three respects. First, data are extended through 1997 (the last year PSID data are available annually), as more data from the PSID have become available. Including this additional data has no impact on the substance of the results. Second, the definition of income is slightly more limited in this paper. Storesletten et al. use total household income, including the labor income of the husband and wife as well as transfer and other income. In this paper, it is important to decompose income into components attributable to husbands and wives. Since transfer and other income cannot be cleanly attributable to one spouse, I limit the definition of income to the labor income of husbands and wives. Including non-labor income in household income has almost no effect on household income-based results. Third and most importantly,

data here are Winsorized (Dixon, 1960). In Storesletten et al. as well as Deaton and Paxson, households with no income are excluded, as the log of zero income is not defined. This approach is sensible because the vast majority of households have strictly positive incomes. While households' total incomes are seldom zero, wives' labor incomes often are. Dropping observations with non-working spouses would not only eliminate too much data, it would eliminate much of the variation of interest. Identifying wives who re-enter the labor force because their husbands receive a negative shock is precisely the point of this paper. Furthermore, eliminating these spouses would introduce problematic sample selection, as the cross-sectional moments would be affected by changes in the composition of the female labor force. To address this concern, income data are Winsorized at the 5 percent and 95 percent levels.³ This means that any variation in income below the 5th and above the 95th percentiles is eliminated from the data, but variation from transitions into and out of very low or high income values is included. Winsorized data include all households, regardless of current labor market participation, so there is no income-based sample selection.⁴

As the measure of “excess” (unpredictable to the econometrician) log income, I use the residual from a regression of the natural log of labor income (for either the husband, the wife, or the sum for household income) on a host of regressors. I follow Storesletten et al. in including as regressors: a cubic in the age of the husband (and also the wife), the number of years of education for the husband (and also the wife), the number of family members in the household, and year fixed-effects. I also include the number of years of marriage.

³All income values below the 5th percentile or above the 95th percentile among the positive values for that variable in a given year are replaced with these 5th and 95th percentile values in that year. Results are unchanged qualitatively and minimally changed quantitatively when Winsorizing is not done for high-end values, and is done at the 2nd percentile for low-end values. Winsorizing is performed before the overlapping panels, described in the previous paragraph, are created.

⁴I also follow Storesletten et al. in building the data set from observations from three-year overlapping panels. Each panel begins in a year, and consists of observations over that year and the next two years when income data is present in all three years and when (Winsorized) income does not increase or decrease by more than a factor of 20 between any two consecutive years.

3 Cross-Sectional Moments

What does variation in cross-sectional moments tell us about past shocks? I frame empirical results with a standard income process for the excess log income of individual i at time t , y_{it} . An individual who weds in year k has an initial income of y_{ik} in year k . I assume that income then evolves in response to “permanent” shocks, ω_{it} , which are zero-mean and i.i.d. in changes. The permanent component of log income, p_{it} , then follows a random walk. Individuals also receive “transitory” shocks, ε_{it} , which are zero-mean, i.i.d. in levels, and uncorrelated with permanent shocks ($E[\omega_{is}\varepsilon_{it}] = 0$):

$$\begin{aligned} y_{it} &= p_{it} + \varepsilon_{it} \\ p_{it} &= y_{ik} + \sum_{s=k+1}^t \omega_{is}. \end{aligned} \tag{1}$$

The “permanent variance” and “transitory variance” refer to the variance of permanent and transitory shocks, $E[\omega_{it}^2] = \sigma_{it}^2$ and $E[\varepsilon_{it}^2] = \tau_{it}^2$, respectively. As in Carroll and Samwick (1997) and others, the variances from this income process measure income volatility – the size of income changes – and proxy for income risk.

3.1 Cross-Sectional Moments Over the Life Cycle

Consider the cohort of couples who wed in year k . I use $var_k(y_t)$ to denote the cross-sectional sample variance of y_{it} in year t over individuals i in cohort k . For example, $var_k(y_k)$ is the dispersion of excess log income in the first year of marriage among couples wed in year k . If everyone in cohort k has the same income process, then the cross-sectional variance of excess log income in year t can be decomposed into three parts:

$$E[var_k(y_t)] = E[var_k(y_k)] + \tau_t^2 + \sum_{s=k+1}^t \sigma_t^2. \tag{2}$$

(i) (ii) (iii)

- (i) heterogeneity in initial permanent income or ability, $var_k(y_k)$;
- (ii) the variance of the current transitory shock, τ_t^2 ;
- (iii) the sum of past variances of permanent shocks, $\sum_{s=k+1}^t \sigma_s^2$.

Over time, some individuals tend to accumulate positive permanent shocks while others tend to accumulate negative ones. As this happens, incomes spread out. Absent trends in the transitory variance, τ^2 , the rate at which the dispersion of income increases over time, $var_k(y_t) - var_k(y_{t-1})$, estimates the variance of permanent income shocks, σ_t^2 . (Deaton and Paxson, 1994)

This is shown in Figure 1, which plots the evolution of the cross-sectional variance of excess log income, $var_k(y_t)$ in equation (2), as a function of year t . When the transitory variance is constant, the slope in this figure – the degree to which the cross-sectional variance increases with the number of years of marriage – is a measure of the permanent variance, σ^2 . The figure’s two panels use different definitions of income: household income for Panel A and husband’s labor income for Panel B. Both panels present years of marriage on the x-axis and the cross-sectional covariance of “excess” log income on the y-axis.⁵ Note that the cross-sectional variance increases by approximately 0.01 per year, from 0.2 to 0.4 over a 20 year period. This implies that permanent income changes have an annual standard deviation of approximately 10 percent. The cross-sectional variance falls slightly in the first few years of marriage, suggesting that the transitory variance falls over this period. Results for all men with age-based cohorts (not just married ones with marriage-based cohorts, as is shown here) yield similar results, though the early drop in the cross-sectional variance is less pronounced.

⁵In each panel, I aggregate data for multiple cohorts to present results for a synthetic “average” cohort. I regress cohort-year cross-sectional sample variances (the cross-sectional variance in year t for the cohort of couples who wed in a given year) against calendar-year dummy variables and dummy variables for the number of years of marriage. Each panel presents a plot of the dummy variables for the number of years of marriage, with these values adjusted by the average of the calendar-year dummy variables. This approach is also used to create synthetic cohorts in Figure 2.

I now consider a joint income process for couples. Using the second sex-chromosome (X for women; Y for men) to aid memory, x_{it} denotes the excess log income of the wife in couple i in year t ; y_{it} denotes her husband's excess log income. $cov_k(x_t, y_t)$ denotes the cross-sectional sample covariance of x_{it} and y_{it} in year t over couples i in cohort k . While I assume that the income of each spouse evolves as in equation (2), couples' income shocks may be correlated. I use δ_{it} to refer to the covariance of couples' permanent shocks, $E[\omega_{xit}\omega_{yit}]$, and ζ_{it} to refer to the covariance of couples' transitory shocks, $E[\varepsilon_{xit}\varepsilon_{yit}]$.⁶

If all couples who wed in year k have the same joint income process, I can construct the covariance analog of equation (2). The cross-sectional covariance of husbands' and wives' excess log incomes in year t can be decomposed into three parts:

$$E[cov_k(x_t, y_t)] = E[cov_k(x_k, y_k)] + \zeta_t + \sum_{s=k+1}^t \delta_s. \quad (3)$$

(i) (ii) (iii)

- (i) assortative mating on income at time of marriage, $cov_k(x_k, y_k)$;
- (ii) the covariance of couples' current transitory shocks, ζ_t ;
- (iii) the sum of past covariances of couples' permanent shocks, $\sum_{s=k+1}^t \delta_s$

Over time, husbands and wives accumulate permanent shocks. The covariance of these shocks, δ_t , is a measure of diversification or coordination in marriage. If permanent shocks are perfectly positively correlated, husbands and wives accumulate the same set of shocks. High-earning wives (who tend to have accumulated positive shocks) will increasingly tend to have high-earning husbands, so the cross-sectional covariance of couples' incomes will increase. If shocks are perfectly negatively correlated, the pattern is reversed. Since wives' and husbands' shocks are of opposite sign, over time high-earning wives will increasingly tend to have low-earning husbands, and *vice versa*. The cross-sectional covariance will fall. Absent trends in transitory covariance, ζ , changes in the cross-sectional covariance of

⁶I assume that $E[\varepsilon_{xis}\omega_{yit}] = E[\varepsilon_{yis}\omega_{xit}] = 0$ for all s and t . Strictly speaking it is not possible for household income, husband's income, and wife's income all to evolve as in equation (1).

income, $cov_k(x_t, y_t) - cov_k(x_{t-1}, y_{t-1})$, measure the covariance of couples' permanent income changes, δ .

Figure 2 plots the evolution of couples' cross-sectional variances and covariances. Panels A, B, and C of Figure 2 show the evolution of the cross-sectional variance of excess log income. Income definitions vary across panels as follows: husband's labor income in Panel A (so that Panel A in Figure 2 exactly repeats Panel B in Figure 1); wife's labor income in Panel B; and, the sum of both spouse's labor income in Panel C. The cross-sectional variance of wives' excess log incomes (Panel B) is very large and falling with the number of years of marriage. This suggests that the transitory variance is also large and falling with years of marriage. Panel D of Figure 2 shows the evolution of the cross-sectional covariance of wives' and husbands' excess log incomes. The cross-sectional covariance falls with the number of years of marriage in the first decade marriage, and then rises slowly later in the marriage. This suggests that permanent income changes are negatively correlated early in marriage and positively correlated later in marriage.

Figure 3 presents two measures of the covariance of couples' permanent income changes, and shows how these evolve with the number of years of marriage. The first measure is calculated from cross-sectional data; the second from panel data. Cross-sectional estimates are obtained as in Panel D of Figure 2 with a quartic in years of marriage replacing calendar-year dummy variables in the procedure described in footnote 5. The derivative of this quartic estimates the permanent covariance and is presented with a dashed line. Panel estimates are obtained using the product of wives' one-year income changes (between years t and $t - 1$) and the five-year changes in their husbands incomes that span the one-year changes (between years $t + 2$ and $t - 3$). This sample moment also identifies the covariance of couples' permanent shocks and is presented with a solid line.⁷ The key point here is that the same marriage-cycle pattern in the covariance of couples' permanent income shocks, δ , is

⁷Meghir and Pistaferri assume an income process very similar to the one in equation (1) and use the following moment to estimate the permanent variance, σ_{it}^2 , from panel data:

$$\sigma_{it}^2 = E[(y_{it+m} - y_{it-1-n})(y_{it} - y_{it-1})].$$

apparent in both panel and cross-sectional data. This indicates that the changes in the cross-sectional covariance provide a sensible estimate of the permanent covariance. Innovations to permanent income are negatively correlated early in marriage and positively correlated later in marriage. Since panel methods are less sensitive to potential attrition problems, the similarity of cross-sectional and panel results suggest that sample selection from cohort attrition does not distort cross-sectional results too severely. This issue is discussed at more length in Section 5.1.

3.2 Cross-Sectional Moments Over the Business Cycle

Section 3.1 considered changes in cross-sectional moments for a given cohort over time. These changes measure the variance of permanent income change, σ^2 , or the covariance of couples' permanent income changes, δ . Given that this estimation can be done separately for each cohort, it seems natural to compare estimates of σ^2 or δ across cohorts. If cohorts differ in their attributes, cross-cohort differences in cross-sectional moments identify the impact of these attributes on the variance or covariance of income changes. For example, if the cross-sectional variance of income grows faster for a cohort of college graduates than for a cohort of high-school drop-outs, it suggests that the variance of permanent shocks, σ^2 , is larger for the highly educated. Of course, we need to worry about the possibility that these cohorts may differ in dimensions other than education (e.g., ability, quality of income data, cohort attrition, etc.) that may drive differences in cross-sectional moments.

This is the approach used by Storesletten et al., who compare the cross-sectional variance of income for cohorts who had worked through periods of greater and lesser economic expansion. They find that the cross-sectional variance of income is greater for cohorts who

Even with the persistence of shocks present in actual data (Abowd and Card, 1989), but assumed away in equation (1), any $m, n \geq 2$ can be used to obtain an unbiased estimate of σ_{it}^2 . In this paper, I use the moment developed in Shore (2006) to estimate the permanent covariance, δ_{it} , from panel data (for $m, n = 2$):

$$\delta_{it} = E[(y_{it+m} - y_{it-1-n})(x_{it} - x_{it-1})].$$

had worked through more recessions. They interpret this to mean that the variance of permanent shocks, σ^2 , is greater in recessions than in booms. I apply the same idea to the cross-sectional covariance of couples' incomes. I compare the cross-sectional covariance of husbands' and wives' excess log incomes for couples whose marriages spanned periods of greater or lesser economic expansion.

This paper uses four measures of economic conditions in year t , g_t (with g_t denoting "growth"). For ease of comparison, all are adjusted so that higher values correspond to better economic conditions:

- a. GDP growth. g_t is real annual log GDP growth from NIPA
- b. unemployment. g_t is *negative 1* times annual unemployment rate from BLS
- c. above-average growth. $g_t = 1/0$ if NIPA GDP growth is above/below average in year t
- d. *not* a recession. $g_t = 1/0$ if the year was not/was an NBER recession⁸

If σ_t^2 and δ_t^2 are affine functions of g_t , then

$$\sigma_t^2 \text{ or } \delta_t^2 = \alpha_{t-k} + \beta g_t. \quad (4)$$

In equation (4), superscripts on α and β vary with the left-hand side variable in question. For example, I will use β^x , β^y , β^{x+y} , and β^{cov} to refer to the effect of economic conditions on the variance of the wife's shocks, the variance of the husband's shocks, the variance of the shocks to household income, and the covariance of couples' shocks, respectively.⁹ Since permanent shocks accumulate over time, the variances and covariances of past shocks will be reflected in the current cross-section. Plugging equation (4) into equations (2) and (3) for cohort k yields:

⁸I use the algorithm from Storesletten et al. to assign NBER recession indicators to calendar years.

⁹While α^x , α^y , α^{x+y} , and α^{cov} may be general functions of the number of years of marriage, $t-k$, β^x , β^y , β^{x+y} , and β^{cov} are assumed to be constant over time and across cohorts. Superscripting follows the same convention for the coefficients γ and b in equation (7), where superscripts refer to the relevant left-hand side variables. x , y , $x+y$, cov refer to the cross-sectional variance of wives', husbands' or households' excess log incomes, or the cross-sectional covariance of couples' excess log incomes, respectively. α and β are assumed to be such that the variance-covariance matrix of couples' income shocks is positive definite for empirically relevant values of g .

$$E[\text{var}_k(y_t)] = E[\text{var}_k(y_k)] + \tau_t^2 + \sum_{s=k+1}^t a_{s-k}^y + \beta^y \sum_{s=k+1}^t g_s \text{ or} \quad (5)$$

$$E[\text{cov}_k(x_t, y_t)] = E[\text{cov}_k(x_k, y_k)] + \zeta_t + \sum_{s=k+1}^t a_{t-k}^{\text{cov}} + \beta^{\text{cov}} \sum_{s=k+1}^t g_s. \quad (6)$$

(i) (ii) (iii.a) (iii.b)

The current cross-sectional variance or covariance, $\text{var}_k(y_t)$ or $\text{cov}_k(x_t, y_t)$, reflects the following terms (from left to right in equations (5) and (6), with order and numbers corresponding to analogous terms in equations (2) and (3)):

- (i) initial conditions at time k , either as $\text{var}_k(y_k)$ or $\text{cov}_k(x_k, y_k)$;
- (ii) current conditions at time t , either as τ_t^2 or ζ_{it} ;
- (iii.a) years of marriage between times k and t , as $\sum_{s=k+1}^t a_{s-k}$; and
- (iii.b) cumulative economic conditions between times k and t , as $\sum_{s=k+1}^t g_s$.

4 Results

The aim of this paper is to identify the effect of economic conditions on the variance of permanent income changes (β^y) and the covariance of couples permanent income changes (β^{cov}). Recall that β^y and β^{cov} are defined in equation (4) and show up in the last terms in equations (5) and (6), respectively. A sensible statistical model reflecting equation (5) or (6) will then include

- (i) controls for initial conditions at time k , denoted X_k ;
- (ii) controls for current conditions at time t , denoted X_t ; and
- (iii.a) controls for the years of marriage between times k and t , denoted X_{t-k}

in order to estimate β , the effect of

- (iii.b) cumulative economic conditions between times k and t , $\sum_{s=k+1}^t g_s$,

on the current cross-sectional variance or covariance, $var_k(y_t)$ or $cov_k(x_t, y_t)$:

$$var_k(y_t) = \gamma_k^y X_k + \gamma_t^y X_t + \gamma_{t-k}^y X_{t-k} + b^y \sum_{s=k+1}^t g_s + e_{kt}^y; \quad (7)$$

$$cov_k(x_t, y_t) = \gamma_k^{cov} X_k + \gamma_t^{cov} X_t + \gamma_{t-k}^{cov} X_{t-k} + b^{cov} \sum_{s=k+1}^t g_s + e_{kt}^{cov}. \quad (8)$$

(i) (ii) (iii.a) (iii.b)

Note that the order and numbering of equations (7) and (8) identifies terms that are analogous to those in equations (2), (3), (5), and (6).

This paper aims to estimate the coefficients in term (iii.b) of equations (7) and (8), the effect of cumulative economic conditions over marriage, $\sum_{s=k+1}^t g_s$, on the current cross-sectional variance of excess log income (b^y) or covariance of couples' excess log incomes (b^{cov}). Assuming that a) labor income evolves as in equation (1), b) controls X_k , X_t , and X_{t-k} are adequate, and c) cohort composition is stable over time, then b from the statistical model in equations (7) and (8) estimates β from the economic model in equations (4), (5), and (6) (i.e., $b = \beta$). $\beta^y > 0$ (or $\beta^{cov} > 0$) implies that the variance of permanent income changes (or covariance of couples' permanent income changes) is greater in booms than in busts.

My preferred choices for controls X_k , X_t , and X_{t-k} in equations (7) and (8) are as follows. In all regressions, X_{t-k} are the set of dummy variables for the number of years of marriage, $t - k$, from zero to 35. This captures life-cycle variation in permanent and transitory risk. In regressions labeled “no controls,” I do not include controls X_k and X_t . This implicitly assumes that all cohorts have the same degree of heterogeneity in ability ($\gamma_k^y = 0$), the same pattern in assortative mating ($\gamma_k^{cov} = 0$), and that transitory risk has been constant over time ($\gamma_t = 0$). In regressions labeled “with controls,” X_t is a linear time trend; X_k are linear controls for economic conditions (real log GDP growth) in year k and also in the three years prior to year k . These control for time- and cohort-specific factors. For example, assortative mating or the dispersion of ability may vary with economic conditions. All regressions are population-weighted by the number of individuals in a given cohort-year observation, with

standard errors clustered at the cohort level.

This reduced-form approach contrasts with the structural model used by Storesletten et al. Their structure imposes assumptions about the dispersion of ability, or equivalently, the initial distribution of income. In particular, they assume that the cross-sectional variance of initial income ($var_k(y_k)$) is the same for all cohorts, effectively imposing $\gamma_k^y = 0$. While this assumption is defensible, its analog for this paper – that the degree of assortative mating ($cov_k(x_k, y_k)$) is the same in each year of marriage, effectively imposing $\gamma_k^{cov} = 0$ – is not. While the dispersion of innate ability may or may not have changed over time, Mare (1991) and Rose (2001) have documented that mating has become more positively assorted over time. This implies that cross-sectional covariance of couples’ incomes is positively correlated with calendar time¹⁰ while GDP growth is negatively correlated with calendar time¹¹. As a result, excluding a control for time will bias the coefficient of interest – the effect of GDP growth on the covariance of couples’ incomes – downward.¹²

4.1 C.S. Variance

Table 1 estimates equation (7), examining the effect of past economic conditions on the cross-sectional variance of excess log income. Income is defined as the husband’s labor income in odd-numbered columns and as the sum of the wife’s and the husband’s labor incomes in even-numbered columns. Table 1 has 4 pairs of columns, with each pair corresponding to one of the four measures of economic conditions described in Subsection 3.2.

Table 1 shows that income risk is countercyclical. The upper “no controls” panel of

¹⁰The cross-sectional variance is also positively correlated with calendar time as is shown in the “with controls” panel of Table 1. An increase over time in transitory volatility that could explain this pattern has been noted in other research. (Gottschalk and Moffitt, 1994; Hacker, 2006)

¹¹While this pattern is less stark in the very first and last years couples in the sample got married, relatively few observations are for marriages that span these years. As a result, a population-weighted regression of g_t on a time trend is strongly and significantly negative.

¹²Furthermore, if the economic benefits of marriage vary with macroeconomic conditions, this could induce business-cycle variation in assortative mating. For example, if the economic benefits of marriage are greater in recessions, then there are economically inefficient marriages – perhaps positively assorted ones (Becker, 1973, 1974) – that are worth undertaking in recessions but not in booms. In this case, given that conditions at the time of marriage are positively correlated with conditions during marriage, failing to control for conditions at time of marriage will bias the coefficient of interest toward zero.

column 7 is a reduced-form representation of the main result from Storesletten et al., though with marriage-based cohorts instead of age-based cohorts. While I do not show results for age-based cohorts here, these are similar. The coefficient $\hat{b}^y = -0.0109$ is strongly and significantly negative, indicating that the variance of excess log income changes falls by 0.0109 in non-NBER recession years compared to NBER recession years. Note that this result is also apparent in the other “no controls” regressions, which use different measures of economic conditions.

The results in Table 1 are sensitive to the inclusion of additional controls. The lower “with controls” panel relaxes the assumption that $\gamma_k = \gamma_t = 0$ to allow the dispersion of initial ability to vary over time and over the business cycle. Including these covariates leads the countercyclical risk result ($\hat{b}^y < 0$) to lessen (become less negative) for the first three measures of economic conditions (columns 1 through 6) though not for the fourth (which is the main focus of Storesletten et al.). As explained on page 15, excluding the time trend – as in the upper “no controls” panel of Table 1 – mechanically biases estimates of b downward.

The key finding from Table 1 is that the relationship between the cross-sectional variance of excess log income and past economic conditions is much stronger for husband’s labor income ($\hat{b}^y < 0$, odd-numbered columns of Table 1) than for total household labor income ($\hat{b}^{x+y} \approx 0$, even-numbered columns of Table 1). Better cumulative economic conditions predict a lower cross-sectional variance of *husbands’* excess log incomes, so that there is less risk to husbands’ incomes in booms. This result is universally weaker and mostly absent when looking at household labor income. Once time trends are included (the lower “with controls” panel), better cumulative economic conditions do not predict a significantly lower cross-sectional variance of *households’* excess log incomes (except in columns 7 and 8 where economic conditions are defined by the number of non-recession years). This paper aims to understand the divergence between results for husbands’ labor incomes (substantial) and the sum of husbands’ and wives’ labor incomes (absent or weaker). The cross-sectional covariance of couples’ incomes sheds light on that divergence.

4.2 C.S. Covariance

Table 2 presents estimates of b^{cov} from equation (8), the effect of past economic conditions on the cross-sectional covariance of couples' excess log incomes. The regressions in Table 2 are identical to the “with controls” results in Table 1, though the dependent variable is the cross-sectional covariance of husbands' and wives' excess log incomes ($cov_k(x_t, y_t)$) not the cross-sectional variance of excess log income ($var_k(y_t)$). $\hat{b}^{cov} > 0$ for all four measures of economic conditions, and strongly and significantly so for the first three of these measures. The cross-sectional covariance of husbands' and wives' excess log incomes is more positive (less negative) for cohorts of couples who have been married through periods of greater economic expansion. Put more informally, there is better risk-sharing – either diversification or dynamic coordination – in bad times than in good. This explains why individual risk increases more in bad times than household risk does. While there may be more risk for individuals in recessions, improved intra-household coordination undoes most or all of this increase in risk.

How large is this effect? If annual GDP growth increases by an additional 1 percent, then the covariance of innovations to couples' permanent incomes will increase by 0.004 ($\hat{b}^{cov} \approx 0.4$ from column 1 of Table 2). To get a rough sense of the scale of this result, the moment used by Meghir and Pistaferri and described in footnote 7 can be used to obtain estimates of the permanent variance for husbands and for wives. These estimates ($\sigma_y^2 = 0.035$ and $\sigma_x^2 = 0.25$, respectively) imply that a 1 percent increase in annual GDP growth increases the correlation of couples' labor income changes (from a baseline of zero, as discussed in Section 5.2) by just over 4 percent ($0.004/\sqrt{0.035 \times 0.250} \approx 0.04$). If the difference in annual GDP growth between good and bad times is 5 percent, the correlation of couples' income changes falls by roughly 21 percent from good times to bad. A very similar result is obtained using the number of above-average years of GDP growth (column 3) as the measure of economic conditions. Here, the coefficient of roughly 0.014 should be interpreted as the impact of an above-average year of GDP growth on the cross-sectional covariance. The scale of this

effect is similar, implying that the correlation of innovations to permanent income rises by 15 percent ($0.014/\sqrt{0.035 \times 0.250} \approx 0.15$) from below average to above average years. Comparable results for unemployment growth and NBER recessions (also shown in Table 2) are of the same order but are somewhat smaller.

Figure 4 shows a (significant) positive relationship between a cohort’s the average cross-sectional covariance (y-axis) and its average GDP growth (x-axis).¹³ This pattern is apparent whether outliers from early years of data are included (Panels A and C) or excluded (Panels B and D). Upper panels (A and B) identify each cohort with a circle proportional to the number of household observations in the data for that cohort; lower panels (C and D) identify each cohort by the year in which it was formed.

Table 3 repeats column 1 of Table 2 in estimating equation (8), but shows alternative specifications for cohort and year controls, X_k and X_t . Again, the cross-sectional covariance of couples’ excess log incomes is predicted with cumulative GDP growth over the marriage to date. While point estimates of b^{cov} vary, they are consistently and significantly positive. For current year controls, X_t , I include dummy variables for year (columns 6 and 7), quadratic time trends (column 3), and current economic conditions (column 5). For cohort or year of cohort controls, X_k , I include economic conditions at time of and just before marriage (columns 4 and 5). However, $\hat{b}^{cov} > 0$ is only apparent after controlling for the time trend in the cross-sectional covariance of income (columns 2 through 8). In column 1 of Table 4, which excludes a time trend, \hat{b}^{cov} is insignificant. As discussed on page 15, excluding this trend biases estimates of b^{cov} downward given the trends in assortative mating over this period.

When controls X_k in equation (8) are dummy variables for cohort year, k (last column in Table 3), results are no longer significant. Here, results are identified by within-cohort

¹³The covariances that are displayed are calculated first by running a population-weighted regression of cohort-year covariances on “years of marriage” dummy variables, X_{t-k} , to remove life-cycle patterns in the cross-sectional covariance nonparametrically. Next, residuals from this regression are collapsed (with population weights) by cohort. These cohort averages are then regressed on a linear time trend to remove trends in assortative mating. The residuals from this population-weighted regression are then plotted against the average demeaned GDP growth over the marriage to date for the average observation in each cohort.

variation over time. The fact that some cohorts both lived through more recessions and have low *levels* of covariance is no longer considered. Instead, only *changes* in the covariance over time for a cohort are used as identifying variation. Since observations for different years from the same cohort share the same macroeconomic history before 1968, this approach exploits variation in economic conditions after but not before 1968. The absence of a significant result here demonstrates the need for variation over longer periods than are spanned by conventional panel data sets (as the cross-sectional methods in this paper do).

The finding that past economic conditions predict the subsequent cross-sectional covariance of couples' excess log incomes is robust to choice of sample. Results (not shown) are qualitatively similar for couples married before and after 1968, when data is taken before and after 1982, for couples married for more and less than 15 years and for couples whose husbands have less than or at least a high school degree. I find $\hat{b}^{cov} > 0$ for all of these sub-samples, though naturally standard errors are larger.

4.3 Decomposing Household Income Risk

Shocks to household permanent income will be driven by shocks to the permanent incomes of husbands and wives. We can then define a measure of the household's permanent income shock, ω_{x+y} :

$$\omega_{x+yit} \equiv s_{it-1}\omega_{xit} + (1 - s_{it-1})\omega_{yit}. \quad (9)$$

The household's permanent shock depends on the wife's income share, s , and each spouse's permanent income shock, ω_{xit} , and ω_{yit} . ω_{x+yit} approximates the change in permanent excess log household income.¹⁴ We can then decompose the proxy for the permanent household

¹⁴If we are willing to define the *level* of permanent income for the wife, husband, and household as $P_{xit} \equiv e^{p_{xit}}$, $P_{yit} \equiv e^{p_{yit}}$, $P_{x+yit} \equiv P_{xit} + P_{yit}$, we can then calculate log household permanent income and the wife's income share as $p_{x+yit} \equiv \ln(P_{x+yit})$, and $s_{it} \equiv P_{xit}/P_{x+yit}$. Change to log permanent household income, $p_{x+yit} = \ln(e^{p_{xit}} + e^{p_{yit}})$, is

$$\begin{aligned} p_{x+yit} - p_{x+yit-1} &= \ln\left(\frac{e^{p_{xit-1}}e^{\omega_{xit}} + e^{p_{yit-1}}e^{\omega_{yit}}}{e^{p_{xit-1}} + e^{p_{yit-1}}}\right) = \ln(s_{it-1}e^{\omega_{xit}} + (1 - s_{it-1})e^{\omega_{yit}}) \\ \ln(s_{it-1}e^{\omega_{xit}} + (1 - s_{it-1})e^{\omega_{yit}}) &\approx s_{it-1}\omega_{xit} + (1 - s_{it-1})\omega_{yit} \equiv \omega_{x+yit} \end{aligned}$$

variance, $\sigma_{x+yit}^2 \equiv E[\omega_{x+yit}]$, as

$$\sigma_{x+yit}^2 = s_{it-1}^2 \sigma_{xt}^2 + (1 - s_{it-1})^2 \sigma_{yt}^2 + 2s_{it-1}(1 - s_{it-1}) \delta_t \quad (10)$$

Total household risk depends on the income risks of husbands and wives, and on the covariance of their income shocks. By approximating s_{it} with its population average, \bar{s} , and differentiating equation (10) with respect to g_t , the following rough decomposition obtains:

$$\beta^{x+y} \approx \bar{s}^2 \beta^x + (1 - \bar{s})^2 \beta^y + 2\bar{s}(1 - \bar{s}) \beta^{cov}. \quad (11)$$

If individual risk decreases in good times ($\beta^x, \beta^y < 0$), then household risk will only remain constant through the business cycle ($\beta^{x+y} \approx 0$) if the covariance of couples' shocks increases in good times ($\beta^{cov} > 0$). This is exactly what we observed in Tables 1 and 2. Table 4 presents these results side-by-side.

Column 1 of Table 4 presents estimates of b^{x+y} (with the caveat in footnote 14), defined as the effect of cumulative economic growth on the cross-sectional variance of excess log *household* income. This small and statistically insignificant result ($\hat{b}^{x+y} = 0.05$) is identical to the one found in the “with controls” panel in column 2 of Table 1. However, cumulative economic conditions predict lower cross-sectional variances of excess log *husband's* income and also *wife's* income, $\hat{b}^y = -0.29$ and $\hat{b}^x = -0.54$ in columns 2 and 3, respectively.¹⁵ Column 4 shows the impact of cumulative economic conditions on the cross-sectional covariance of couples' excess log incomes, $\hat{b}^{cov} = 0.39$, which merely repeats results from column 1 of Table 2 (or column 4 of Table 3). Assuming an empirically relevant $s = 1/4$, the left

Note that the approximation in the last line will hold with equality only when the wife's income share, s , is very close to 0 or to 1. Only at these limits, when s does not change over time, will ω_{x+yit} be i.i.d when ω_{xit} and ω_{yit} . Otherwise, ω_{x+yit} is just a linear approximation of $p_{x+yit} - p_{x+yit-1}$. The methods used in this paper will identify the permanent variance of household income only approximately, regardless of whether it is defined as $E[(p_{x+yit} - p_{x+yit-1})^2]$ or as $E[\omega_{x+yit}^2]$.

¹⁵Results differ very slightly from Table 1 because the censoring procedure for the final three columns of this table eliminates all observations with missing data or extreme changes in either the wife's or the husband's income, while Table 1 dropped only observations with missing data or extreme changes in the husband's income.

and right hand sides of equation (11) – calculated from \hat{b} and $s = 1/4$ as 0.05 and -0.05 , respectively – are similar in magnitude and statistically indistinguishable from zero and each other. In other words, all or nearly all of the increased risk for individuals in periods of low growth is offset by improved risk-sharing by couples at these times. The hypothesis that this offset is perfect cannot be rejected.

5 Alternative Explanations

5.1 Selection

So far, this paper has documented that the cross-sectional covariance of couples' incomes depends on past economic conditions, with estimates of equation (8) shown in Tables 2 and 3 and discussed in Section 4.2. Recall from the discussion of equation (8) that in order to infer the covariance of couples income changes from the subsequent cross-sectional covariance, the composition of a given cohort must be stable over time. If the kinds of couples who divorce in a recession differs from those who divorce in a boom, this may affect the cross-sectional covariance of couples' incomes.¹⁶

Imagine that couples with similar incomes (or couples whose incomes moved together) were more like to stay married in booms and get divorced in busts. In this case, cohorts whose marriages spanned periods of greater economic expansion would have higher cross-sectional covariances. In results not shown, I investigate this hypothesis directly with divorce hazard data from the PSID. I predict divorce with current or lagged (three prior years, cumulative) GDP growth, the product of husbands' and wives' incomes (to proxy for assortative mating, with controls for the income level of each spouse), and the interaction of these measures. The coefficient on the interaction term is small, insignificant, and robust

¹⁶As pointed out by an anonymous referee, this potentially problem is particularly worrisome if the economic benefits of marriage vary over the business cycle. Note that as long as controls for economic conditions at time of marriage, X_k , are sufficient, the kind of couples who marry in good times can differ from the kind who marry in bad times.

to sample or additional controls. This suggests that positively assorted couples are not more likely to get divorced than other couples in good or bad times. I perform the same exercise, replacing the product of couples' incomes with the product of couples' one-year income changes, to proxy for the covariance of the couples' income changes. Again, the coefficient on the interaction term is small, insignificant, and robust to sample or additional controls. I find no evidence that business cycles impact who gets divorced.

5.2 Procyclic Covariance \leftrightarrow Procyclic Correlation

When variances vary with economic conditions, covariances can vary mechanically even when correlations do not. In particular, the finding that covariances fall in recessions could be explained solely by increased variances in recessions if the correlation (assumed to be constant in this thought experiment) were sufficiently negative.¹⁷ This cannot explain the results presented here. First, the finding that there is more risk in recessions for individuals but not households is only consistent with a constant correlation of couples' income changes if that correlation is -1 . Second, the covariance of couples' income changes is approximately zero in the data. This is apparent by inspection of Figure 3, which shows that the covariance of couples' permanent shocks are negative early in marriage and positive later in marriage. More rigorously, Shore (2006) shows that the absolute value of the correlation of couples' income shocks is small (7 percent or less). This fact is confirmed and exploited in Shore (2007), which documents heterogeneity in covariance by exploiting the fact that covariances are close to zero on average. When $\delta = 0$, then $\beta^{cov} = \frac{d\rho}{dg}\sigma_x\sigma_y$ and changes in covariance map directly to changes in correlation. Therefore, a procyclic covariance implies a procyclic

¹⁷Let $\rho \equiv \delta/\sqrt{\sigma_x^2\sigma_y^2}$ denote the correlation of couples' permanent income changes. Differentiating with respect to g , rearranging terms, and substituting from equation (4) yields:

$$\beta^{cov} = \frac{d\rho}{dg}\sigma_x\sigma_y + \frac{1}{2}\delta\left(\frac{\beta^y}{\sigma_y^2} + \frac{\beta^x}{\sigma_x^2}\right).$$

Table 1 (and previous research) documents that $\beta^y < 0$. Note that $\beta^y < 0$, $\frac{d\rho}{dg} = 0$ and $\beta^{cov} > 0$ are not inconsistent so long as δ is sufficiently negative.

correlation. When correlations are close to zero, changes in the variance will have little impact on the covariance.

5.3 Interpretation

The main drawback of the cross-sectional methods used here is that behavior is not observed in booms or busts. We can only see the echo of past behavior in the current cross-section. While this paper identifies countercyclic risk-sharing between spouses, it is not clear what type of risk-sharing behavior is counter-cyclic. One possibility is that the diversification is countercyclic; the covariance of couples' exogenous shocks is more negative in bad times than good. For example, in bad times couples might avoid pairs of jobs with highly correlated shocks. Another possibility is that dynamic coordination, the response of one spouse to a shock to the other spouse, is countercyclic.

One rough way to distinguish between shocks (diversification) and choices (dynamic coordination) is to decompose earnings into wages and hours. By definition, income is merely the product of the hourly wage and the number of hours worked. Therefore, a change in log permanent income is just the sum of an hours change and a wage change, $\omega_t = \omega_t^{wage} + \omega_t^{hours}$. A stylized model might consider wage changes to be exogenous shocks and hours changes to be chosen. However, there are many settings where hours changes are involuntary (e.g., layoffs) or where wage changes are chosen (e.g., changing to a job with a different compensating wage differential).

Since husbands are typically employed full-time in the long run, their permanent variation in hours is small. Therefore, I limit the hours-wage decomposition to wives. If the covariance of couples' income change is procyclic ($\beta^{cov} > 0$), then either the covariance of husbands' incomes and wives' hours is procyclic ($\beta_{xhours}^{cov} > 0$), the covariance of husbands' incomes and wives' wages is procyclic ($\beta_{xwage}^{cov} > 0$), or some combination ($\beta^{cov} = \beta_{xwage}^{cov} + \beta_{xhours}^{cov} > 0$).

Table 6 presents this decomposition, using estimates of b from the statistical model in equation (8) to measure β from the economic model in equations (4) and (6). The regressions

here mirror those from previous tables, though the cross-sectional covariance of husbands' excess log incomes and wives' hours or wives' wages are used as dependent variables. While the decomposition of the covariance is estimated imprecisely, the covariance of husbands' incomes both with wives' wages and with wives' hours seems to be more negative when growth is low. The total effect for wives' income ($\hat{b}^{cov} = 0.39$, column 1) is decomposed into the effect for wives' hours ($\hat{b}_{xhours}^{cov} = 0.17$, column 2) and wives' wages ($\hat{b}_{xwage}^{cov} = 0.20$, column 3), where the latter two effects sum roughly to the first.¹⁸ The wage effect is significantly different from zero, while the hours effect is not. Both wives' wages and hours seem to covary with husbands' incomes more in booms than in busts.

The interpretation of this result is unclear. Perhaps wives whose husbands receive negative shocks are more likely to increase work hours in recessions. This leads to increased skills and attachment to the labor force, so that their subsequent wages and hours are higher. Another possibility is that in recessions, wives whose husbands receive negative shocks are more likely to choose jobs with larger compensating wage differentials. In other words, they respond to negative shocks in recessions by working harder or in more unpleasant jobs, even if they do not work substantively more hours.

These findings still leave important unanswered questions about the mechanism of action. For example, results could be driven entirely by business-cycle variation in fertility. If couples' were more likely to have children in recessions and parents of young children have more negatively correlated income changes (the need for home production means that one spouse working more likely entails the other working less), this could explain the procyclical covariance of couples' income changes. Since this sort of coordination is not risk-reducing, this explanation would have substantially different welfare implications. In practice, this explanation is unlikely. National fertility is quite stable at high frequencies; Mocan (1990) argues that, if anything, fertility is procyclic.

¹⁸The decomposition is not exact because the Winsorizing procedure is performed separately for each variable and therefore does not bind for exactly the same observations for income, wage and hours.

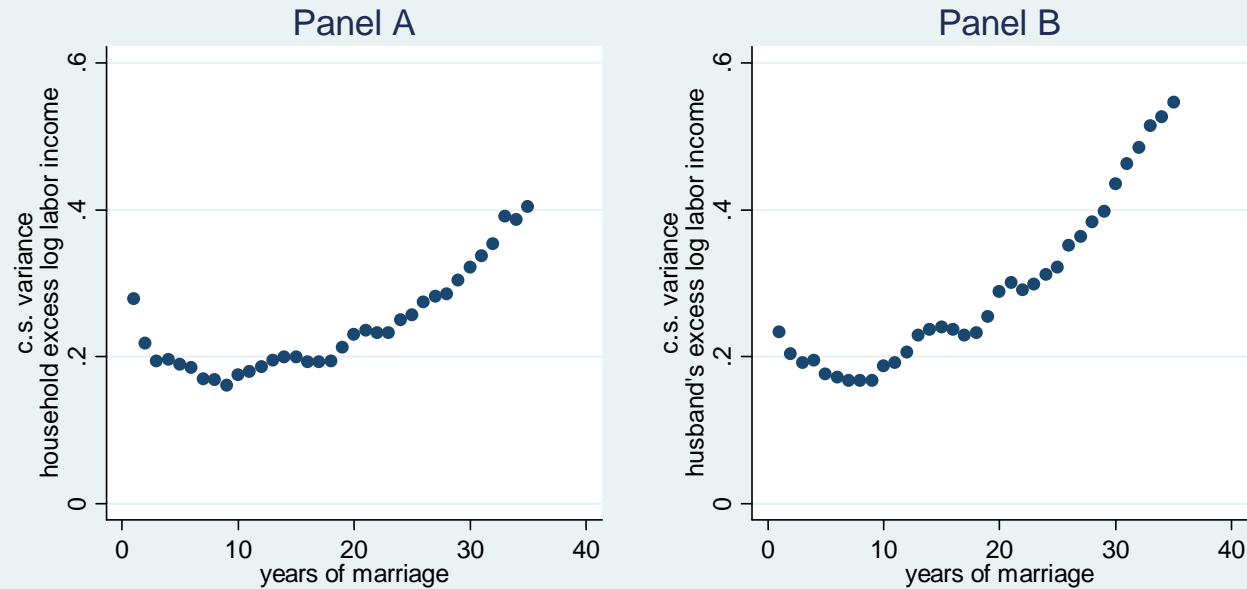
6 Conclusion

Husbands face more idiosyncratic labor income risk in recessions. However, after controlling for trends in cross-sectional variance over time, this pattern is absent or substantially weaker for household's (husband's plus wife's) labor income risk. This divergence can be explained by improved intra-household risk-sharing in recessions; husbands and wives have more negatively correlated innovations to permanent income in recessions than in booms. This effect is quite large. While the average correlation of couples' innovations to permanent labor income is close to zero, this masks large life-cycle and business-cycle patterns. The correlation of couples' innovations to permanent labor income falls by roughly 20 percent in recessions relative to booms.

References

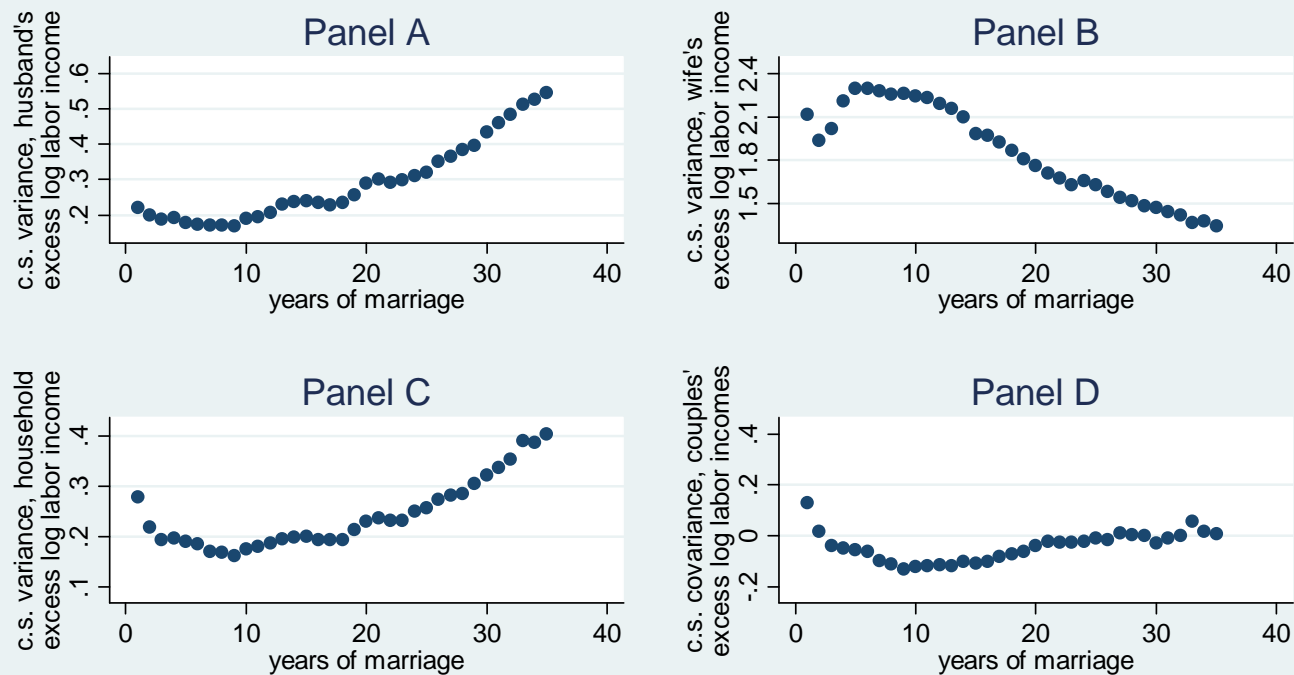
- ABOWD, J., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–445.
- BECKER, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813–846.
- (1974): "A Theory of Marriage: Part II," *Journal of Political Economy*, 82(2/2), S11–S26.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40(1), 41–71.
- CULLEN, J. B., AND J. GRUBER (2000): "Does Unemployment Insurance Crowd Out Spousal Labor Supply?," *Journal of Labor Economics*, 18(3), 546–72.
- DEATON, A., AND C. PAXSON (1994): "Intertemporal Choice and Inequality," *Journal of Political Economy*, 102(3), 437–467.
- DIXON, W. J. (1960): "Simplified Estimation from Censored Normal Samples," *The Annals of Mathematical Statistics*, 31, 385–391.
- GOTTSCHALK, P., AND R. MOFFITT (1994): "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 0(2), 217–254.
- HACKER, J. (2006): *The Great Risk Shift: The Assault on American Jobs, Families, Health Care, and Retirement; And How You Can Fight Back*. Oxford University Press, Oxford.
- JUHN, C., AND S. POTTER (2007): "The Impact of Working Spouses on Cyclical Fluctuations in the Labor Force," Presentation at the 2007 Society of Labor Economists Annual Meeting.
- LONG, C. (1953): "Impact of effective demand on the labor supply," *American Economic Review Papers and Proceedings*, 43, 458–467.
- LUNDBERG, S. (1985): "The Added Worker Effect," *Journal of Labor Economics*, 3(1), 11–37.
- MARE, R. (1991): "Five Decades of Educational Assortative Mating," *American Sociological Review*, 56(1), 15–32.
- MEGHIR, C., AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.
- MOCAN, N. H. (1990): "Business Cycles and Fertility Dynamics in the United States: A Vector Autoregressive Model," *Journal of Population Economics*, 3(2), 125–46.
- PARKER, S. W., AND E. SKOUFIAS (2004): "The Added Worker Effect over the Business Cycle: Evidence from Urban Mexico," *Applied Economics Letters*, 11(10), 625–630.
- ROSE, E. (2001): "Marriage and Assortative Mating: How Have the Patterns Changed?," working paper, economics department, University of Washington.
- SHORE, S. H. (2006): "The Co-Movement of Couples' Incomes," University of Pennsylvania Working Paper.
- (2007): "Testing for Correlated Heterogeneity: Theory and Evidence from Couples' Incomes," Johns Hopkins University.
- STEPHENS, M. (2002): "Worker displacement and the added worker effect," *Journal of Labor Economics*, 20(3), 504–37.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695–717.

Figure 1: C.S. Variance of Excess Log Income Over the Life-Cycle




Note: Each panel plots averages for each year of experience or marriage against the cross-sectional variance of income. Panel A uses log excess total household labor income as the measure of income; Panel B uses log excess labor income of the household head as the measure of income. Year-specific averages are found by regressing the cross-sectional moments on calendar year dummy variables and dummy variables for the number of years of marriage. Regressions are population weighted. The y-axis presents the dummy variables for years of work experience or marriage, where all dummy variables are adjusted by the average of the calendar year dummy variables to keep scaling comparable across panels.

Figure 2: Couples' C.S. Moments over Life-Cycle of Marriage



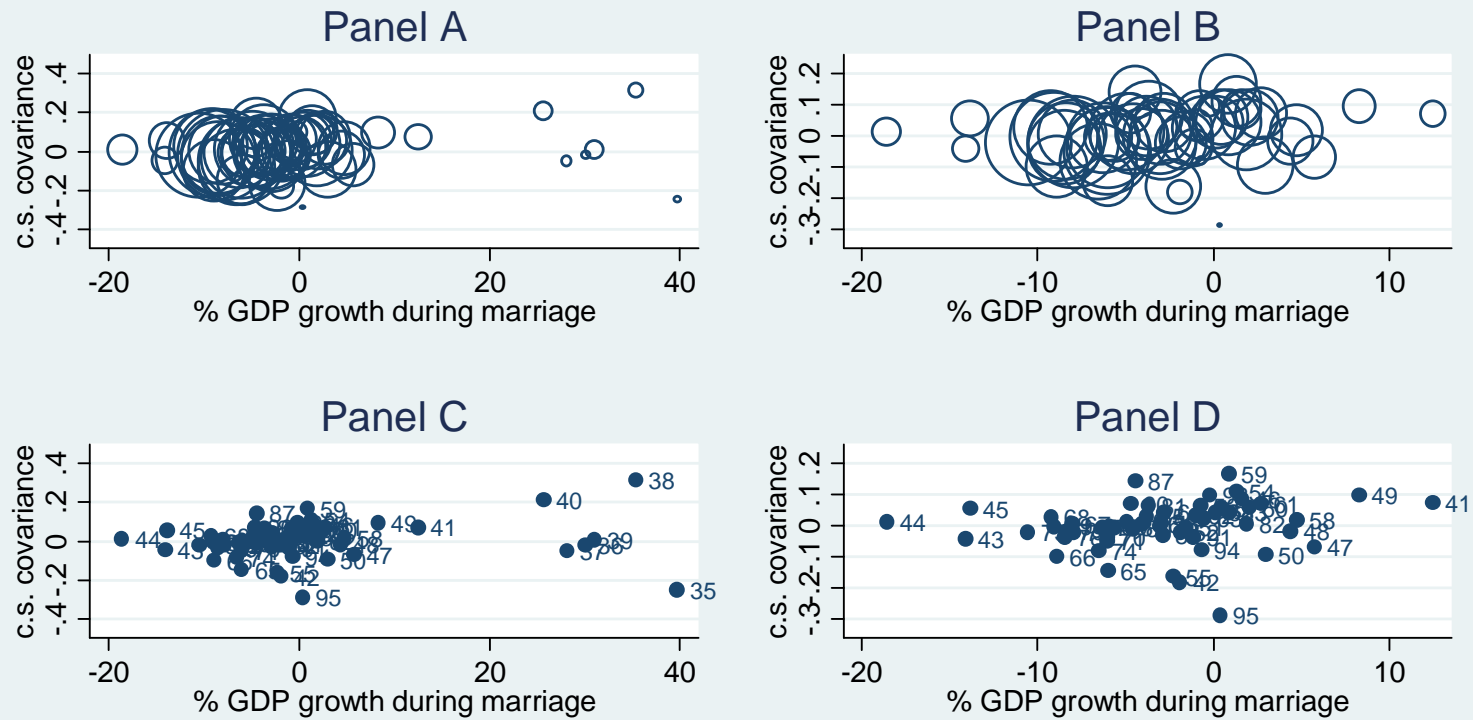
Note: Each panel plots averages for each year of marriage against the cross-sectional variance or covariance of income. The sample is the set of all married couples in the husband's first marriage who never divorce in the data. Panel A presents the cross-sectional variance of log excess income for husbands' labor income; Panel B presents wives' labor income; Panel C presents the sum of the couples' labor incomes. Panel D presents the cross-sectional covariance of the log excess income of husband and wife. Year-specific averages are found as in Figure 1.

Figure 3:
Covariance of Couples' Innovations to Permanent Income



Note: This figure plots estimates of the cross-sectional covariance of innovations to permanent excess income for a given number of years of marriage. The x-axis indicates the number of years of marriage. The y-axis shows the covariance. Given that the variance of permanent innovations to husbands' and wives' incomes are 0.035 and 0.25 respectively using the moment developed by Meghir and Pistaferri (2004), divide covariances by 0.09 to obtain a rough estimate of the correlation. The blue (top on the left side) line presents an estimate from cross-sectional data, the derivative of coefficients from a quartic in the number of years of marriage used to predict the cross-sectional variance. The red (bottom on the left side) line presents estimates from panel data. See text for details.

Figure 4: GDP Growth and C.S. Covariance



Note: Cross-sectional covariances are population-weighted averages for each cohort. Cross-sectional covariances are adjusted for years-of-marriage non-parametrically, with dummy variables, before they are collapsed at the cohort level. A linear time trend in year of marriage is then removed from these cohort averages, so that plots are adjusted for year of marriage and years of marriage. In Panels A and C, circle size is proportional to the number of observations for that cohort. In Panels B and D, number labels refer to the year of marriage for that cohort. Panels A and B include data for all cohorts, while Panels C and D include only cohorts married after 1940.

Table 1: Impact of Cumulative Economic Conditions on Cross-Sectional Variance of Income
Individual versus Household Income

	Dependent Var.	Cross-Sectional Variance of Income							
	Income Definition	Individual	Household	Individual	Household	Individual	Household	Individual	Household
	Economic Measure	GDP Growth		-1 x Unemployment Rate		GDP Growth Above Average (1 or 0)		Not an NBER Recession (1 or 0)	
	Cohort Definition	Marriage-Based							
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
No Controls	Cumulative Conditions	-0.637 (0.097)**	-0.199 (0.059)**	-0.412 (0.037)**	-0.191 (0.027)**	-0.039 (0.004)**	-0.016 (0.003)**	-0.026 (0.007)**	-0.008 (0.003)*
	Observations	1,073	1,073	1,048	1,048	1,073	1,073	1,073	1,073
	R ²	0.305	0.250	0.423	0.331	0.353	0.287	0.193	0.228
With Controls	Cumulative Conditions	-0.297 (0.084)**	0.055 (0.063)	-0.183 (0.063)**	-0.002 (0.047)	-0.013 (0.005)**	0.002 (0.004)	-0.021 (0.005)**	-0.006 (0.003)*
	Year	0.006 (0.001)**	0.004 (0.001)**	0.005 (0.001)**	0.004 (0.001)**	0.006 (0.001)**	0.004 (0.001)**	0.007 (0.001)**	0.004 (0.000)**
	Initial Conditions	-0.146 (0.162)	-0.089 (0.122)	0.027 (0.164)	-0.101 (0.119)	-0.136 (0.152)	-0.096 (0.120)	-0.121 (0.163)	-0.130 (0.119)
	Prior Conditions	-0.290 (0.241)	-0.134 (0.220)	0.077 (0.217)	-0.191 (0.203)	-0.111 (0.224)	-0.175 (0.212)	-0.067 (0.234)	-0.227 (0.209)
	Observations	1,073	1,073	1,048	1,048	1,073	1,073	1,073	1,073
	R ²	0.476	0.407	0.484	0.408	0.469	0.407	0.487	0.411

Note: Standard errors, clustered at the cohort level, are in parentheses. “*” and “**” indicate significance at the 5% and 1% level, respectively. The economic measure over the course of marriage excludes the first year of marriage. The measure of economic conditions are, in order: log GDP growth; the inverse of the unemployment rates in each year; whether the year had above average GDP growth; and, whether the year was not an NBER recession, using the definition from Storesletten et al. (2004) to map NBER recessions to years. Cumulative Conditions is the sum of the economic measures over the marriage. Initial Conditions is the GDP growth in the 1st year of marriage; Prior Conditions is the GDP growth for the 3 Yrs. before marriage. Cohorts consist of married couples (with the husband in his first marriage, who do not subsequently divorce in the sample) who were married in the same year. In odd-numbered columns, the measure of income is total (husband plus wife) labor income. In even-numbered columns, the measure of income is the husband’s labor income. Moments are calculated as the cross-sectional variance of excess log income, as described in the text. Regressions also include dummy variables for the number of years that the cohort has been married.

Table 2: Impact of Cumulative Economic Conditions on the Cross-Sectional Covariance of Couples' Incomes

Dependent Var.	Cross-Sectional Covariance of Couples' Incomes			
Economic Measure	GDP Growth	-1 x Unemployment Rate	GDP Growth Above Average (1 or 0)	Not an NBER Recession (1 or 0)
Cohort Definition	Marriage-Based			
	(1)	(2)	(3)	(4)
Cumulative Conditions	0.393 (0.106)**	0.180 (0.080)*	0.014 (0.006)*	0.007 (0.005)
Year	0.006 (0.001)**	0.006 (0.001)**	0.006 (0.001)**	0.004 (0.001)**
Initial Conditions	-0.253 (0.230)	-0.418 (0.231)	-0.289 (0.232)	-0.361 (0.231)
Prior Conditions	0.359 (0.335)	-0.133 (0.316)	0.085 (0.311)	-0.043 (0.307)
Observations	1,073	1,048	1,073	1,073
R ²	0.242	0.238	0.228	0.221

Note: Standard errors, clustered at the cohort level, are in parentheses. “*” and “**” indicate significance at the 5% and 1% level, respectively. The dependent variable is the cross-sectional covariance of husbands' excess log incomes and wives' excess log incomes from a given calendar year for the cohort of couples married in the same year. The economic measure over the course of marriage excludes the first year of marriage. The measure of economic conditions are, in order: log GDP growth; the inverse of the unemployment rates in each year; whether the year had above average GDP growth; and, whether the year was not an NBER recession, using the definition from Storesletten et al. (2004) to map NBER recessions to years. Cumulative Conditions is the sum of the economic measures over the marriage. Initial Conditions is the GDP growth in the 1st year of marriage; Prior Conditions is the GDP growth for the 3 Yrs. before marriage. Cohorts consist of married couples (with the husband in his first marriage, who do not subsequently divorce in the sample) who were married in the same year. The measure of income is labor income of the husband or wife. Moments are calculated as the cross-sectional covariance of husbands' and wives' excess log incomes, as described in the text. Regressions also include dummy variables for the number of years that the cohort has been married.

Table 3: Impact of Cumulative Economic Conditions on the Cross-Sectional Covariance of Couples' Incomes
Robustness Checks: Specification

Dependent Var.	Cross-Sectional Covariance of Couples' Incomes							
Economic Measure	GDP Growth							
Cohort Definition	Marriage-Based							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cumulative Conditions	-0.027 (0.087)	0.344 (0.100)**	0.310 (0.109)**	0.393 (0.106)**	0.442 (0.112)**	0.459 (0.137)**	0.610 (0.137)**	0.023 (0.174)
Year		0.006 (0.001)**	0.001 (0.008)	0.006 (0.001)**	0.006 (0.001)**			
(Year) ²			0.000 (0.000)					
Initial Conditions				-0.253 (0.230)	-0.240 (0.231)		-0.178 (0.240)	
Prior Conditions				0.359 (0.335)	0.418 (0.337)		0.630 (0.351)	
Current GDP Growth					-0.045 (0.159)			-0.256 (0.179)
Dummy Variables	numyears	numyears	numyears	numyears	numyears	numyears, endyear	numyears, endyear	numyears, begyear
Observations	1,073	1,073	1,073	1,073	1,073	1,073	1,073	1,073
R ²	0.134	0.232	0.233	0.242	0.249	0.257	0.272	0.474

Note: Standard errors, clustered at the cohort level, are in parentheses. “*” and “**” indicate significance at the 5% and 1% level, respectively. The dependent variable is the cross-sectional covariance of husbands' excess log incomes and wives' excess log incomes from a given calendar year for the cohort of couples married in the same year. The economic measure over the course of marriage excludes the first year of marriage (and the last year of marriage for columns 5 and 8). The measure of economic conditions is log GDP growth. Cumulative Conditions is the sum of the economic measures over the marriage. Initial Conditions is the GDP growth in the 1st year of marriage; Prior Conditions is the GDP growth for the 3 Yrs. before marriage. Cohorts consist of married couples (with the husband in his first marriage, who do not subsequently divorce in the sample) who were married in the same year. The measure of income is labor income of the husband or wife. Moments are calculated as the cross-sectional covariance of husbands' and wives' excess log incomes, as described in the text. As denoted by “numyear” in the “dummy variable” row, regressions also include dummy variables for the number of years that couples in a given cohort have been married. “Endyear” or “Begyear” in this column denotes that dummy variables for the year of marriage (begyear) or current year (endyear) have also been included in the regression. Column 4 from this table repeats column 4 from Table 2.

Table 4: Impact of Cumulative Economic Conditions on Cross-sectional moments
Decomposing the variance of total household income

Dependent Var.	Cross-Sectional Variance			Cross-Sectional Covariance
	Household	Husband	Wife	Husband, Wife
Economic Measure	GDP Growth			
Cohort Definition	Marriage Based			
	(1)	(2)	(3)	(4)
Cumulative Conditions	0.055 (0.063)	-0.286 (0.087)**	-0.544 (0.398)	0.393 (0.106)**
Year	0.004 (0.001)**	0.006 (0.001)**	-0.046 (0.003)**	0.006 (0.001)**
Initial Conditions	-0.089 (0.122)	-0.136 (0.162)	-0.502 (0.796)	-0.253 (0.230)
Prior Conditions	-0.134 (0.220)	-0.287 (0.242)	-2.027 (1.139)	0.359 (0.335)
Observations	1,073	1,073	1,073	1,073
R ²	0.407	0.471	0.562	0.242

Note: This table repeats column 4 of Table 3 for various subsamples. Standard errors, clustered at the cohort level, are in parentheses. “*” and “**” indicate significance at the 5% and 1% level, respectively. The economic measure over the course of marriage excludes the first year of marriage. The measure of economic conditions is log GDP growth. Cumulative Conditions is the sum of the economic measures over the marriage. Initial Conditions is the GDP growth in the 1st year of marriage; Prior Conditions is the GDP growth for the 3 Yrs. before marriage. Cohorts consist of married couples (with the husband in his first marriage, who do not subsequently divorce in the sample) who were married in the same year. The measure of income is labor income of the husband or wife, or the sum. Moments are calculated as the cross-sectional variance of the excess log income for the first three columns. Income is the total labor income of the husband and wife in column 1, the labor income of the husband in column 2, the labor income of the wife in column 3. Column 4 uses the cross-sectional covariance of husbands’ and wives’ excess log incomes, as described in the text. Regressions also include dummy variables for the number of years that the cohort has been married.

Table 5: Impact of past economic conditions on cross-sectional moments
Decomposing the cross-sectional variance of couples' incomes

Dependent Var.	Cross-Sectional Covariance of Husband's Income with Wife's Variable		
	Income	Hours	Wage
Economic Measure	GDP Growth		
Cohort Definition	Marriage based		
	(1)	(2)	(3)
Cumulative Conditions	0.393 (0.106)**	0.169 (0.091)	0.197 (0.058)**
Year	0.006 (0.001)**	0.004 (0.001)**	0.002 (0.001)**
Initial Conditions	-0.253 (0.230)	-0.266 (0.180)	0.036 (0.097)
Prior Conditions	0.359 (0.335)	0.267 (0.260)	0.028 (0.163)
Observations	1,073	1,073	1,069
R ²	0.242	0.169	0.151

Note: This table repeats column 4 of Table 3 for various subsamples. Standard errors, clustered at the cohort level, are in parentheses. “*” and “**” indicate significance at the 5% and 1% level, respectively. The economic measure over the course of marriage excludes the first year of marriage. The measure of economic conditions is log GDP growth. Initial Conditions is the GDP growth in the 1st year of marriage; Prior Conditions is the GDP growth for the 3 Yrs. before marriage. Cohorts consist of married couples (with the husband in his first marriage, who do not subsequently divorce in the sample) who were married in the same year. The dependent variable is the cross-sectional covariance of the excess log labor income of the head and wife (column 1), the cross-sectional covariance of the excess log labor income of the husband and the excess log hours of the wife (column 2), and the cross-sectional covariance of the excess log labor income of the husband and the excess log wage of the wife (column 3). Regressions also include dummy variables for the number of years that the cohort has been married.