The Co-Movement of Couples' Incomes

Stephen H. Shore*†‡

First Draft: July 15, 2006
This Draft: October 12, 2006

Abstract

How much do couples' incomes move together? In this paper, I estimate the parameters of a joint labor income process for couples. I exploit the relationship between short-term changes in wives' incomes and the long-term changes in their husbands' incomes that span them. The product of these changes cleanly estimates the covariance of a couple's permanent innovations to income without contamination from their correlated measurement errors. I develop a "wife-swap bootstrap" test that rejects non-parametrically the hypothesis that couples' incomes move independently.

Innovations to couples' permanent incomes have a correlation of -10 percent on average. Co-movement is driven in large part by changes in wives' labor force participation that are negatively correlated with changes in their husbands' incomes. This negative correlation is present throughout the range of income changes; when a wife's income increases, it indicates a shift downward in the whole distribution of her husband's income. There is excess variation in sample estimates of co-movement, which is explained by positive assortative mating in income volatility. While there is no evidence of any additional latent variation in co-movement, there is strong life-cycle variation. In particular, innovations to a husband's income and his wife's income are strongly negatively correlated early in marriage but become more positively correlated throughout marriage.

*University of Pennsylvania, Wharton School. Email: sshore@wharton.upenn.edu. Please contact author at: 3620 Locust Walk, 3012 Steinberg Hall-Dietrich Hall, Philadelphia, PA 19104; 215-898-7770.

†JEL Classification: D13 - Household Production and Intrahousehold Allocation; D31 - Personal Income, Wealth, and Their Distributions.

‡Preliminary draft. Comments very welcome. Do not cite without permission. I thank Shelly Lundberg, Shane Jensen, and Dylan Small, as well as seminar participants at the Wharton School applied economics seminar, for helpful comments. I thank Michael Pan and Seth Zeleznik for excellent research assistance. 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.

Contents

List of Figures

List of Tables

1 Introduction

Efforts to model individuals' income processes have focussed on measuring income volatility, the variance of changes to income. Research in this area is motivated by the desire to understand how much income risk (proxied by income volatility) people face, what affects the amount of such risk, and how it impacts behavior (e.g., precautionary saving). Here, I examine the joint income process for married couples. In particular, I estimate the co-movement of couples' incomes – the covariance of changes to husbands' incomes and changes in their wives' incomes.

In finance, co-movement is important because it identifies the degree to which diversification reduces risk. As long as husbands' exogenous shocks are not perfectly correlated with their wives' shocks, marriage helps to diversify labor income risk; the lower that correlation, the greater the diversification benefits of marriage. This diversification will be seen in the co-movement of spouses' wages or involuntary layoffs. Co-movement may also reflect joint endogenous choices. These coordinated changes are motivated by a wife's leisure being a complement with her husbands', or her production being a substitute for her husbands'. It will be seen in the co-movement of effort or hours worked. Finally, co-movement may reflect the endogenous response of one spouse to exogenous shocks to the other. One example is the added worker effect, the labor supply response of one spouse to their partner's unemployment (Lundberg, 1985; Cullen and Gruber, 2000). This effect will be seen in the co-movement of one spouse's effort or hours worked with the other spouse's wages or involuntary unemployment. Whether reflecting joint exogenous shocks (diversification), joint endogenous choices (coordination), or the endogenous response of one spouse to the exogenous shocks of another (added worker effect), the co-movement of couples' incomes identify economic benefits of marriage and their dynamics.

Differentiating permanent changes in income from transitory ones poses a challenge when estimating income volatility. The variance of permanent shocks is useful as a proxy for risk, while the variance of transitory shocks is problematic because it is observationally equivalent to measurement error. Identifying moments that measure the variance of permanent shocks without contamination from transitory ones has been a major innovation in recent research (Carroll and Samwick, 1997; Meghir and Pistaferri, 2004). In this paper, I develop an estimator to measure the covariance of couples' permanent income changes without contamination from correlated measurement error.

To estimate this permanent covariance, I examine the relationship between the short-term change in a wife's income (e.g., between last year and this year) and the long-term change in her husband's income that spans this short-term change (e.g., between three years ago and two years from now).

Both the long- and short-term changes incorporate permanent innovations to income over the short-term period. But if correlated measurement error is transitory enough, all other components of these long- and short-term changes will be uncorrelated. Under fairly weak assumptions, the product of these two changes then provides an unbiased estimate – for a couple at an instant in time – of the permanent covariance, the covariance of couples' permanent innovations to income.[1].

While this "long-short product estimator of the permanent covariance" can provide a reliable estimate of the permanent covariance given enough data, any single long-short product is an extremely noisy estimate of the permanent covariance. The large variation in estimates poses a challenge for hypothesis testing. Estimation error will be correlated over time, while not all parameters that determine the degree of this correlation are estimated. As a result, testing the null of no co-movement on average is not straightforward. And since the income process itself may be mis-specified, it is not clear how much variation in these product estimates would be expected in the absence of heterogeneity. As a result, it is also not straightforward to test for heterogeneity, as distinct from estimation error. To address this concern, I develop a "wife-swap bootstrap" to test the hypothesis that couples' incomes move independently. This block bootstrap randomization first pairs all husbands and wives at random and then computes long-short product estimates from the observed income series of these random pairs. Doing this repeatedly builds a reference distribution, the distribution of estimates that would be expected if husbands' and wives' incomes moved independently but exactly as they do individually in the data. This overcomes the problem that any test about co-movement for couples is actually a joint test of the hypothesis of interest and the hypothesis that income processes are specified properly. Using this test, I reject non-parametrically both the hypothesis that couples incomes move independently on average (i.e. that the mean of co-movement parameters is zero), and also the hypothesis that there is no heterogeneity.

I estimate the permanent covariance from the Panel Study of Income Dynamics (PSID). I

[1]In the variance analog to this long-short product estimator, Meghir and Pistaferri (2004) identify the variance of permanent income from the product of a short-term change in income and the longer-term change in income that spans this short-term change. Since transitory shocks that affect the short-term change will be uncorrelated with those that affect the long-term change, while permanent shocks affect both long- and short-term changes in income, this moment identifies the variance of innovations to permanent income and is unaffected by transitory income on average.

This idea is related to Stephens (2002), who examines the long-term effect on wives' employment of their husbands' job displacement, the loss of a job due to firing, lay-offs, plant closings in a given year. Relative to Stephens (2002), I map the comparison of couples' short- and long-term changes to a parameter of their joint income process and examine co-movement beyond what could be attributed to a husbands' job displacement.

Carroll and Samwick (1997) identify the average variance of innovations to permanent income from the difference between average squared long-term changes in income and average squared medium-term changes in income. Since both long- and medium-term changes in income are affected equally by the transitory shocks at either end of these differences, but long-term changes in income reflect more permanent shocks, this moment identifies the average lifetime variance of permanent shocks when transitory shocks are constant. This paper also examines the moment formed from the average product of couples' long-term changes in income minus the average product of couples' medium-term changes in income. This "long-short difference estimator of the average covariance" provides similar estimates.

compute predicted log income for husbands' and wives', Winsorize the residuals, and examine the time-series properties of couples' Winsorized residuals. I identify four major stylized facts:

First, innovations to permanent income are negatively correlated on average, with a correlation of roughly -10 percent. Short-term increases in wives' incomes tend to coincide with long-term declines in husbands' incomes on average. The wife-swap bootstrap confirms that this finding is highly statistically significant. This negative correlation is present when looking at the relationship between long-term changes in husbands' incomes and short-term changes in wives' incomes, wives' hours, wives' incomes conditional on being in the labor force, and wives' labor force participation.[2]

Second, the negative correlation between innovations to couples' incomes occurs throughout the income distribution. Quantile regressions reveal a negative relationship between short-term changes in a wife's income and the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles of the long-term changes in her husband's income that span them. When a wife's income increases in a given year, the entire distribution of the permanent component of her husband's income falls. Negative co-movement is not limited to coordination over extreme positive or negative shocks. A similar pattern is apparent when looking at the relationship between short-term changes in wives' labor force participation and the distribution of long-term changes in husbands' incomes that span them.

Third, the long-short product estimates of the permanent covariance exhibit more variation than would be expected if income followed the same process for all couples at all times. This excess variation implies that large absolute short-term changes in wives' incomes tend to coincide with large absolute long-term changes in their husbands' incomes. This could reflect positively correlated magnitudes of shocks, a *positive covariance of couples' variance moments*; alternatively, it could reflect heterogeneity in the degree of co-movement, a *positive variance in couples' covariance moments*. In the first case, there is regime-switching between a high-volatility regime for both spouses and a low-volatility regime for both spouses. In the second case, there is regime-switching between positively correlated shocks and negatively correlated shocks. These two possibilities are observationally equivalent without explanatory variables to predict which regime an observation will be in. I document that nearly all of this excess variation can be explained by assortative mating in income volatility, the tendency of husbands with high income volatility to be married to wives with

[2] This result differs from other research which finds a that couples' short-term changes in income are positively correlated. Unlike most other research (Hess, 2004), this paper removes predictable (to the econometrician) income changes and looks at the evolution of "excess" income. Since life-cycle patters and assortative mating cause the predictable component of couples' incomes to move together, failing to control for these predictable changes will bias estimates of couples' correlations upward. Also, this paper separates innovations to permanent income from innovations to transitory income. Transitory income shocks are positively correlated, perhaps reflecting correlated measurement error. Since only one spouse answers survey questions for the whole household in the PSID and most other household surveys, correlated measurement error should be expected. Standard measures of co-movement include this transitory covariance, which will bias estimates of co-movement upward.

high income volatility. There is no evidence of latent heterogeneity in couples permanent covariance moments once this assortative mating in income volatility is accounted for.

Fourth, there is a strong life-cycle pattern in co-movement. Couples have strongly negatively correlated innovations to permanent income early in marriage (-30 percent) but positively correlated innovations to permanent income later in marriage (10 percent). This life-cycle variation is apparent in the relationship between long-term changes in husbands' incomes and short-term changes in wives' incomes, wives' hours, wives' incomes conditional on being in the labor force, and wives' labor force participation. This pattern may reflect changes over time in the relative importance of spouses' production being substitutes (giving rise to negative co-movement and likely important early in life) and their leisure being compliments (giving rise to positive co-movement and likely important nearing retirement).

The remainder of the paper is organized as follows: Section 2 presents data and methods; Section 3 presents average results and tests for excess variation from the average; Section 4 decomposes this excess variation; Section 5 concludes. More detail about subsections, tables, and figures is given in the Table of Contents.

2 An Income Process for Couples

Here, I write down a joint process for the log excess incomes of husbands and wives. The income process I apply to either spouse is very standard; with slight variants, it has been used in many papers (e.g., Hall and Mishkin (1982); Carroll and Samwick (1997); Meghir and Pistaferri (2004)). Most of these use this income process to model the evolution of total household income. The novelty here is to write down a *joint* process to model the evolution of husbands' and wives' incomes. I then develop estimators to identify co-movement, the covariance of couples' innovations to income.

2.1 Data

Data for this paper is drawn from the Panel Study of Income Dynamics (PSID). The PSID is a nationally representative panel of U.S. households that has tracked families since 1968. It includes annual data on households, including the education, income, hours worked, employment status, and age of husbands and wives. I restrict the sample to married couples, to observations for which both the husband and wife are between the ages of 22 and 60, and for which the couple has been married

Table 1: Regressions Predicting Log Income for Husbands and Wives

Dependent Variable: Log Income of Spouse

Spouse	Husband	Wife
# of years married	0.0490 (14.40)**	-0.0378 (7.10)**
(# of years married)^2	-0.0020 (8.68)**	0.0018 (4.81)**
(# of years married)^3	0.0000 (6.76)**	0.0000 (3.62)**
Total Family Size	-0.0303 (5.74)**	0.0200 (2.22)*
Any Kids Age 0-1?	0.0400 (1.24)	-0.2260 (3.46)**
# of Kids Age 0-1	0.0026 (0.09)	-0.0278 (0.44)
Any Kids Age 2-6?	0.0252 (2.14)*	0.0140 (0.64)
# of Kids Age 2-6	0.0143 (1.52)	-0.2074 (11.78)**
Any Kids Age 7-20?	-0.0126 (1.44)	-0.0557 (3.79)**
# of Kids Age 7-20	0.0206 (3.44)**	-0.0955 (9.23)**
Year Dummies?	Yes	Yes
(HD & WF educ. Dummy) x (HD & WF age cubic)?	Yes	Yes
Observations	55,487	39,619
R^2	0.497	0.402

Data comes from the PSID, 1968-2001. This table presents the results of an OLS regression of the log labor income of the husband or wife on the covariates shown. The sample is limited to couples with both spouses between the ages of 22 and 60, and to observations for which the couple has been married for no more than 35 years. Observations are excluded when the spouse in question reports an annual labor income of less than $\$1,000$ (2001 dollars). Education categories are: none, elementary, junior high, some high school, high school, some college, college, and graduate school. The education level is set to the maximum for that individual in the sample when it is missing. Absolute value of t-statistics are in parentheses. "*" indicates significance at the 5% level; "**" indicates significance at the 1% level.

for no more than 35 years.[3] I use the annual labor income as a measure of income.

As the measure of "excess" log income, I use the residual from a regression of the natural log of labor income (for either the husband or the wife) on host of regressors. So that log income results are not dominated by income values close to zero, I limit the regression sample to individuals who earn at least $\$1,000$ (in 2001 dollars). In each regression, I use as regressors: a cubic in age for each level of educational attainment (none, elementary, junior high, some high school, high school, some college, college, graduate school) for both husband and wife, a cubic in the number of years the couple has been married, the presence and number of infants, young children, and older children in the household, the total number of family members in the household, and dummy variables for each calendar year. Including calendar year dummy variables eliminates the need to convert nominal income to real income explicitly. Table 1 presents the coefficients on these regressors, with coefficients on year dummy variables and the (education) x (age cubic) omitted for parsimony.

The residuals from this regression are Winsorized at the 5^{th} and 95^{th} percentiles, so that residuals below the 5^{th} percentile are replaced by the 5^{th} percentile value and those above the 95^{th} percentile are replaced by the 95^{th} percentile value. At the same time, values omitted from the initial regression because real annual income was below $\$1,000$ are given the 5^{th} percentile residual value. The vast majority of these initially omitted values have an income of exactly zero. This reduces selection bias by including extreme values, while at the same time limiting the degree to which such outlier drive the results. Even more important, it allows us to exploit variation coming from transitions into and out of the labor force. While Winsorizing below some level (here, the 5^{th} percentile, though the exact level will not affect the results) is important for these reasons, Winsorizing above the 95^{th} percentile is done for symmetry and has no substantive impact on the results. Results are unaffected qualitatively by adjusting the minimum income value of $\$1,000$. By Winsorizing after computing residuals and not before, the change in income is zero by construction for any individual who remains below the minimum excess income level or above the maximum excess income level. Thus results cannot be driven by changes over time in the Winsorizing bound.[4]

Table 2 presents data on the distribution of Winsorized excess log income. I use $y^s_{i,t}$ to refer to the (Winsorized excess) log labor income for spouse $s \in \{H \text{ (husband)}, W \text{ (wife)}\}$ in couple i in year t. The first two columns show the distribution of excess log income for husbands and wives,

[3] Results are robust to excluding any couple who subsequently got divorced in the sample, or to couples who remain in the sample for relatively few years.

[4] The drawback of this order is that neither the mean value of the residuals nor their changes will be equal to zero, so that squared values will not equal variances exactly. This adjustment is second-order; results are qualitatively identical and quantitatively similar with or without correcting for this explicitly.

Table 2: Distribution of Winsorized Excess Log Income for Husbands and Wives

Excess Log Income	Level		One-Year Change		Five-Year Change	
Variable	y^s		$y^s_{i,t} - y^s_{i,t-1}$		$y^s_{i,t+2} - y^s_{i,t-3}$	
Spouse (s)	Husband	Wife	Husband	Wife	Husband	Wife
Mean	-0.0823	-0.6033	-0.0125	0.0004	-0.0431	0.0400
St. Dev.	0.5520	1.0345	0.3409	0.5976	0.4751	0.9389
Observations	61,198	61,198	53,439	53,439	39,976	39,976
Minimum	-1.0399	-1.7265	-1.8514	-2.7934	-1.8514	-2.7934
1^{st} Percentile	-1.0399	-1.7265	-1.2197	-2.0872	-1.5074	-2.5264
5^{th} Percentile	-1.0399	-1.7265	-0.5757	-1.0443	-0.9368	-1.8600
25^{th} Percentile	-0.4365	-1.7265	-0.1011	-0.0796	-0.2294	-0.2026
50^{th} Percentile	0.0050	-0.5203	0	0	0	0
75^{th} Percentile	0.3248	0.3675	0.0965	0.0938	0.1795	0.3440
95^{th} Percentile	0.7857	0.9612	0.4984	1.0107	0.6933	1.8075
99^{th} Percentile	0.8115	1.0669	0.9973	1.9187	1.2320	2.4485
Maximum	0.8115	1.0669	1.8514	2.7934	1.8514	2.7934

This table presents the distribution of the residuals obtained from the regression detailed in Table 1. Those residuals were Winsorized at the 5 percent and 95 percent level, with observations for real annual labor income below $\$1,000$ (which were excluded from the regression in Table 1) given the 5 percent Winsorized value. The first two columns of present the distribution of Winsorized excess log income for the husband and wife, respectively. The next two columns present the distribution of one-year changes for husbands and wives. The last two columns present the distribution of five-year changes. A value of zero for the change moments indicates that income remained below the 5 percent level or above the 95 percent level in both years over which the difference was taken.

y^H and y^W. The mean of excess log income is slightly negative, -0.08, for husbands and strongly negative, -0.60, for wives. Individuals with income below $\$1,000$ are excluded from the regression but assigned the 5^{th} percentile Winsorized value, which explains why the mean Winsorized residuals are negative. The inter-quartile range of excess log income is -0.44 to 0.32 for husbands, which with a log-linearized approximation corresponds to income increasing 76 percent when moving from the 25^{th} percentile to the 75^{th} percentile. The inter-quartile range is -1.73 to 0.37 for wives; roughly one-third of women report little or no labor income and therefore have an excess log income at the lower bound of -1.73. The third and fourth columns present the distribution of one-year changes in excess log income for husbands and wives. Naturally, the mean of one-year changes in y is close to zero. The inter-quartile range of one-year changes for husbands is -0.10 to 0.10 and is -0.08 to 0.09 for wives; excess income does not change more than 10 percent from year to year for most individuals. However, there are extreme changes in income, so the standard deviation of changes to income is 0.34 for husbands and 0.60 for wives. The fifth and sixth columns repeat the results for the third and fourth columns, but present five-year changes instead of one-year changes. These long-term changes have only slightly higher standard deviations of 0.48 for husbands and 0.94 for

Table 3: Autocorrelation of Changes in Excess Log Income for Husbands and Wives

Autocorrelation for Husbands

correlation	$y^H_{i,t} - y^H_{i,t-1}$
$y^H_{i,t-1} - y^H_{i,t-2}$	-0.2752
$y^H_{i,t-2} - y^H_{i,t-3}$	-0.0554
$y^H_{i,t-3} - y^H_{i,t-4}$	-0.0322
$y^H_{i,t-4} - y^H_{i,t-5}$	-0.0189

Autocorrelation for Wives

correlation	$y^W_{i,t} - y^W_{i,t-1}$
$y^W_{i,t-1} - y^W_{i,t-2}$	-0.2044
$y^W_{i,t-2} - y^W_{i,t-3}$	-0.0726
$y^W_{i,t-3} - y^W_{i,t-4}$	-0.0437
$y^W_{i,t-4} - y^W_{i,t-5}$	-0.0361

Autocorrelation of one-year changes in husbands' and wives' incomes at one- through four-year lags. One-year changes in income are the Winsorized excess log incomes shown in Table 2.

wives, suggesting some mean-reversion in income.[5]

Mean-reversion in income can also be seen in Table 3, which shows the autocorrelation of $y^s_{i,t} - y^s_{i,t-1}$, the autocorrelation in one-year changes in husband's and wives' excess log incomes incomes. One-year auto-correlation in income is strongly negative (-0.28 and -0.20, respectively, for husbands and wives); income tends to increase in the year following an decrease. This negative autocorrelation between consecutive one-year changes in income is a central feature of data on earnings dynamics, and any process for labor income must accommodate this feature of the data. Auto-correlation at longer lags is negative but small and decreasing. The non-zero values for higher-order lags are driven primarily by the inclusion of individuals without income. Such individuals leave the labor force (income falls) and then subsequently re-enter the labor force (income rises). Excluding these observations eliminates higher-order autocorrelation entirely for women and mostly for men. The next two sections develop a process for labor income to map these changes in income to parameters of a labor income process.

2.2 Individual Income Process

Here, I characterize a standard single income process for $y^s_{i,t}$, the excess log income for spouse s in couple i at time t.[6] I assume that excess log income is composed of permanent (p) and transitory (ε) components:

$$\begin{aligned} y^s_{i,t} &= p^s_{i,t} + \varepsilon^s_{i,t}; \\ p^s_{i,t} &= p^s_{T^i_0} + \sum_{\tau=T^i_0+1}^{t} \omega^s_{i,\tau}. \end{aligned} \tag{1}$$

[5] If changes to income were uncorrelated, the standard deviation would increase with the square of the length of the sample period, in this case by a factor of $\sqrt{5} \approx 2.2$ from one- to five-year changes.

[6] This process could apply to log total household income, $s \in \{S \text{ (sum)}\}$, or to the log income of each spouse, $s \in \{H(\text{husband}), W(\text{wife})\}$. However, any non-degenerate version of this income process cannot simultaneously apply to the husband's log income, the wife's log income, and the log of total household income, as the income process is additive in *logs* while household income is the sum of the *levels* of the spouses' incomes.

Transitory income, $\varepsilon_{i,t}^{s}$, is i.i.d. with variance $\left(\sigma_{i,t}^{\varepsilon,s}\right)^2$; permanent income, $p_{i,t}^{s}$, has a unit root so that innovations to permanent income, $p_{i,t}^{s} - p_{i,t-1}^{s} = \omega_{i,t}^{s}$, are i.i.d. with variance $\left(\sigma_{i,t}^{\omega,s}\right)^2$.[7,8] Subsequently, "transitory variance" refers to the variance of transitory income, $\left(\sigma_{i,t}^{\varepsilon,s}\right)^2$; "permanent variance" refers to the variance of innovations to permanent income, $\left(\sigma_{i,t}^{\omega,s}\right)^2$. A slightly more general process in the same class is:

$$
\begin{aligned}
y_{i,t}^{s} &= p_{i,t}^{s} + \sum_{\tau=t-\epsilon+1}^{t} \phi_{t-\tau}\varepsilon_{i,\tau} \\
p_{i,t}^{s} &= p_{T_0^i}^{s} + \sum_{\tau=0}^{t-\Omega} \omega_{i,\tau}^{s} + \sum_{\tau=t-\Omega+1}^{t} \theta_{t-\tau}\omega_{i,\tau}^{s}.
\end{aligned} \tag{2}
$$

Here, permanent shocks come into effect over Ω periods (with the vector θ representing the rate at which this happens), and transitory shocks fade completely after ϵ periods (with the vector ϕ representing the rate at which they fade). Equation 2 reduces to 1 when $\epsilon = 1$ and $\Omega = 0$, with ϕ_0 set to 1 without loss of generality.

2.3 Identifying Volatility, σ^2

Here, I briefly discuss three types of estimators of income volatility: a raw estimator of variance, a product ($\times$) estimator of permanent variance, and a difference ($-$) estimator of average permanent variance. These are summarized in the top panel of Table 5. I follow standard notation in placing a hat, "$\wedge$", above a parameter to indicate a sample estimator of that parameter. To the right of such an expression, I describe the estimator in angular brackets, "$\langle\ \ \rangle$"; the subscript within these

[7] Transitory income has an individual-specific mean (set to zero by adjusting initial permanent income without loss of generality) and is uncorrelated across time, $E\left[\varepsilon_{i,t}^{s}\right] = E\left[\varepsilon_{i,t}^{s}\varepsilon_{i,r}^{s}\right] = 0$ for all $t \neq r$. It has a finite variance, $E\left[\left(\varepsilon_{i,t}^{s}\right)^2\right] = \left(\sigma_{i,t}^{s}\right)^2$. Innovations to permanent income $p_{i,t}^{s} - p_{i,t-1}^{s} = \omega_{i,t}^{s}$ are mean zero and uncorrelated across time, $E\left[\omega_{i,t}^{s}\right] = E\left[\omega_{i,t}^{s}\omega_{i,r}^{s}\right] = 0$ for all $t \neq r$. Innovations to permanent income also have a finite variance, $E\left[\left(\omega_{i,t}^{s}\right)^2\right] = \left(\sigma_{i,t}^{\omega,s}\right)^2$. $p_{T_0^i}^{s}$ is the initial permanent income, the permanent income in year T_0^i.

In practice, $E\left[\omega_{i,t}^{s}\right]$ and $E\left[\varepsilon_{i,t}^{s}\right]$ will not be mean-zero. The Winsorizing procedure and attrition allows expected changes in income to differ from zero. Adjusting for this is trivial if one is willing to de-mean excess log income in each year or Winsorize income prior to the predictive regression (while excluding attriting data from the predictive regression). However, in doing so we lose the appealing interpretation of $y_t - y_{t-m} = 0$ as a case where income remained above or below the Winsorizing bound over the interval. This would leave open the possibility that results are driven by changes over time in the level at which Winsorizing binds. In practice, the impact of such adjustments is tiny, and has no impact on the substance of the results. OLS regressions presented in the paper implicitly de-mean while other techniques in the paper do not; the two approaches recover qualitatively identical and quantitatively nearly identical results.

[8] While this model provides a useful stylized process for income, income does not evolve exactly according to this process. Most obviously, labor income will be zero – and log income undefined – when an individual is not in the labor force; transitions into and out of employment the labor force are not modeled. While Winsorizing allows zeros from the data to be accommodated, the income process does not account for them explicitly. Subsequent data analysis will examine changes in labor force participation, transitions into and out of this zero state, explicitly.

brackets denotes the type of estimator (raw, product, or difference) while the superscript denotes relevant individuals or intervals used for this estimator. These variance estimators are neither novel nor the focus of this paper. However, I discuss them before introducing their covariance analogs, which are the focus of this paper. Also, estimates of variance are necessary if covariance estimates are to be scaled into correlations.

The simplest estimator of income volatility is merely the squared change in income over a given interval, $\left(y^s_{i,t} - y^s_{i,t-m}\right)^2$. I refer to this as the "*raw* estimator of variance," since it makes no effort to separate permanent variance from transitory variance, or to account for θ and ϕ.

Given the labor income process in equation 1 (or equation 2 for $m, n \geq \Omega,\epsilon$), Meghir and Pistaferri (2004) propose the following moment to estimate the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$:

$$E\left[\left(y^s_{i,t+n} - y^s_{i,t-1-m}\right)\left(y^s_{i,t} - y^s_{i,t-1}\right)\right] = \left(\sigma^{\omega,s}_{i,t}\right)^2 .^{9} \tag{3}$$

This moment presents the product of a short-term change in income, $y^s_{i,t} - y^s_{i,t-1}$, and the long-term change in income that surrounds it, $y^s_{i,t+n} - y^s_{i,t-1-m}$. For this reason, I refer to its sample analog,

$$\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_\times\rangle \equiv \left(y^s_{i,t+n} - y^s_{i,t-1-m}\right)\left(y^s_{i,t} - y^s_{i,t-1}\right), \tag{4}$$

as the "*long-short product* estimator of permanent variance." The subscripted "$\times$" in brackets, a multiplication sign, refers to this product estimator. By assumption ($m, n \geq \Omega,\epsilon$), the transitory conditions at the beginning and end of the long-term change are uncorrelated with the transitory conditions at the beginning and end of the short-term change. As a result, the expectation of the product of long- and short-term changes will be unaffected by the transitory variance, $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$. Since both the short- and long-term change in income share the same permanent shock between years $t-1$ and t, the expectation of $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_\times\rangle$ will reflect the permanent variance over this interval, $\left(\sigma^{\omega,s}_{i,t}\right)^2$.[10]

For the expectation of $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_\times\rangle$ to recover the permanent variance, the key identifying assumption is that permanent shocks come fully into effect after Ω years and transitory shocks have faded fully after ϵ years. This assumption has been tested at length for processes for household income and is generally accepted for $\Omega, \epsilon \geq 1$ (Abowd and Card, 1989). For this reason, $m = n = 2$ is used throughout this paper.[11]

[9] This assumes that $\left(\sigma^{\omega,s}_{i,\tau}\right)^2 = \left(\sigma^{\omega,s}_{i,t}\right)^2$ for all $t - \Omega \leq \tau \leq t$ and that $E\left[\varepsilon^s_{i,\tau}\omega^s_{i,\tau}\right]$ is constant for all $t - \epsilon \leq \tau \leq t$.

[10] When the permanent variance is constant and the covariance of permanent and transitory shocks is constant, $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_\times\rangle$ provides an unbiased estimate of the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$. When the permanent variance changes over time, $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_\times\rangle$ provides a weighted average of recent permanent variances.

[11] However, this choice is undercut here by the higher-order autocorrelation induced by the Winsorizing procedure

Table 4: Distribution of Individual Income Volatility Estimates

Distribution of Husbands' Volatility Estimates

Type of Estimates	Long-Short Product Estimates of the Permanent Variance	Raw Estimates of Variance		
Estimator	$\left(\hat{\sigma}_{i,t}^{\omega,H}\right)^2 \langle\times\rangle$	$\left(y_{i,t}^H - y_{i,t-1}^H\right)^2$	$\left(y_{i,t}^H - y_{i,t-3}^H\right)^2$	$\left(y_{i,t}^H - y_{i,t-5}^H\right)^2$
Mean	0.0293	0.1060	0.1882	0.2308
St. Dev.	0.2275	0.3175	0.4407	0.4784
Observations	38,120	36,678	36,678	36,678

Distribution of Wives' Volatility Estimates

Type of Estimates	Long-Short Product Estimates of the Permanent Variance	Raw Estimates of Variance		
Estimator	$\left(\hat{\sigma}_{i,t}^{\omega,W}\right)^2 \langle\times\rangle$	$\left(y_{i,t}^W - y_{i,t-1}^W\right)^2$	$\left(y_{i,t}^W - y_{i,t-3}^W\right)^2$	$\left(y_{i,t}^W - y_{i,t-5}^W\right)^2$
Mean	0.1430	0.3109	0.6566	0.9070
St. Dev.	0.7967	0.8815	1.3854	1.6371
Observations	38,120	36,678	36,678	36,678

Estimators are formed from the changes in Winsorized excess log income, presented in Table 2. The first column of numbers presents the distribution of long-short product estimates of the permanent variance, $\left(\hat{\sigma}_{i,t}^{\omega,s}\right)^2 \langle\times\rangle \equiv \left(y_{i,t+n}^s - y_{i,t-1-m}^s\right)\left(y_{i,t}^s - y_{i,t-1}^s\right)$, from equation 4 following Meghir and Pistaferri (2004). The second, third, and fourth columns present distribution of the raw estimates of variance, $\left(y_{i,t}^s - y_{i,t-m}^s\right)^2$, for $m = 1$, 3, and 5. So that differences across m do not reflect sample selection, estimtes in columns 2, 3, and 4 are only included when $\left(y_{i,t}^s - y_{i,t-m}^s\right)^2$ is non-missing for all values of s and m for given values of i and t.

Column 1 of Table 4 presents the mean and standard deviation of long-short product estimates of the permanent variance, $\left(\hat{\sigma}_{i,t}^{\omega,s}\right)^2 \langle\times\rangle$, for both husbands and wives. Means of $\left(\hat{\sigma}_{i,t}^{\omega,s}\right)^2 \langle\times\rangle$ provides plausible estimates of the permanent variances, 0.029 for husbands (17 percent annual standard deviation of permanent changes in income) and 0.143 for wives (38 percent annual standard deviation). While the inter-quartile ranges are modest (not shown in table, -0.008 to 0.028 and 0 to 0.042, respectively), variation in $\left(\hat{\sigma}_{i,t}^{\omega,s}\right)^2 \langle\times\rangle$ is enormous, with standard deviations of 0.23 and 0.80, respectively. This variation will reflect both heterogeneity in the volatility parameter, $\left(\sigma_{i,t}^{\omega,s}\right)^2$, and noise in estimating it, $\left(\hat{\sigma}_{i,t}^{\omega,s}\right)^2 \langle\times\rangle - \left(\sigma_{i,t}^{\omega,s}\right)^2$; differentiating between the two is the focus of Meghir and Pistaferri (2004) and (with a more general income process) Alvarez, Browning, and Ejrnaes (2001).

Given the same process for income, Carroll and Samwick (1997) propose the following moment

and documented in Table 3. In particular, permanent income appears to be slightly less than permanent, reflecting occasional entry into and exit from the labor force that is excluded from most other studies. This problem will not affect estimates of the permanent covariance given the assumptions described and tested in Subsection 2.4, and will have only a second-order impact on the mapping from permanent covariances to permanent correlations.

to estimate the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$, for $m > n > \Omega, \epsilon$:

$$E\left[\left(y^s_{i,t_1} - y^s_{i,t_1-m}\right)^2\right] - E\left[\left(y^s_{i,t_2} - y^s_{i,t_2-n}\right)^2\right] = (m-n)\left(\sigma^{\omega,s}_{i,t}\right)^{2\,12} \tag{5}$$

This moment presents the difference between squared long-term changes in income and squared shorter-term (but $n > \Omega, \epsilon$) changes in income. For this reason, I refer to its sample analog,

$$\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \left\langle{}^{mn}_{-}\right\rangle \equiv \left(y^s_{i,t_1} - y^s_{i,t_1-m}\right)^2 - \left(y^s_{i,t_2} - y^s_{i,t_2-n}\right)^2 \tag{6}$$

as the "*long-short difference* estimator of the permanent variance". The subscript "−" in brackets, a subtraction sign, refers to this difference estimator; the superscript numbers in brackets, mn, refer to the pair of interval lengths, m and n, over which squared changes are calculated. As long as points are far enough apart, increasing the interval length will increase the number of permanent shocks faced over that interval but leave constant at two the number of transitory shocks, one at the beginning of the interval and one at the end.[13] As a result, differences between long- and shorter-term squared changes (which contain the same number of transitory variances, while the long-term squared changes include more permanent variances) reveal the permanent variance.

The drawback of this difference estimator, $\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \left\langle{}^{mn}_{-}\right\rangle$, is that its constituent parts, $\left(y^s_{i,t_1} - y^s_{i,t_1-m}\right)^2$ and $\left(y^s_{i,t_2} - y^s_{i,t_2-n}\right)^2$, cannot span identical time-periods. Therefore, the expectation of the long-short difference estimator in equation 6 will not generally recover a useful parameter when volatility parameters vary. For this reason, $\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \left\langle{}^{mn}_{-}\right\rangle$ should not be used to estimate variation over time in $\left(\sigma^{\omega,s}_{i,t}\right)^2$. Instead, authors typically use lifetime (or full-sample) averages of the raw estimates of variance, $\left(y^s_{i,t} - y^s_{i,t-m}\right)^2$, over all t for all possible values of m to obtain an average estimate of $\left(\sigma^{\omega,s}_{i,\cdot}\right)^2$ for household i or the population as a whole:

$$\begin{aligned} mean_i\left[\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \left\langle{}^{mn}_{-}\right\rangle\right] &\equiv mean_i\left[\left(y^s_{i,t} - y^s_{i,t-m}\right)^2\right] - mean_i\left[\left(y^s_{i,t} - y^s_{i,t-n}\right)^2\right] \text{ or} \\ mean\left[\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \left\langle{}^{mn}_{-}\right\rangle\right] &\equiv mean\left[\left(y^s_{i,t} - y^s_{i,t-m}\right)^2\right] - mean\left[\left(y^s_{i,t} - y^s_{i,t-n}\right)^2\right].^{14} \end{aligned} \tag{7}$$

---

[12] This assumes that $\left(\sigma^{\omega,s}_{i,t}\right)^2$, $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$, and $E\left[\varepsilon^s_{i,t}\omega^s_{i,t}\right]$ are time-invariant.

[13] Formally, far enough part means that transitory shocks at the beginning and end of the period are uncorrelated (*i.e.* $n > \epsilon$) and permanent shocks at the beginning of the period have come fully into effect by the end of the period ($n > \Omega$).

[14] This is typically operationalized by calculating the raw estimator of variance, $\left(y^s_{i,t} - y^s_{i,t-m}\right)^2$, for all possible combinations of t and $m > 2$, then regressing these sample moments on m. When parameters are constant over time and income follows the stylized process in equation 1 (but not the more general process in equation 2), the

Columns 2, 3, and 4 of Table 4 present the distribution of raw estimates of variance, $\left(\hat{\sigma}^{s}_{i,t}\right)^2 \langle^{m}_{raw}\rangle$, for various interval lengths, $m = 1$, 3, and 5. Equation 5 implies when $\Omega, \epsilon < 3$, the difference estimator provides an estimate of the average permanent variance, $mean\left[\left(\hat{\sigma}^{\omega,H}_{i,\cdot}\right)^2 \langle^{53}_{-}\rangle\right] = 0.0213$ and $mean\left[\left(\hat{\sigma}^{\omega,W}_{i,\cdot}\right)^2 \langle^{53}_{-}\rangle\right] = 0.1252$, respectively for husbands and wives. These estimates are similar to but modestly lower than the average long-short product estimates of the permanent variance, $mean\left[\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle_{\times}\rangle\right]$, shown in column 1 of Table 4.

2.4 Identifying Co-Movement of Couples' Incomes, δ

The goal of this paper is to understand the co-movement of couples' incomes. Formally, I estimate the covariance of innovations to couples' permanent incomes, $E\left[\omega^H_{i,t}\omega^W_{i,t}\right] \equiv \delta^{\omega}_{i,t}$, which I subsequently refer to as the "permanent covariance." I aim to separate this permanent covariance from the covariance of couples' transitory incomes, $E\left[\varepsilon^H_{i,t}\varepsilon^W_{i,t}\right] \equiv \delta^{\varepsilon}_{i,t}$, which I will refer to as the "transitory covariance." Panel data on couples' incomes, including data from the PSID, is typically obtained from a single survey asked of only one member of a household about the entire household. As a result, correlated measurement error in couples' reported incomes is likely. While economically uninteresting, such correlated measurement error presents a problem for the econometrician because it is observationally equivalent to the transitory covariance. Transitory income risk will have a much smaller impact on expected lifetime utility than permanent income risk with a similar variance, so transitory co-movement is arguably less interesting even in the absence of correlated measurement error. Because the transitory covariance may be of less interest and is impossible to differentiate from measurement error, I focus on the permanent covariance. Here, I present three types of co-movement estimators: a raw estimator of covariance, a product ($\times$) estimator of the permanent covariance, and a difference ($-$) estimator of the average permanent covariance. These are summarized in the bottom panel of Table 5. Each of these covariance estimators has a variance analog that was introduced in Subsection 2.3 and summarized in the top panel of Table 5.

The simplest estimator of co-movement – and the one used in other research (e.g., Hess (2004))

---

intercept from this regression recovers an unbiased estimate of $2\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$. The coefficient from this regression reveals an estimate of $\left(\sigma^{\omega,s}_{i,t}\right)^2$. Provided that variation in $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$ is either time-invariant or linear in time, sample averages will recover an unbiased estimate of the average value of $\left(\sigma^{\omega,s}_{i,t}\right)^2$. Since this average estimate places a negative weight on $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$ for t at the beginning and end of the sample but a positive weight on $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$ in the middle of the sample, the moment will be biased downward (upward) for individuals whose transitory income volatility is high (low) early and late in life but lower (higher) in middle age. As a result, this average estimator will be biased ($E\left[mean_i\left[\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2 \langle^{mn}_{-}\rangle\right]\right]$ will generally not equal $mean_i\left[\left(\sigma^{\omega,s}_{i,t}\right)^2\right]$) unless unless the transitory variance is an affine function of time.

Table 5: Summary of Estimators

Individual Variance (Volatility) Estimators

Estimator Name	Definition	Eq'n #	Graphic
Raw Estimator of Variance	$\left(y^s_{i,t} - y^s_{i,t-m}\right)^2$	(N/A)	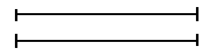
Long-Short Product Estimator of the Permanent Variance	$\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2 \langle{}_{\times}\rangle \equiv \left(y^s_{i,t+n} - y^s_{i,t-1-m}\right) \times \left(y^s_{i,t} - y^s_{i,t-1}\right)$	(4)	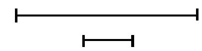
Long-Short Difference Estimator of the Permanent Variance	$\left(\hat{\sigma}^{\omega,s}_{i,.}\right)^2 \langle{}^{mn}_{-}\rangle \equiv \left(y^s_{i,t_1} - y^s_{i,t_1-m}\right)^2 - \left(y^s_{i,t_2} - y^s_{i,t_2-n}\right)^2$	(6)	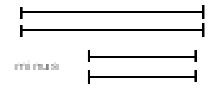

Couples' Coariance (Co-Movement) Estimators

Estimator Name	Definition	Eq'n #	Graphic
Raw Estimator of Covariance	$\left(\hat{\delta}_{i,t}\right)\langle{}^{m}_{raw}\rangle \equiv \left(y^H_{i,t} - y^H_{i,t-m}\right) \times \left(y^W_{i,t} - y^W_{i,t-m}\right)$	(8)	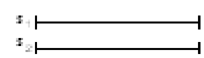
Long-Short Product Estimator of the Permanent Covariance	$\left(\hat{\delta}^{\omega}_{i,t}\right)\langle{}^{s_1 s_2}_{\times}\rangle \equiv \left(y^{s_1}_{i,t+n} - y^{s_1}_{i,t-1-m}\right) \times \left(y^{s_2}_{i,t} - y^{s_2}_{i,t-1}\right)$	(11)	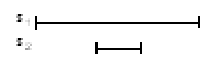
Long-Short Difference Estimator of the Permanent Covariance	$\left(\hat{\delta}^{\omega}_{i,.}\right)\langle{}^{mn}_{-}\rangle \equiv \left(\hat{\delta}_{i,t_1}\right)\langle{}^{m}_{raw}\rangle - \left(\hat{\delta}_{i,t_2}\right)\langle{}^{n}_{raw}\rangle$	(13)	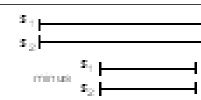

See relevant equations and discussion in text for details.

– is the product of concurrent changes in couples' incomes over a given interval,

$$\left(\hat{\delta}_{i,t}\right)\langle{}^{m}_{raw}\rangle \equiv \left(y^H_{i,t} - y^H_{i,t-m}\right)\left(y^W_{i,t} - y^W_{i,t-m}\right) \tag{8}$$

The subscript "raw" refers to this "*raw* estimator of covariance"; the superscript m in brackets denotes the interval length over which the concurrent changes are taken. Even given the stylized income process in equation 1,

$$E\left[\left(y^H_{i,t} - y^H_{i,t-m}\right)\left(y^W_{i,t} - y^W_{i,t-m}\right)\right] = \sum_{\tau=t-m+1}^{t} \delta^{\omega}_{i,\tau} + \delta^{\varepsilon}_{i,t} + \delta^{\varepsilon}_{i,t-m}; \tag{9}$$

raw estimates contain both transitory and permanent covariances. Given the more realistic income process in equation 2, the raw estimates will be impacted by the rate at which correlated transitory shocks damp out, ϕ, or correlated permanent shocks come into effect, θ.

Here, I propose an estimator of the permanent covariance, $\delta^{\omega}_{i,t}$, that avoids contamination from $\delta^{\varepsilon}_{i,t}$, θ, and ϕ. Broadly speaking, I examine the (bivariate) joint distribution of $y^{s_1}_{i,t+n} - y^{s_1}_{i,t-m-1}$ and $y^{s_2}_{i,t} - y^{s_2}_{i,t-1}$, the short-term changes in spouse s_2's income and the long-term changes in spouse s_1's income that span these short-term changes. As long as the correlated component of transitory income is not too persistent, the product of these two changes – equivalent to the covariance of these two changes when both are zero in expectation – reveals an unbiased estimate of the permanent covariance. When couples' incomes both follow the process detailed in equation 2 (for $m, n \geq \Omega, \epsilon$), the following moment identifies the permanent covariance, $\delta^{\omega}_{i,t}$:

$$E\left[\left(y^{s_1}_{i,t+n} - y^{s_1}_{i,t-m-1}\right)\left(y^{s_2}_{i,t} - y^{s_2}_{i,t-1}\right)\right] = \delta^{\omega}_{i,t}.^{15} \tag{10}$$

Its sample analog, the "*long-short product* estimator of the permanent covariance,"

$$\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{s_1 s_2}_{\times}\right\rangle \equiv \left(y^{s_1}_{i,t+n} - y^{s_1}_{i,t-m-1}\right)\left(y^{s_2}_{i,t} - y^{s_2}_{i,t-1}\right), \tag{11}$$

is merely the product of the short-term change in a wife's (or husband's) excess long income and the long-term change in her husband's (or his wife's) income that spans it. As before, the subscript "$\times$", a multiplication sign, refers to this product estimator. The superscript pair of letters, $s_1 s_2 \in \{HW, WF\}$, denote whether long-term changes are taken for the husband or the wife; the long-term change in income is taken for the first spouse in the pair, s_1, with the short-term change taken from the second spouse in the pair, s_2. The difference between spouse s_1's income in years $t-m-1$ and $t+n$ is merely his (or her) transitory income in year $t+n$ minus his transitory income in year $t-m-1$ plus the $m+n$ shocks to permanent income he received between years $t-m-1$ and $t+n$. The difference between the spouse s_2's income in years $t-1$ and t is merely her (or his) transitory income in year t minus her transitory income in year $t-1$ plus the permanent income shock she received between years $t-1$ and t. Transitory conditions for one spouse in years $t+n$

[15] This is the covariance analog of the long-short product used by Meghir and Pistaferri (and shown in equation 3) to estimate the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$.

This assumes that $E\left[\omega^{s_1}_{i,\tau}\omega^{s_2}_{i,\tau}\right] \equiv \delta^{\omega}_{i,t}$ for all $t-\Omega \leq \tau \leq t$ and that $E\left[\omega^{s_1}_{i,t}\varepsilon^{s_2}_{i,t}\right]$ is constant for all $t-\epsilon \leq \tau \leq t$, for $s_1, s_2 \in \{H, W\}$ and $s_1 \neq s_2$. The precise form that equation 10 takes when covariance parameters are time-varying is available from the author and was included in an earlier draft of this paper; see the subsequent paragraph for a discussion of this moment when parameters are time-varying.

Table 6: Testing for Persistence of Correlated Measurement Error: Correlation of One-Year Changes in Income with Spouse's Lagged One-Year Changes in Income

Correlation: One-Year Husbands' Changes vs. Lagged One-Year Wives' Changes

correlation	$y^H_{i,t} - y^H_{i,t-1}$
$y^W_{i,t} - y^W_{i,t-1}$	0.0235
$y^W_{i,t-1} - y^W_{i,t-2}$	0.0025
$y^W_{i,t-2} - y^W_{i,t-3}$	0.0026
$y^W_{i,t-3} - y^W_{i,t-4}$	0.0110
$y^W_{i,t-4} - y^W_{i,t-5}$	0.0002

Correlation: One-Year Wives' Changes vs. Lagged One-Year Husbands' Changes

correlation	$y^W_{i,t} - y^W_{i,t-1}$
$y^H_{i,t} - y^H_{i,t-1}$	0.0235
$y^H_{i,t-1} - y^H_{i,t-2}$	-0.0062
$y^H_{i,t-2} - y^H_{i,t-3}$	-0.0132
$y^H_{i,t-3} - y^H_{i,t-4}$	0.0099
$y^H_{i,t-4} - y^H_{i,t-5}$	-0.0033

Correlation of one-year Husbands' or Wives' (left and right panel, respectively) one-year changes in income on contemporaneous and lagged one-year changes in Wives' or Husbands' (left and right panel, respectively) incomes. Lags are one- through four-year lags. One-year changes in income are the Winsorized excess log incomes shown in Table 2.

and $t-m-1$ will be uncorrelated by assumption (discussed and supported in the next paragraph) with the transitory conditions for the other spouse in years t and $t-1$. As a result, the expectation of the product of these differences will be unaffected by the transitory covariance. But since each spouse's income difference is affected by permanent changes in income between years $t-1$ and t, the expectation of the product reveals the permanent covariance.

When the permanent covariance is constant, the long-short product estimator of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\langle^{s_1 s_2}_{\times}\rangle$, provides an unbiased estimate of the permanent covariance, $\delta^{\omega}_{i,t}$. The estimator will be unbiased even when transitory shocks for one spouse are correlated with the permanent shocks of the other, as long as this covariance is constant.[16] When the permanent covariance changes over time, this product recovers an unbiased estimate of a weighted average of recent values of $\delta^{\omega}_{i,t}$. This estimator is unbiased and unaffected by the level of or change in the transitory covariance, or by θ and ϕ. The identifying assumption here is weaker than for the corresponding estimator of the permanent variance, $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2\langle_{\times}\rangle$ (from equation 4); only the correlated component of permanent shocks must come fully into effect within Ω periods and only the correlated component of transitory income must have disappeared completely within ϵ periods. Uncorrelated shocks in permanent or transitory income can follow any process.[17] This identifying assumption is supported by Table 6, which shows that wives' one-year changes in income are roughly

[16] This is important when comparing this estimator to the long-short product estimator of permanent variance. That transitory and permanent shocks for a given spouse are uncorrelated merely reflects the definition of permanent and transitory when ϕ and θ can vary. However, it seems entirely possible that one spouse's transitory shocks might be correlated with the other's permanent shocks. For example, if one spouse becomes temporarily unemployed (with no loss in future income), this may induce the other spouse to temporarily return to work and then decide to continue working after the temporary shock has passed. This estimator of permanent covariance will be unaffected by such dynamics when their importance is time-invariant.

[17] Though in this case covariances will not map cleanly to correlations as variance estimates will be biased.

uncorrelated with their husbands' non-contemporaneous one-year changes in income, and *vice versa*. As a result, persistent correlated transitory shocks are not present so that the permanent covariance can be estimated with the long-short product, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{s_1 s_2}_{\times}\right\rangle$.

Alternatively (for $m > n > \Omega, \epsilon$), this moment estimates the permanent covariance:

$$E\left[\left(y^H_{i,t} - y^H_{i,t-m}\right)\left(y^W_{i,t} - y^W_{i,t-m}\right)\right] - E\left[\left(y^H_{i,r} - y^H_{i,r-n}\right)\left(y^W_{i,r} - y^W_{i,r-n}\right)\right] = (m-n)\,\delta^{\omega}_{i,t} \quad (12)$$ [18]

Its sample analog, the "*long-short difference* estimator of the permanent covariance,"

$$\begin{aligned}\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle &\equiv \left(\hat{\delta}^{\omega}_{i,t_2}\right)\left\langle^{m}_{raw}\right\rangle - \left(\hat{\delta}_{i,t_2}\right)\left\langle^{n}_{raw}\right\rangle \\ &= \left(y^H_{i,t_1} - y^H_{i,t_1-m}\right)\left(y^W_{i,t_1} - y^W_{i,t_1-m}\right) - \left(y^H_{i,t_2} - y^H_{i,t_2-n}\right)\left(y^W_{i,t_2} - y^W_{i,t_2-n}\right),\end{aligned} \quad (13)$$

measures the rate at which raw estimates of covariance increase with the length of the interval over which they are taken. Again, the subscript "−", a subtraction sign, refers to this difference estimator; mn refers to the pair of interval lengths that are compared. Increasing the number of years over which the squared difference is taken increases the number of correlated permanent shocks; however, the number of correlated transitory shocks remains constant at two, one at the beginning and one at the end of the interval. As with $\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2\left\langle^{mn}_{-}\right\rangle$, this estimator will be problematic when model parameters vary over time. However, the average value of $\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle$, calculated from the averages of $\left(\hat{\delta}_{i,t}\right)\left\langle^{m}_{raw}\right\rangle$ and $\left(\hat{\delta}_{i,t}\right)\left\langle^{n}_{raw}\right\rangle$ over all possible values of t (and perhaps all possible mn pairs), reveals an unbiased estimate of the average permanent covariance for a given couple i or for the population as a whole:

$$\begin{aligned}mean_i\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle\right] &\equiv mean_i\left[\left(\hat{\delta}_{i,t}\right)\left\langle^{m}_{raw}\right\rangle\right] - mean_i\left[\left(\hat{\delta}_{i,t}\right)\left\langle^{n}_{raw}\right\rangle\right] \text{ or} \\ mean\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle\right] &\equiv mean\left[\left(\hat{\delta}_{i,t}\right)\left\langle^{m}_{raw}\right\rangle\right] - mean\left[\left(\hat{\delta}_{i,t}\right)\left\langle^{n}_{raw}\right\rangle\right]\end{aligned} \quad (14)$$ [19]

Since this difference estimator, $\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2\left\langle^{mn}_{-}\right\rangle$, is not well-suited to identify within-couple variation over time in $\delta^{\omega}_{i,t}$, this paper will focus primarily on the product estimator, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{s_1 s_2}_{\times}\right\rangle$.

[18] This moment is the covariance analog of the long-short difference of squared changes used by Carroll and Samwick (and shown in equation 5) to estimate the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$.

This moment assumes that $m > n > \Omega, \epsilon$ and that $\delta^{\omega}_{i,t}$, $\delta^{\varepsilon}_{i,t}$, $E\left[\varepsilon^H_{i,t}\omega^W_{i,t}\right]$ and $E\left[\varepsilon^W_{i,t}\omega^H_{i,t}\right]$ are constant. The combination of model parameters revealed by this moment when parameters are time-varying was included in an earlier draft of this paper and is available from the author upon request.

[19] As in footnote 14 for $\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2\left\langle^{mn}_{-}\right\rangle$, this average estimator will be biased ($E\left[mean_i\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle\right]\right]$ will not generally equal $mean_i\left[\hat{\delta}^{\omega}_{i,t}\right]$ and $E\left[mean\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{mn}_{-}\right\rangle\right]\right]$ will not generally equal $mean\left[\hat{\delta}^{\omega}_{i,t}\right]$) unless the transitory covariance is an affine function of time.

Table 7: Distribution of Couples' Income Co-Movement Estimates

Type of Estimates	Long-Short Product Estimates of the Permanent Covariance		Raw Estimates of Covariance		
Estimator	$\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{\text{HW}}_{\times}\right\rangle$	$\left(\hat{\delta}^{,\omega}_{i,t}\right)\left\langle{}^{\text{WH}}_{\times}\right\rangle$	$\left(\hat{\delta}_{i,t}\right)\left\langle{}^{1}_{raw}\right\rangle$	$\left(\hat{\delta}_{i,t}\right)\left\langle{}^{3}_{raw}\right\rangle$	$\left(\hat{\delta}_{i,t}\right)\left\langle{}^{5}_{raw}\right\rangle$
Mean	-0.0065	-0.0075	-0.0024	-0.0033	-0.0101
St. Dev.	0.2993	0.3234	0.2221	0.4119	0.5105
Observations	38,120	38,120	36,678	36,678	36,678
Minimum	-4.0575	-3.9392	-5.1717	-5.1717	-5.1717
1^{st} Percentile	-1.0591	-1.1479	-0.6855	-1.3790	-1.7791
5^{th} Percentile	-0.3076	-0.3589	-0.1603	-0.4349	-0.6802
25^{th} Percentile	-0.0106	-0.0188	-0.0029	-0.0183	-0.0415
50^{th} Percentile	0	0	0	0	0
75^{th} Percentile	0.0092	0.0144	0.0040	0.0167	0.0351
95^{th} Percentile	0.2653	0.3108	0.1483	0.3786	0.5878
99^{th} Percentile	0.9436	1.0016	0.6065	1.3867	1.7799
Maximum	5.1717	4.6346	5.1717	5.1717	5.1717

The first two columns present the distribution of the long-short product estimate of the permanent covariance, $\left(\hat{\delta}^{,\omega}_{i,t}\right)\left\langle{}^{s_1 s_2}_{\times}\right\rangle \equiv \left(y^{s_1}_{i,t+n} - y^{s_1}_{i,t-m-1}\right)\left(y^{s_2}_{i,t} - y^{s_2}_{i,t-1}\right)$, for $s_1 s_2 = HW$ in column 1 and $s_1 s_2 = WF$ in column 2. These are taken as the product of the short-term change in one spouses income and the long-term change in the other's income that spans this short-term change. Columns 3, 4, and 5 present the distribution of raw estimates of the covariance, $\left(\hat{\delta}_{i,t}\right)\left\langle{}^{m}_{raw}\right\rangle \equiv \left(y^{H}_{i,t} - y^{H}_{i,t-m}\right)\left(y^{W}_{i,t} - y^{W}_{i,t-m}\right)$, for $m = 1$, 3, and 5. These show the product of contemporaneous one-, three-, and five-year changes in couples' incomes. One-half the difference between the mean of estimators in columns 4 and 5 recovers the expected average covariance parameter, $\delta^{\omega}_{i,t}$. Estimators in columns 3, 4, and 5 are only included when $\left(\hat{\delta}_{i,t}\right)\left\langle{}^{m}_{raw}\right\rangle$ is non-missing for all values of s and m for given values of i and t. This ensures that differences in sample moments do not reflect sample selection as those without moments for large values of m were more likely to have left the sample early. The moment in column 3 is positive on average when this sample selection criteria is ignored and all observations are included in the average.

3 Estimates of Co-Movement, $\hat{\delta}$

Table 7 presents the distribution of co-movement estimators. The first column presents the distribution of one set of long-short product estimates of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle \equiv \left(y^{H}_{i,t+n} - y^{H}_{i,t-m-1}\right)\left(y^{W}_{i,t} - y^{W}_{i,t-1}\right)$. The second column shows the distribution of the other such set of estimates, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{WH}_{\times}\right\rangle \equiv \left(y^{W}_{i,t+n} - y^{W}_{i,t-m-1}\right)\left(y^{H}_{i,t} - y^{H}_{i,t-1}\right)$. The third, fourth, and fifth columns present the distribution of the raw estimates of covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{m}_{raw}\right\rangle \equiv \left(y^{H}_{i,t} - y^{H}_{i,t-m}\right)\left(y^{W}_{i,t} - y^{W}_{i,t-m}\right)$, for $m = 1$, 3, and 5, respectively.

3.1 Average Co-Movement, $mean\left[\hat{\delta}\right]$

The first row of the first column of Table 7 shows the average long-short product estimate of the permanent covariance, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right] = -0.0065$. Given average estimates of the permanent variance for husbands and wives ($mean\left[\left(\hat{\sigma}^{\omega,H}_{i,t}\right)^2\left\langle_{\times}\right\rangle\right] = 0.029$ and $mean\left[\left(\hat{\sigma}^{\omega,W}_{i,t}\right)^2\left\langle_{\times}\right\rangle\right] = 0.14$ from Table 4), this implies a permanent correlation of -10 percent on average.[20] The (Winsorized excess log) permanent incomes of husbands and wives tend to move in opposite directions. The second column of Table 7 shows the distribution of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{WH}_{\times}\right\rangle$, which is similar and has $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{WH}_{\times}\right\rangle\right] = -0.0074$. This implies a permanent correlation of -12 percent on average. Equation 14 provides an alternative (difference) estimate of the average permanent covariance, $mean\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{53}_{-}\right\rangle\right] = -0.0034$, which implies a permanent correlation of -6.6 percent.[21] [22]

Negative co-movement can also be seen in Table 8. The left panel shows the regression of husbands' long-term changes in income, $y^H_{i,t+n} - y^H_{i,t-m-1}$, on wives' short-term changes in income, $y^W_{i,t} - y^W_{i,t-1}$.

$$\begin{aligned} y^H_{i,t+n} - y^H_{i,t-m-1} &= \alpha + \beta\left(y^W_{i,t} - y^W_{i,t-1}\right) + error_{i,t} \\ \hat{\beta}_{OLS} &= cov\left(y^H_{i,t+n} - y^H_{i,t-m-1}, y^W_{i,t} - y^W_{i,t-1}\right)/var\left(y^W_{i,t} - y^W_{i,t-1}\right) \\ \hat{\beta}_{OLS} &\approx mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]/mean\left[\left(\hat{\sigma}^{W}_{i,t}\right)^2\left\langle^{1}_{raw}\right\rangle\right] = 0.02 \end{aligned} \tag{15}$$

The coefficient from this univariate linear regression will equal the average permanent covariance estimate, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]$, divided by wives' average one-year raw variance estimate, $mean\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]$, which implies the coefficient of -0.02.[23] Given that the raw one-year

[20] $-0.0065/\sqrt{0.029 \bullet 0.14} \approx -0.10$. In the precence of heterogeneity in $\left(\sigma^{\omega,H}_{i,t}\right)^2$, $\left(\sigma^{\omega,H}_{i,t}\right)^2$, and $\delta^{\omega}_{i,t}$, – particularly correlated heterogeneity – this mapping of average covariance and variance into average correlation will be approximate at best, and should be viewed merely as suggestive.

[21] For symmetry, I use the analagous estimates of the average values of $mean\left[\left(\hat{\sigma}^{\omega,H}_{i,\cdot}\right)^2\left\langle^{53}_{-}\right\rangle\right] = 0.0213$ and $mean\left[\left(\hat{\sigma}^{\omega,W}_{i,\cdot}\right)^2\left\langle^{53}_{-}\right\rangle\right] = 0.1252$, respectively ($-0.066 = -0.0034/\sqrt{0.0213 \bullet 0.1252}$). As with $mean\left[\left(\hat{\sigma}^{\omega,s}_{i,\cdot}\right)^2\left\langle^{53}_{-}\right\rangle\right]$, $mean\left[\left(\hat{\delta}^{\omega}_{i,\cdot}\right)\left\langle^{53}_{-}\right\rangle\right]$ is prone to bias in the presence of concavity in the life-cycle pattern in $\delta_{i,t}$.

[22] Note that the finding of negative co-movement is different from results from other papers (e.g., (Hess, 2004)). This difference in sign reflects in part use of excess income and not raw income in this paper. In particular, couples of similar ages and educational backgrounds tend to have predicted income paths that move together. For example, young people and well-educated people tend to marry one another. Since these groups have relatively steep labor income profiles, the correlation of incomes will be positive even if the non-deterministic component of incomes are negatively correlated. The difference in sign also reflects the use of different moments to capture co-movement. This paper uses moments that identify the permanent covariance without contamination from posively correlated measurement error or other correlated transitory shocks.

[23] Note that the average product of two random variables is equivalent to the sample covariance of these variables when they are both mean-zero. By construction, changes in excess log income are mean-zero, so that $cov\left(y^H_{i,t+n} - y^H_{i,t-m-1}, y^W_{i,t} - y^W_{i,t-1}\right)$ is $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]$ and $var\left(y^W_{i,t} - y^W_{i,t-1}\right)$ is $mean\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]$.

Table 8: Co-Movement OLS Regression: Long-Term Changes in One Spouse's Income vs. the Short-Term Changes in the Other's Income They Span

Dependent Variable: $y^H_{i,t+2} - y^H_{i,t-3}$, Long-Term Change in Husband's Income

$y^W_{i,t} - y^W_{i,t-1}$	-0.018 (4.36)**	-0.019 (4.43)**
$\left(y^W_{i,t} - y^W_{i,t-1}\right)^2$		-0.004 (1.45)
Constant	-0.043 (17.62)**	-0.042 (16.06)**
Observations	38,120	38,120
R^2	0.0005	0.0006

Dependent Variable: $y^W_{i,t+2} - y^W_{i,t-3}$, Long-Term Change in Wife's Income

$y^H_{i,t} - y^H_{i,t-1}$	-0.069 (4.58)**	-0.066 (4.36)**
$\left(y^H_{i,t} - y^H_{i,t-1}\right)^2$		0.024 (1.47)
Constant	0.041 (8.49)**	0.039 (7.54)**
Observations	38,120	38,120
R^2	0.0005	0.0006

This table presents results from regressions that predict the five-year changes in Winsorized excess log income of one spouse with the one-year changes in their spouse's Winsorized excess log income in the middle of each five-year interval. The left panel predicts five-year changes for wives; the right panel predicts five-year changes for husbands. The first column in each panel presents a linear functional form; the second column presents a quadratic functional form. Changes in income are as shown in Table 2. Absolute value of t-statistics are in parentheses. "*" indicates significance at the 5% level; "**" indicates significance at the 1% level.

estimates of variance for the wife must be positive, this OLS regression provides another way to see the negative relationship between permanent innovations to couples' incomes. This result implies that when a wife's income doubles in a given year her husband's income falls by 2 percent more on average than it would otherwise in the five-year period surrounding this change. The second set of results in this column show OLS regressions of $y^H_{i,t+n} - y^H_{i,t-m-1}$ on a quadratic in $y^W_{i,t} - y^W_{i,t-1}$. The right panel shows the regression of wives' long-term changes in income, $y^W_{i,t+n} - y^W_{i,t-m-1}$ on husbands' short-term changes in income, $y^H_{i,t} - y^H_{i,t-1}$. These reveal a parallel relationship between couples' changes in income. The hypothesis that the higher-order coefficients are equal to zero cannot be rejected. The relationship between short-term changes in wives' incomes and spanning long-term changes in husbands' average incomes seems to be relatively linear.

It is worth noting that the relationship between husbands' (Winsorized, excess) log incomes and wives' (Winsorized, excess) log incomes is also present when looking at husbands' log incomes and a variety of work-related variables for wives. Table 9 presents the distribution of $\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)\left(x^W_{i,t} - x^W_{i,t-1}\right)$ for a several work-related variables for wives, x^W. The previ-

In practice, average changes in income will not be exactly mean-zero due to sample selection and Winsorizing effects. These effects are second-order. De-meaned results provided in regressions replicate almost exactly results based on the (not de-meaned) long-short product estimator. From Table 2, $mean\left[y^W_{i,t} - y^W_{i,t-1}\right] = 0.0004$ and $mean\left[y^H_{i,t+n} - y^H_{i,t-m-1}\right] = -0.0431$. The product of these two is trivially small, so de-meaning changes in income is quantitatively unimportant. As a result, regression results (wich de-mean) align almost perfectly with results using $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]$ (which do not).

Table 9: Distribution of Estimates of the Permanent Covariance of Husbands' Excess Log Incomes with Wives' Excess Hours, Excess Log Incomes, and Labor Force Participation

Long-Short Product Estimates of the Permanent Covariance: $\left(y^H_{i,t+2} - y^H_{i,t-3}\right)\left(x^W_{i,t} - x^W_{i,t-1}\right)$

Husband's Variable: y	Log Labor Income			
Wife's Variable: x	Log Labor Income	Hours Worked	Log Labor Income if in Labor Force	In Labor Force? one or zero
Implied Correl.	-0.1008	-0.1269	-0.1749	-0.0980
Mean	-0.0065	-5.7	-0.0076	-0.0029
St. Dev.	0.2993	242.3	0.2809	0.1734
Observations	38,120	38,120	23,139	38,120
Minimum	-4.0575	-3309.9	-3.6556	-1.8514
1$^{\text{st}}$ Percentile	-1.0591	-864.6	-0.9867	-0.6758
5$^{\text{th}}$ Percentile	-0.3076	-283.9	-0.3223	-0.0756
25$^{\text{th}}$ Percentile	-0.0106	-13.1	-0.0310	0.0000
50$^{\text{th}}$ Percentile	0.0000	0.0	0.0000	0.0000
75$^{\text{th}}$ Percentile	0.0092	10.8	0.0306	0.0000
95$^{\text{th}}$ Percentile	0.2653	244.6	0.2886	0.0355
99$^{\text{th}}$ Percentile	0.9436	766.6	0.8415	0.6195
Maximum	5.1717	3227.3	3.6701	1.8514

Each column presents the distribution of a long-short product estimator of the permanent covariance, $\left(y^H_{i,t+2} - y^H_{i,t-3}\right)\left(x^W_{i,t} - x^W_{i,t-1}\right)$, for a different variable, x. In each case, y refers to the Winsorized excess log income of the husband presented in Table 2. See text for construction of each x variable for the wife. The first row presents the implied correlation. This is calulated as the ratio of the mean long-short product estimator (shown in the second row of this table) to the square root of the product of the mean long-short product estimators of the permanent variance for husbands and wives.

ous results examined the relationship between changes in the excess log incomes of husbands and the excess log incomes of wives. Here, we look also at changes in excess (levels, not logs) hours worked by wives, changes in excess log incomes for wives who remain working (where Winsorizing does not bind, dropping observations that had previously be Winsorized), and changes in labor force participation for wives.[24] Note that the covariances are negative for all work-related variables for wives, with relatively similar correlations for each work-related variable for wives.[25] Examining long-term changes in these work-related variables for husbands would not be fruitful, since the vast majority of husbands remain fully employed in the long-run and therefore exhibit no useful variation not captured by variation in income. When wives' hours increase or when wives enter the labor

[24] Excess hours are calculated just as excess log income but in levels and not logs, with Winsorizing at the 5th and 95th percent levels. Excess log income for wives who work are just as excess log income, but with any observations below the 5th percentile or above the 95th percentile dropped. Changes in labor force participation are -1 if wives leave the labor force, 0 if they remain in or out of the labor force during the period, and 1 if they enter the labor force. A wife is considered in the labor force if her income exceeds the 5th percentile level, so that it provides a complement to the previous variable. Unfortunately, hours data is too noisy to examine wives' wages, which are measured as the ratio of income to hours worked. This is problematic when hours worked are zero.

[25] But since these variables will not evolve as in equation 2, estimators do not map to model parameters.

Table 10: Quantile Regression of Long-Term Changes in Husband's Excess Log Income on Short-Term Changes in Wife's Excess Log Income

Linear Quantile Regression
Dependent Variable: Long-Term Change in Husband's Income, $y^H_{i,t+2} - y^H_{i,t-3}$

Percentile	5$^{\text{th}}$ Percentile	25$^{\text{th}}$ Percentile	75$^{\text{th}}$ Percentile	95$^{\text{th}}$ Percentile
$y^W_{i,t} - y^W_{i,t-1}$	-0.029 (1.40)	-0.014 (1.94)	-0.022 (6.08)**	-0.027 (2.02)*
Constant	-0.938 (89.49)**	-0.233 (77.92)**	0.184 (55.08)**	0.696 (64.67)**

Quadratic Quantile Regression
Dependent Variable: Long-Term Change in Husband's Income, $y^H_{i,t+2} - y^H_{i,t-3}$

Percentile	5$^{\text{th}}$ Percentile	25$^{\text{th}}$ Percentile	75$^{\text{th}}$ Percentile	95$^{\text{th}}$ Percentile
$y^W_{i,t} - y^W_{i,t-1}$	-0.026 (3.21)**	-0.020 (4.23)**	-0.014 (2.06)*	-0.035 (2.15)*
$\left(y^W_{i,t} - y^W_{i,t-1}\right)^2$	-0.029 (2.83)**	-0.017 (5.32)**	0.008 (2.20)*	0.041 (4.39)**
Constant	0.181 (79.55)**	0.680 (57.14)**	-0.227 (63.21)**	-0.928 (55.15)**

Results from quantile regressions. The dependent variable is $y^H_{i,t+2} - y^H_{i,t-3}$. The explanatory variables are a linear (top panel) or quadratic (bottom panel) functions of $y^W_{i,t} - y^W_{i,t-1}$. The probability that all coefficients from the top panel are the same is 0.505. The probability that all quadratic coefficients from the bottom panel are the same is 0.000, as is the probability that all linear coefficients from the bottom panel are equal to zero. Absolute value of t-statistics are in parentheses. "*" indicates significance at the 5% level; "**" indicates significance at the 1% level. Results obtained from parallel regressions in which husbands and wives are switched obtain results which are identical in sign and statistical significance.

force, it indicates a long-term drop in their husbands' incomes on average. This suggests that co-movement reflects not just diversification – correlated exogenous shocks – but also coordination – joint endogenous decisions.

What type of changes in income – large or small – are moving together? This can be seen in the top panel Table 10, which shows the linear relationship between a short-term change in a wife's income, $y^W_{i,t} - y^W_{i,t-1}$, and the 5$^{\text{th}}$, 25$^{\text{th}}$, 75$^{\text{th}}$, and 95$^{\text{th}}$ percentiles of long-term changes her husband's income that span them, $y^H_{i,t+2} - y^H_{i,t-3}$.[26] The relationship between $y^W_{i,t} - y^W_{i,t-1}$ and $y^H_{i,t+2} - y^H_{i,t-3}$ is negative at all quantiles and similar (statistically indistinguishable) across quantiles. The hypothesis that coefficients for all quantiles are jointly equal to zero is strongly rejected. When a wife increases her income in a given year, it suggests a shift downward in the whole distribution of her husband's long-term changes in income. Her husband's income is more likely to have been falling

[26] These quantile regressions mirror the "mean" (OLS) regressions in Table 8 of the same functional form.

Figure 1: Quantile (Linear) Regression: Distribution of One Spouse's Long-Term Changes in Income vs. Short-Term Changes in the Other's Income

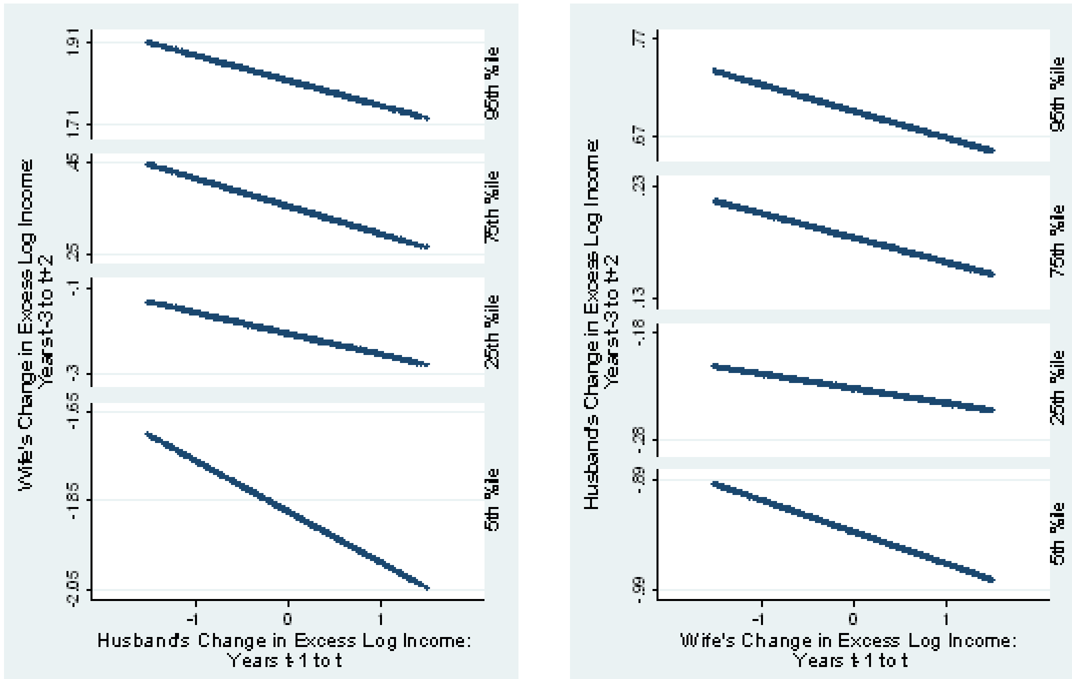

These figures show the relationship between short-term changes in wives' incomes (left panel, husbands' incomes in right panel) and the distribution of long-term changes in husbands' incomes (left panel, wives' incomes in right panel) that span these short-term changes. The left panel shows graphically the results from the top panel of Table 10; the right panel depicts the same regression with the husband and wife reversed. These present the impact of one-year changes in one spouse's income on the the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles in the five-year changes in the other spouse's income that span these one-year changes. In each panel, the hypothesis that all slopes are the same cannot be rejected.

slightly, more likely to have been falling dramatically, less likely to have been increasing slightly, and less likely to have been falling dramatically. This same effect can be seen in the left panel of Figure 1, which presents these results graphically; it shows the negative linear relationship predicted between $y^W_{i,t} - y^W_{i,t-1}$ and the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles of $y^H_{i,t+2} - y^H_{i,t-3}$. The scale for each quantile is the same, so that the slopes are comparable within the panel. The right panel flips husbands and wives, looking at the relationship between $y^H_{i,t} - y^H_{i,t-1}$ and the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles of $y^W_{i,t+2} - y^W_{i,t-3}$. Results are similar, with the negative relationship apparent and similar at all quantiles. The covariance of couples' incomes reflects negative co-movement throughout the distribution of possible outcomes. It is not limited to wives responding to very bad outcomes as in the added worker effect literature (e.g., the effects of husbands' unemployment (Lundberg, 1985) or the long-term effect of husbands' job displacement (Stephens, 2002)).

Figure 2: Joint Distribution of the Husband's Change in Log Excess Income vs. the Wife's Change in Labor Force Participation

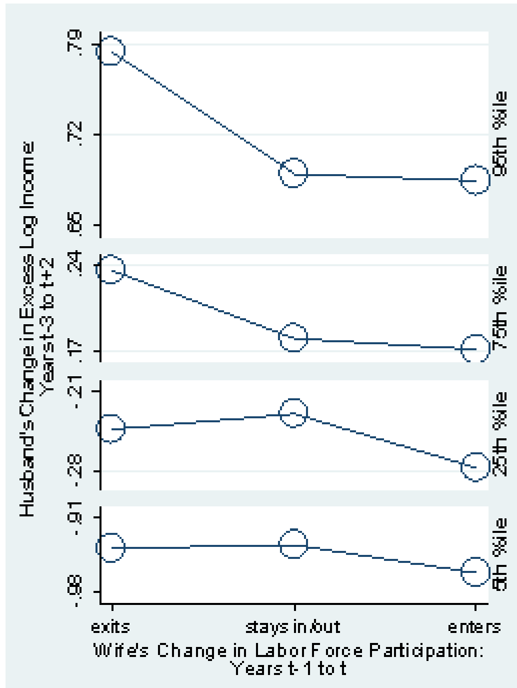

This figure presents the the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles of $y^H_{i,t+2} - y^H_{i,t-3}$. These are partioned by short-term changes in wives' labor force participation, between years $t-1$ and t. As described in the text, labor force participation is defined as "in " or "out " according to whether income is below or above the 5^{th} percentile Winsorizing bound. All observations are then categorize according to whether the wife "enters", "exits", or "stays in/out" of (no change) the labor force.

The negative relationship between short-term changes for wives and the distribution of changes in their husbands' incomes is also present when looking at changes in wives' labor force participation instead of changes in wives' incomes. Figure 2 presents the distribution of long-term changes in husbands' incomes, $y^H_{i,t+n} - y^H_{i,t-m-1}$, given short-term changes in wives' labor force participation. In each year, I define a wife as being in or out of the labor force if her income is below the 5^{th} percentile Winsorizing bound (which would include incomes of zero or real annual incomes below $\$1,000$). Over any one year period, a wife either "exits" the labor market, "stays in/out" of the labor force if there is no change, or "enters" the labor force. The distribution of $y^H_{i,t+n} - y^H_{i,t-m-1}$ is shown for each of these three possible states for wives for the 5^{th}, 25^{th}, 75^{th}, and 95^{th} percentiles of $y^H_{i,t+n} - y^H_{i,t-m-1}$. When a wife leaves the labor force, it suggests the whole distribution of her husband's changes in long-term income is higher. When a wife enters the labor force in a given year, it suggests the whole distribution of her husband's changes in long-term income is marginally lower.

3.2 "Wife-Swap Bootstrap" Test of Couples' Independence

The negative mean of the long-short product estimator of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$ (presented in column 1 of Table 7) – or equivalently, the negative relationship between $y^W_{i,t} - y^W_{i,t-1}$ and $y^H_{i,t+2} - y^H_{i,t-3}$ (shown in Table 8 or the top panel of Table 10) – indicates that innovations to couples' incomes are negatively correlated *on average*. However, there is dramatic variation in the long-short product estimates of the permanent covariance; $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right] = -0.0065$ is much smaller in absolute value than $s.d.\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right] = 0.299$. How should we interpret this variation? Is it large enough to make the average effect statistically insignificant? Can the large sample variance be explained solely by measurement error? In general, these questions are difficult to answer when estimation error is correlated over t, as it will be when we compute a panel of values for $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$. If the income processes for husbands and wives were understood and modelled perfectly, the correlated measurement error in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$ could be predicted analytically and these hypotheses could be tested. Doing this is not possible for two reasons. First, it requires estimates not only of the permanent variance, $\left(\sigma^{\omega,s}_{i,t}\right)^2$ – which can be estimated cleanly using the moment specified in equation 3 – but also a host of other parameters, θ, ϕ and $\left(\sigma^{\varepsilon,s}_{i,t}\right)^2$, which have no moment to which they map cleanly. In fact, the reason to use the long-short product, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$, to estimate the permanent covariance, $\delta^{\omega}_{i,t}$, is that it is unbiased regardless of the values for these other parameters. If we must obtain estimates of these parameters to make inferences about the permanent covariance, $\delta^{\omega}_{i,t}$, then there is no reason to use this estimator. Just as important, the process for income may not evolve exactly as in equation 2; there is higher-order autocorrelation that is incompatible with equation 2. As a result, any test is a joint test of the hypothesis of interest and the hypothesis that the individual income processes are specified properly.

Here, I propose a randomization test (Fisher (1935); Romano (1989); Kennedy (1995)) of the hypothesis that husbands' incomes evolve independently from their wives' incomes. This technique makes no assumptions about the labor income process of the husband or the wife, but instead exploits (by block bootstrapping) the observed income processes for husbands and wives in the data. First, I randomly assign a wife from the sample to each husband in the sample (or equivalently husbands to each wife). Pairs are random within the set of couples with the same number of years of data. For each random husband-wife pair, I place this pair's two income series side-by-side and calculate all possible long-short product estimates of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$, for

Table 11: "Wife-Swap Bootstrap" Test of the Independence of Couples' Income Processes: Comparing the Distribution of Permanent Covariance Estimates for Actual and Wife-Swapped Data

Distribution Feature	Actual	Wife-Swap Bootstrap Mean	Med.	St. Dev.	Min.	Max.
Mean	-0.0065	-0.0006	-0.0006	0.0012	-0.0045	0.0032
St. Dev.	0.2993	0.2764	0.2765	0.0045	0.2627	0.2886
St. Dev. of Lifetime Average	0.1197	0.1053	0.1052	0.0037	0.0920	0.1177
Mean Time-Trend	0.0017	-0.0001	0.0000	0.0006	-0.0023	0.0016
St. Dev. of Time-Trend	0.0566	0.0567	0.0562	0.0052	0.0433	0.0786
Autocorrelation Coef.	-0.1794	-0.1815	-0.1816	0.0139	-0.2274	-0.1375
Observations	38,120	36,260	36,259	31	36,174	36,376

The first column presents the distribution of long-short estimates of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle \equiv \left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)\left(y^W_{i,t} - y^W_{i,t-1}\right)$. The rows present various moments from this distribution. The next five columns present results from the bootstrap procedure described in the paper. In particular, all observations for each husband and wife with the same number of years of data are randomly paired, and $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$ is computed for each random pair in each year. This forms one bootstrap draw, and each feature of the distribution of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$ found in the rows is computed for each bootstrap draw. This procedure is repeated 1,000 times. The next five columns present the mean, median, standard deviation, minimum and maximum values over the 1,000 draws for the features of the distribution given in the rows. The first and second rows present the mean and standard deviation of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$. The third row presents the observation-weighted standard deviation of couple-specific lifetime averages of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$. The fourth row presents the average linear time trend for $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$. The fifth row is formed by computing a couple-specific linear time trend of permanent covariance estimates, and presents the standard deviation of these estimates. The sixth row presents the coefficient from a regression of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$ on $\left(\hat{\delta}^{\omega}_{i,t-1}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$. The final row presents the number of observations on $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{\mathrm{HW}}_{\times}\right\rangle$ used for these calculations. It is lower for the bootstrapped distributions because couples tend to have missing data in the same years, so randomly pairing couples increases the number of years over which at least one has missing data.

this pair. Doing this for all pairs provides a panel of synthetic $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$ values from which I obtain summary statistics. I repeat this process $1,000$ times, and present the distribution of these summary statistics. This provides a reference distribution, the distribution of summary statistics that would be expected under the null hypothesis that couples' incomes moved independently but individuals' incomes evolved just as they do in the data. This technique could be given the slightly risque name "wife-swap bootstrap," as the reference distribution is obtained by bootstrapping full income series from the actual data for randomly swapped husband-wife pairs.[27]

Table 11 compares summary statistics from the observed data with the distribution of these summary statistics from the wife-swap bootstrap. The first row shows the mean of the permanent

[27] The closest application of a randomization test to couples is (Jepsen and Jepsen, 2002), who study assortative mating by comparing "actual couples with randomly created (artificial) couples to see if actual couples are more similar or more different than the random pairings."

covariance estimates, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]$; the first column shows features of the true observed data; as a result, the first row of the first column of Table 11, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]_{true} = -0.0065$, merely repeats the results from the first row of the first column of Table 7. The subsequent columns show features of the reference distribution. For example, the first row shows the distribution of mean permanent covariance estimates, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]_{wsbs}$, obtained over the $1,000$ "wife-swap bootstrap" draws. The reference distribution has a mean of -0.0006, a standard deviation of 0.0012 and in $1,000$ draws varies from -0.0045 to 0.0032. In other words, the average estimate of the permanent covariance in the observed data is significantly more negative than the null provided by the reference distribution, rejecting the hypothesis that $mean\left[\delta^{\omega}_{i,t}\right] = 0$.

The second row presents the standard deviation of co-movement estimates, $s.d.\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]$, both the value from observed data (column 1) and distribution of values from the "wife-swap bootstrap" reference distribution (subsequent columns). The standard deviation in the observed data, $s.d.\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]_{true} = 0.299$, is significantly greater than values from the reference distribution. $s.d.\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right]_{wsbs}$ has a mean of 0.276, a standard deviation of 0.005, a minimum of 0.263, and a maximum of 0.289 in $1,000$ draws.[28] There is "too much" variation in sample estimates to be explained solely by measurement error; there is less variation in estimates from randomly paired couples, whose incomes move independently by construction. There must be some underlying heterogeneity in model parameters driving excess variation in parameter estimates. In the next section, I present this "excess variance" graphically and decompose it parametrically.

4 Excess Variation in $\hat{\delta}$

Variation in the long-short product estimator of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$, can be decomposed using the definition of variance:

$$var\left(\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right) \equiv E\left[\left(y^{W}_{i,t} - y^{W}_{i,t-1}\right)^{2}\left(y^{H}_{i,t+n} - y^{H}_{i,t-m-1}\right)^{2}\right] - E\left[\delta\right]^{2}. \tag{16}$$

As shown in the last section, there is too much variation in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$ to be explained by measurement error alone; there must be some heterogeneity between couples or within a couple over time.[29]

[28] The same excess variation is present when looking at the comparable moment based on the co-movement of husbands' log incomes with wives' hours.

[29] The "wife-swap bootstrap" test showed that $var\left(\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right)$ was larger than it would be under the null that couples' incomes moved independently. If one is willing to forgo confidence intervals, this can be shown without the "wife-swap bootstrap" test. If couples' incomes moved independently, then

$$var\left(\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right) = E\left[\left(\hat{\sigma}^{H}_{i,t}\right)^{2}\left\langle^{m+n+1}_{raw}\right\rangle\right] E\left[\left(\hat{\sigma}^{W}_{i,t}\right)^{2}\left\langle^{1}_{raw}\right\rangle\right].$$

Here, I provide an alternative way to see this heterogeneity in the data. Note from the definition of covariance that:

$$cov\left(\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2,\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\right) \equiv E\left[\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2\right] - E\left[\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2\right]E\left[\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\right]. \tag{17}$$

Both the variation in co-movement (variance of long-short product estimates of the permanent covariance, $var\left(\left(\hat{\delta}^\omega_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right)$) and correlated heterogeneity in volatility (the covariance of couples' raw variance estimates, $cov\left(\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2,\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\right)$) are identified from the same moment, $E\left[\left(y^W_t-y^W_{t-1}\right)^2\left(y^H_{t+n}-y^H_{t-m-1}\right)^2\right]$. Combining equations 16 and 17 and given that $E\left[\delta^\omega_{i,t}\right]^2$ is negligibly small,

$$var\left(\left(\hat{\delta}^\omega_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right) > E\left[\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2\right]E\left[\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\right] \text{ implies} \tag{18}$$

$$cov\left(\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2,\left(y^W_{i,t}-y^W_{i,t-1}\right)^2\right) > 0. \tag{19}$$

If there is too much variation in co-movement estimates, $\left(\hat{\delta}^\omega_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle$, large (absolute; formally, squared) changes in income for husbands and wives will tend to coincide.

The tendency of large absolute changes in couples' incomes to coincide is apparent in the bottom panel of Table 10. This panel presents quantile regressions of long-term changes in husbands' incomes, $y^H_{i,t+2}-y^H_{i,t-3}$, on a second-order polynomial (quadratic) in the short-term changes in their wives' incomes, $y^W_{i,t}-y^W_{i,t-1}$. The noteworthy features of this panel are the negative coefficients on quadratic terms for quantiles below the median (5^{th} and 25^{th}) alongside positive coefficients on these terms for quantiles above the median (75^{th} and 95^{th} percentiles). This statistically significant finding is apparent from the bowing away from the mean in Figure 3, which presents predicted values from these regressions graphically. The gap between 5^{th} and 95^{th} or 25^{th} and 75^{th} percentiles is greatest when the change in the wife's income is large in absolute value. When a wife's income has a substantial one-year change, her husband is much more likely to have a substantial long-term change in income (either positive or negative) surrounding this period.

However,

$$var\left(\left(\hat{\delta}^\omega_{i,t}\right)\left\langle^{HW}_{\times}\right\rangle\right) = 0.2993^2 = 0.0896 > mean\left[\left(\hat{\sigma}^H_{i,t}\right)^2\left\langle^{5}_{raw}\right\rangle\right] mean\left[\left(\hat{\sigma}^W_{i,t}\right)^2\left\langle^{1}_{raw}\right\rangle\right] = 0.3109 \bullet 0.2308 = 0.0718.$$

The variance of co-movement estimates is too large to be consistent with the hypothesis that couples' incomes move independently.

Figure 3: Quantile (Quadratic) Regression: Long-Term Changes in Husband's Income vs. Short-Term Changes in Wife's Income

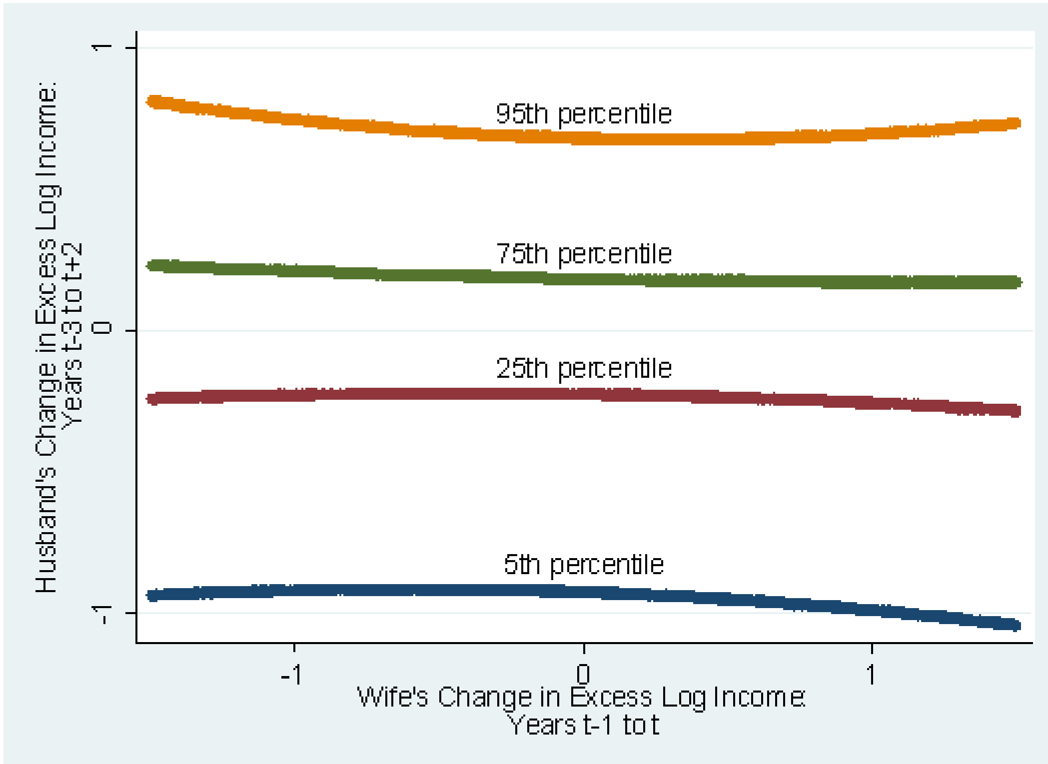

Figure plots the predicted values from the quadratic quantile regression in the bottom panel of Table 10. See text and Table 10 for details. The hypotheses that all quadratic terms are equal to zero or equal to each other are both strongly rejected.

The same pattern is apparent when switching husbands and wives in the estimation. In quantile regression results that are not shown, there is excess variation in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{WH}_{\times}\right\rangle$ (or bowing away from the mean). Figure 2 shows that the same patterns of correlated absolute changes is also apparent when looking at short-term changes in wives' labor force participation and the distribution of long-term changes in husbands' log incomes. The distribution of long-term changes in husbands' incomes is more diffuse when these changes span intervals in which their wives changed labor force participation. Large (absolute) changes in husbands' incomes tend to coincide with both entry into and exit from the labor force for their wives.

4.1 Decomposing Excess Variation in $\hat{\delta}$

What should we make of the finding that there is "too much" variation in the long-short product estimates of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$, or equivalently that large absolute changes in husbands' and wives' incomes tend coincide? If the labor income process for both spouses follows equation 1, non-contemporaneous shocks are independent, shocks are conditionally normal, then

$$cov\begin{pmatrix} \left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2, \\ \left(y^W_{i,t} - y^W_{i,t-1}\right)^2 \end{pmatrix} = cov\left(E\left[\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2\right], E\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]\right) \quad (20)$$

$$+2\left(var\left(\rho_{i,t}\right) + E\left[\rho_{i,t}\right]^2\right) E\left[\left(\sigma^{\omega,H}_{i,t}\right)^2 \left(\sigma^{\omega,W}_{i,t}\right)^2\right].^{30}$$

I use $\rho_{i,t} \equiv \delta^{\omega}_{i,t} / \sqrt{\left(\sigma^{\omega,W}_{i,t}\right)^2 \left(\sigma^{\omega,H}_{i,t}\right)^2}$ to indicate the "permanent correlation." The "excess variance" in the long-short product estimates can be decomposed into two types of regime switching, as shown in equation 20. First, the *expected magnitude* of couples' shocks may be correlated:

$$cov\left(E\left[\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2\right], E\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]\right) > 0. \quad (21)$$

In this case, observations are either in a high-volatility regime for both spouses or low-volatility regime for both spouses. If the expected magnitude of shocks is correlated, then in high-volatility regimes large changes of either sign in both spouses' incomes will be likely. In low-volatility regimes, small changes of either sign in both spouses' incomes will be likely. As a result, there will be an excess of observations where either both or neither spouse experiences a large change in income, and a dearth of observations in which exactly one spouse experiences a large change in income.

Second, there may be heterogeneity in the correlation between spouse's income changes:

$$var\left(\rho_{i,t}\right) > 0.^{31} \quad (22)$$

Couples' innovations switch between regimes (either over time or between couples) in which they are positively correlated and regimes in which they are negatively correlated. In regimes in which innovations are positively correlated, both spouses will have a large positive change in income, both spouses will have little change in income, or both spouses will experience a large negative change in income. In regimes in which innovations are negatively correlated, one spouse will experience a large positive change in income while the other experiences a large negative change or both will experience very little change. Averaging over these two regimes, there will be an excess of observations where either both or neither spouse experiences a large change in income, and a dearth of observations in which exactly one spouse experiences a large change in income. This gives a pattern in the data

[30] Assuming a different process for income will lead to a qualitatively similar decomposition.

[31] $E\left[\rho_{i,t}\right]^2 \approx (-0.1)^2 = 0.01$ is too small to explain more than a tiny fraction of the covariance in equation 20.

Figure 4: Stylized Joint Distribution of Couples' Changes in Income

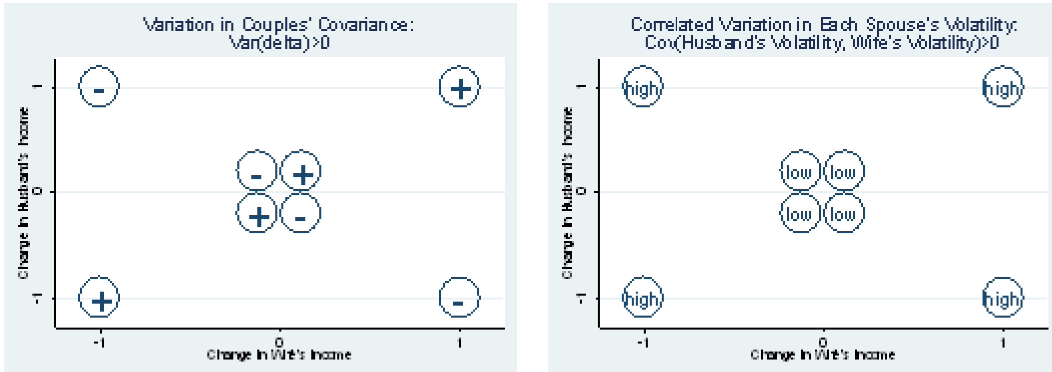

Stylized depiction of the data. See text for details.

that is qualitatively the same as the previous explanation.

This decomposition of "excess variance" into *variance in the covariance parameters* (equation 22) and the *covariance of couples' variance parameters* (equation 21) can be seen clearly in Figure 4. The two panels of Figure 4 merely presents a stylized depiction of the joint distribution of realized changes in each spouses income, showing the results visible in Figure 3 and the lower panel of Table 10 in a stylized way. The main feature of this joint distribution is that observations are either found at the corners, in which case both spouses had a large (absolute) change in income, or in the middle, in which case neither spouse had a large change in income. There are few observations (and in this stylized graph, none) in which exactly one spouse had a large (absolute) change in income.

The left panel of Figure 4 shows one of the two possible explanations for this pattern: variation in the covariance parameter, δ (from equation 22). In this panel, some observations are drawn from a regime or distribution in which couples income changes are positively correlated (here, perfectly so). These observations are labeled with a "+" for their correlation and are found in a line from lower-left to upper-right. Other observations in this panel are drawn from a regime or distribution in which couples income changes are negatively correlated (here, perfectly so). These observations are labeled with a "-" and are found in a line from the upper-left to lower-right. This variation in δ could reflect differences across couples (e.g., some couples always have more highly correlated incomes than others) or within couples over time (e.g., couples' incomes become more highly correlated over time). Note that income volatility for each spouse is constant in this example.

The right panel of Figure 4 shows the other possible explanation for this pattern: the true income volatility parameters for husbands and wives vary and are positively correlated (equation 21). In this panel, some observations are drawn from a regime or distribution with high volatility for both spouses. These observations are labeled with a "high" and are found at the corners of the panel, as each spouse will have a large absolute change in income under this regime. Other observations are drawn from a regime or distribution with low volatility for both spouses. These observations are labeled with a "low" and are found in the middle of the panel, as each spouse will have only a negligible change in income under this regime. This correlated variation in income volatility could reflect innate between-couple differences in income volatility (e.g., assortative mating on income risk) or within-couple variation over time (e.g., clustering of risks in time). Note that the covariance of couples incomes is constant in this example.

While either explanation or both could in theory explain the general shape of the joint distribution of couples' incomes, there is too little variation in $\rho_{i,t}$ (equation 22) to explain the magnitude of the effect. Even if ρ (improbably) alternated between -1 and 1, so that $var\left(\rho_{i,t}\right) + E\left[\rho_{i,t}\right]^2 = 1$, this could explain less than half of the excess variation if the expected magnitudes of couples income shocks are uncorrelated.[32] Section 4.2 shows that a plausible degree of assortative mating in risk – correlated between-couple variation in the magnitude of shocks, the tendency of individuals with volatile incomes to be married to one another – could explain the covariance of couples' volatility measures (or equivalently, the excess variance in $\hat{\delta}$).

Without being able to identify changes in regimes (e.g., high or low correlation, high volatility for both spouses or low volatility for both spouses) ex-ante, it is impossible to determine whether results are being driven by heterogeneity in δ^{ω} (variance in the covariance parameter) or correlated heterogeneity in $(\sigma^s)^2$ (the covariance of couples' variance parameters). However, it is possible to identify explanatory variables that affect ρ, so that heterogeneity in these variables explains heterogeneity in the degree of co-movement. It is also possible to identify explanatory variables that move both $E\left[\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2\right]$ and $E\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]$ in the same direction, so that heterogeneity in these explanatory variables explains the co-movement of couples' expected volatilities. Put another way, it is possible to differentiate the two types of regime-switching if, but only if, we can identify variables that predict which regime we are in.

[32] Here, I combine equations 16 and 17. Use sample moments (from column 1 in Table 7, $var\left(\hat{\delta}^{HW,\omega}_{i,t}\right) = 0.2993^2 = 0.0896$ and $mean\left[\hat{\delta}^{HW,\omega}_{i,t}\right] = -0.0065$; from the first row of columns 4 and 6 of Table 7 as $\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2 = 0.2308$ and $\left(y^W_{i,t} - y^W_{i,t-1}\right)^2 = 0.3109$) gives $cov\left(\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2, \left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right) \approx 0.018$.

Equation 20, $cov\left(E\left[\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2\right], E\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]\right) = 0$, $var\left(\rho_{i,t}\right) + E\left[\rho_{i,t}\right]^2 = 1$, $\widehat{mean}\left(\sigma^{\omega,H}_{i,t}\right) = 0.029$, and $mean\left(\sigma^{\omega,W}_{i,t}\right) = 0.14$, imply $cov\left(\left(y^H_{i,t+n} - y^H_{i,t-m-1}\right)^2, \left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right) = 0.0084$.

Table 12: Assortative Mating on Income Volatility: The Covariance of Couples' Lifetime Average Volatility Estimates

Wife \ Husband	Average One-Year Raw Variance	Average Five-Year Raw Variance	Average Long-Short Product Variance
Average One-Year Raw Variance	0.0074 (13.27%)	0.0128 (11.75%)	0.0013 (3.62%)
Average Five-Year Raw Variance	0.0102 (8.30%)	0.0364 (15.04%)	0.0025 (3.04%)
Average Long-Short Product Variance	0.0015 (3.69%)	0.0035 (4.46%)	0.0005 (2.02%)

For each individual, three moments are calculated. The first, $mean_i\left[\left(y^s_{i,t} - y^s_{i,t-1}\right)^2\right]$, is the sample average of squared one-year change in excess log income. The second, $mean_i\left[\left(y^s_{i,t} - y^s_{i,t-5}\right)^2\right]$, is the sample average of squared five-year change in excess log income. The third, $mean_i\left[\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2\langle\times\rangle\right]$, is the sample average of the product of one-year changes in excess log income and the five-year changes that surround them, as described in equation 4. The table presents the covariance across couples of these individual-specific measures, the degree of assortative mating in various income estimators. These are shown in the columns for husbands and rows for wives. Numbers in parentheses denote the sample correlations. Covariances are weighted by the number of observations used to form the individual-specific averages for each couple.

4.2 Assortative Mating: Between-Couple Variation in σ^2

Couple-specific indicator variables are obvious candidates to predict either heterogeneity in co-movement, δ^ω, or correlated heterogeneity in volatility, $(\sigma^{\omega,s})^2$. In Subsection 4.3, I test the hypothesis that couples differ in their co-movement parameters, δ^ω (or ρ). Individuals will almost certainly differ in their lifetime-average income volatility; in this subsection, I show that the lifetime-average volatilities of married couples are correlated. Put another way, there is assortative mating in income volatility, with individuals with volatile incomes married to each other.[33]

To investigate assortative mating in income volatility, I compute three lifetime-average individual-specific estimates of income volatility for both husbands and wives: $mean_i\left[\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2\langle\times\rangle\right]$, $mean_i\left[\left(y^s_{i,t+n} - y^s_{i,t-m-1}\right)^2\right]$ and $mean_i\left[\left(y^s_{i,t} - y^s_{i,t-1}\right)^2\right]$. Next, I compute the between-couple covariance (with correlation in parentheses) of average variance estimates for husbands and average variance estimates for wives.

The between-couple covariance of average five-year variances for husbands and average one-year variances for wives, $cov\left(mean_i\left[\left(y^H_{i,t+2} - y^H_{i,t-3}\right)^2\right], mean_i\left[\left(y^W_{i,t} - y^W_{i,t-1}\right)^2\right]\right)$, is presented in the first row of the second column of Table 12. This is the relevant covariance when trying to explain

[33] While I use the phrase "assortative mating in risk" to refer to this correlation, it may not reflect a decision by high volatility individuals to marry each other. For example, this could reflect "contagious" income volatility (high volatility people make their spouse's high volatility). Since this matching is on transitory volatility, it could merely reflect the tendency of people with substantial uncertainty (and therefore measurement error) about their own income to also have uncertainty about their spouse's income.

excess variation in co-movement parameters, which are identified as the product of five-year changes for husbands and one-year changes for wives (equation 21). These volatilities are highly positively correlated (covariance of 0.0128, which corresponds to a correlation of 12 percent). Husbands with high five-year volatility have wives with high one-year volatility.[34] Note that this correlated person-level variation in income volatility explains most of the excess variation in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$ (equivalent to $cov\left(\left(y^{H}_{i,t+n}-y^{H}_{i,t-m-1}\right)^{2},\left(y^{W}_{i,t}-y^{W}_{i,t-1}\right)^{2}\right)\approx 0.018$). The bowing shape seen in Figure 3 and the "excess variance" in the sample permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$, can be attributed mostly to the high correlation between the volatilities of husbands' and wives' incomes. This reflects correlated transitory, not permanent, variances.[35]

4.3 Between-Couple Variation in δ

In addition to the between-couple correlated heterogeneity in volatility discussed in Subsection 4.2, an alternative source of between-couple variation is heterogeneity in couples' average permanent covariance, $mean_i\left[\delta^{\omega}_{i,t}\right]$. Some couples may have more co-movement than others. For example, couples who work in the same firm or industry will likely have highly positive co-movement (Shore and Sinai, 2006). To investigate between-couple variation in δ^{ω}, I calculate the average long-short product estimate of the permanent covariance for each couple, $mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]$. The between-couple variation in this estimate, $var\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)$, will reflect variation in the permanent covariance parameter, $var\left(mean_i\left[\delta^{\omega}_{i,t}\right]\right)$, and the variance of estimation error, $var\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]-mean_i\left[\delta^{\omega}_{i,t}\right]\right)$.

To test the null hypothesis that there is no between-couple variation in the true parameters, $var\left(mean_i\left[\delta^{\omega}_{i,t}\right]\right)=0$, I compare the sample value of $s.d.\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)_{true}$ in the data to the distribution of sample values of $var\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)_{wsbs}$ obtained from the wife-swap bootstrap. These results are presented in the third row of Table 11. We reject the hypothesis that $s.d.\left(mean_i\left[\delta^{\omega}_{i,t}\right]\right)=0$. There is more cross-couple variation in estimates from the actual data ($s.d.\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)_{true}=0.1197$, observation-weighted) than from the wife-swap bootstrap reference distribution ($s.d.\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)_{wsbs}$ ranges from 0.1053 to 0.1177 in 1,000 draws) for which all couples have $mean_i\left[\delta^{\omega}_{i,t}\right]=0$ by construction. Here, we can reject the hypothesis that

[34] While this finding could be contaminated slightly with correlated mis-measurement stemming from heterogeneity in ρ over time within couples, this will be negligible.

[35] The covariance of average long-short product estimates of couples' permanent variances (given in the lower-right corner of Table 12),

$$cov\left(mean_i\left[\left(\hat{\sigma}^{\omega,H}_{i,t}\right)^{2}\langle\times\rangle\right],mean_i\left[\left(\hat{\sigma}^{\omega,W}_{i,t}\right)^{2}\langle\times\rangle\right]\right),$$

is small, with a correlation of only 2 percent. Since permanent volatilities are not highly correlated, the correlation in couples' squared changes in income must be attributed to correlated transitory volatility.

Table 13: "Wife-Swap Bootstrap" Test with Assortative Matching in Income Volatility: Comparing the Distribution of Permanent Covariance Estimatesfor Actual and Volatility-Matched Wife-Swapped Data

Distribution Swapping Feature	Actual Data	Wife-Swap Bootstrap			
		Random		Similar Volatility	
		Mean	St. Dev.	Mean	St. Dev.
Mean	-0.0072	-0.0006	0.0013	-0.0009	0.0014
St. Dev.	0.2862	0.2683	0.0048	0.2842	0.0042
St. Dev of Lifetime Average	0.0745	0.0719	0.0021	0.0760	0.0020
Observations	35,628	32,918	34	33,072	45

The first column presents features of the distribution of long-short product estimates of the permanent covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$, for all couples with at least 10 years of data. Data is as in the first column of Table 11, though here couples with too few observations are excluded. The next two columns present the mean and standard deviation (over 1,000 bootstrap draws) of various moments from a wife-swap bootstrap, where couples are matched randomly subject to the constraint that both members of a pair must have a similar (within the same 5-year bin) number of years of data. The final two columns present the mean and standard deviation of various moments from a volatility-matched wife-swap bootstrap. Couples are paired at random among the set of possible partners within the same bin based on volatility and years of data. Bins are computed by separating couples into groups similar (within the same 5-year bin) number of years of data. Within these year-based groups, couples are then grouped according to the lifetime-average raw variance quintile of each spouse. The first and second rows present the mean and standard deviation of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$, $mean\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]$ and $s.d.\left(\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right)$. The third row presents the observation-weighted standard deviation of couple-specific lifetime averages of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$, $s.d.\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle\right]\right)$.

$s.d.\left(mean_i\left[\delta^{\omega}_{i,t}\right]\right)=0$, so there is significant between-couple variation in the permanent covariance.

The problem with this approach is that it ignores the impact of assortative mating on estimation error. By construction, the bootstrapped distributions are formed from randomly paired couples so that couples' volatilities will be independent. Correlated volatilities in the actual data could make estimation error larger in the actual data than in the wife-swapped data. To look for excess variation in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle {}^{HW}_{\times}\right\rangle$ beyond what would be expected given the pattern of assortative mating in income volatility, I change the cohort within which couples are swapped in the bootstrap. In Subsection 3.2, husbands and wives were paired at random. Here, I separate all couples with similar numbers of years of data into 25 bins based on the lifetime-average income volatility of the husband and wife.[36] Random pairing is done within these year and volatility bins. Once the wife-swap bootstrap randomly pairs husbands with wives whose raw variances similar to their actual wife and

[36] First, I separate couples into bins based on the number of years of data available for the couples. I exclude all couples with fewer than 10 observations, and then place all couples with 11 to 15 years of data together, all couples with 16 to 20 years of data together, and so on. Within these years-of-data-based bins, I compute the 20^{th}, 40^{th}, 60^{th}, and 80^{th} percentiles for $mean_i\left(y^H_{i,t+n}-y^H_{i,t-m-1}\right)^2$, $mean_i\left(y^W_{i,t}-y^W_{i,t-1}\right)^2$. Within these years-of-data-based bins, I separate couples into 25 bins based on each spouse's volatility quintile (5 for the husband x 5 for the wife).

vice versa, variation from the actual data in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$ or $mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle\right]$ is now well within the range of the new conditional reference distribution.

Table 13 shows that there is no evidence of latent heterogeneity in the permanent covariance or its couple-specific average, $\delta^{\omega}_{i,t}$ or $mean_i\left[\delta^{\omega}_{i,t}\right]$, once we account for the tendency of high-volatility individuals to be married to one-another. The first column presents the distribution of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$ for actual data. The first and second rows present the mean and standard deviation of covariance estimates, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$. The third row presents the observation-weighted standard deviation of couple-specific lifetime averages, $s.d.\left(mean_i\left[\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle\right]\right)$. The second and third columns present the mean and standard deviation, respectively, from 1,000 bootstrap draws where couples are paired at random among the set with similar years of data; there is no volatility-based matching. The fourth and fifth columns present the mean and standard deviation, respectively, from 1,000 bootstrap draws with year- and volatility-based matching. The first row shows that accounting for volatility-based matching will not explain the negative average permanent covariance. The average permanent covariance estimate in actual data is strongly negative, similar to previous results (though not identical since couples with fewer than 10 years of data are excluded), and well outside the reference distribution provided by either bootstrap. However, the second and third rows show that the standard deviation of permanent covariance estimates and lifetime-average permanent covariance estimates in the actual data are almost exactly what would be expected given the degree of volatility-based matching in the data. We therefore fail to reject the null hypothesis that of no heterogeneity in the permanent correlation, $var\left(\rho_{i,t}\right)=0$.[37] This does not mean that there is no heterogeneity, or that no variables will predict variation in $\delta^{\omega}_{i,t}$ or $mean_i\left[\delta^{\omega}_{i,t}\right]$, only that such variables account for too little variation in the data given the power of the wife-swap bootstrap test.

4.4 Within-Couple Variation in δ

Besides the latent variation in δ^{ω} examined in Subsection 4.3, the degree of co-movement varies systematically based on observable variables. Table 14 presents results from regressions that predict variation in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$. The first four columns present life-cycle patterns in permanent covariance estimates, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$, as shown by the positive coefficient on the number of years of marriage.

[37] Hess (2004) argues that between-couple variation in $mean_i\left[\delta_{i,t}\right]$ predicts divorce, and uses the sample mean, $mean_i\left[\left(y^H_{i,t}-y^H_{i,t-m}\right)\left(y^H_{i,t}-y^H_{i,t-m}\right)\right]$ (instrumenting for this sample value) to proxy for $mean_i\left[\delta_{i,t}\right]$. Since I fail to reject the hypothesis that there is between-couple variation in $\delta^{\omega}_{i,t}$, it suggests that the between-couple variation in $mean_i\left[\delta_{i,t}\right]$ that predicts divorce in Hess (2004) is variation in the transitory covariance. If this is true, it admits the possibility that Hess's results are driven by a tendency of couples with (predictably, using instruments) more positively correlated measurement error having higher divorce rates. This is sensible to the degree that the tendency of a spouse to mis-state his (or her) income relative to his partner contains information about marriage quality.

Table 14: Determinants of Permanent Covariance

Dependent Variable: Long-Short Product Estimator of the Permanent Covariance, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
# of Years Married	0.0018 (2.13)*	0.0025 (2.03)*	0.0034 (3.24)**	0.0032 (2.12)*	0.0016 (1.57)	0.0028 (2.01)*	0.0017 (1.68)
(# of Years Married)^2	0.0000 (1.41)	-0.0001 (1.52)	-0.0001 (2.40)*	-0.0001 (1.38)	0.0000 (1.09)	0.0000 (1.22)	0.0000 (1.40)
Year	0.0005 (2.42)*	0.0004 (1.93)			0.0005 (2.29)*		0.0005 (2.38)*
# of Kids Age 0-1					-0.0086 (1.92)	-0.0034 (0.65)	-0.0122 (2.66)**
# of Kids Age 2-6					-0.0001 (0.03)	0.0041 (1.19)	0.0035 (1.35)
# of Kids Age 7-20					-0.0005 (0.34)	0.0015 (0.58)	-0.0001 (0.07)
Husband's # of Years of Education							-0.0012 (1.68)
Wife's # of Years of Education							-0.0002 (0.20)
Fixed Effects?	no	no	yes	yes	no	yes	no
Predicted Product?	no	yes	no	yes	no	no	no
Observations	38,120	38,120	38,120	38,120	38,120	38,120	38,010
R^2	0.0000	0.0010	0.0000	0.0000	0.0010	0.0010	0.0010

Table shows results from regressions that predict $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$ with covariates. Some regressions include as covariates a quadratic in the number of years of marriage, the calendar year, the number of children in the household of various ages, the number of years of education for the husband and wife, couple-specific fixed effects, and controls for the product of predicted values of one-year changes in the wife's income and five-year changes in the husband's income. Absolute value of t-statistics in parentheses. "*" indicates significance at the 5% level; "**" indicates significance at the 1% level.

Incomes co-move more and more positively over time for first twenty-five years of marriage. This is apparent graphically in Figure 5. Figure 5 is obtained by regressing long-short product estimates of the permanent covariance and variances, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$, $\left(\hat{\sigma}^{\omega,s}_{i,t}\right)^2$, and $\left(\hat{\sigma}^{\varepsilon,s}_{i,t}\right)^2$, separately on six-degree polynomials in the number of years of marriage. These coefficients are used to obtain predicted covariance and variance values for each year of marriage. Figure 5 plots the implied correlation for each year of marriage obtained from this procedure, with confidence intervals obtained using the "delta method." Permanent innovations to income are strongly negatively correlated early in marriage, with a correlation of nearly -40 percent. This correlation increases nearly monotonically for the first 25 years of marriage and reaches its maximum correlation of roughly 10 percent. The correlation falls later in marriage and becomes negative again after 30 years of marriage and continues to fall further. This finding is consistent with results from Shore (2006), which uses repeated observations on the cross-sectional covariance of couples' incomes to show that couples incomes are

Figure 5: Co-Movement of Couples' Incomes Over the Life-Cycle

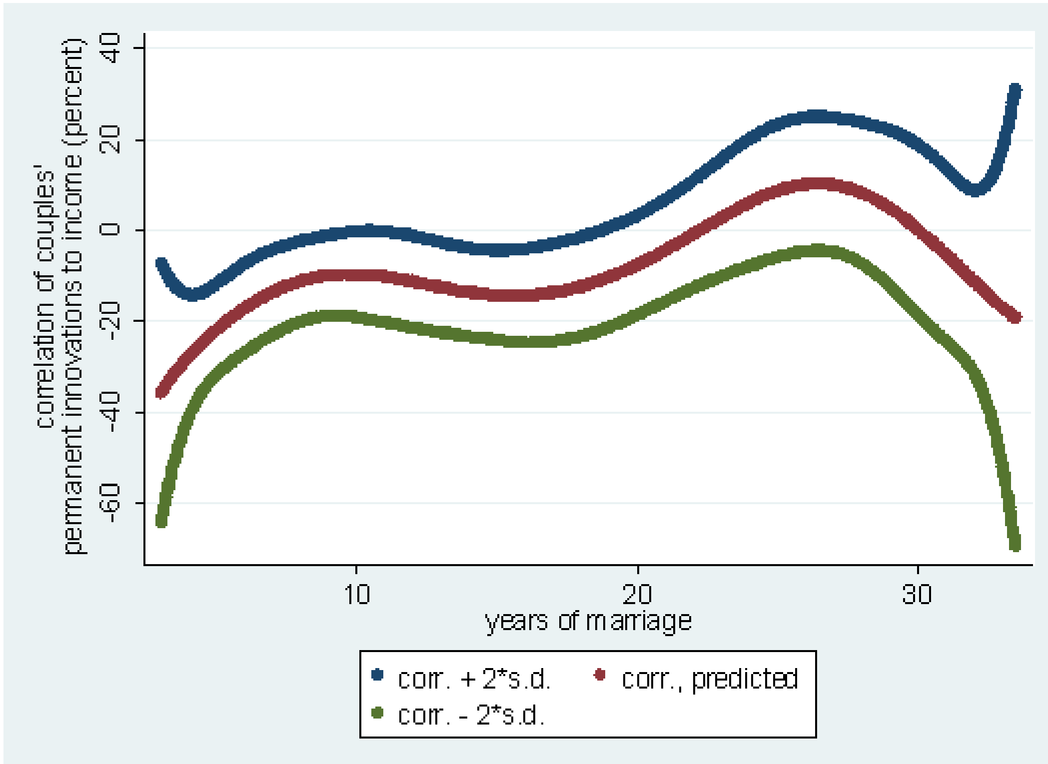

Figure plots the predicted correlation of permanent innovations to income as a function of the number of years of marriage. These are calculated as follows. First, $\left(\hat{\sigma}_{i,t}^{\omega,H}\right)^2 \left\langle_\times\right\rangle$, $\left(\hat{\sigma}_{i,t}^{\omega,W}\right)^2 \left\langle_\times\right\rangle$, and $\left(\hat{\delta}_{i,t}^{\omega}\right) \left\langle_\times^{HW}\right\rangle$ are calculated for each observation as in equations 3 and 10. These are each regressed on a six-degree polynomial in the number of years of marriage, and a predicted value of each is then computed for each possible year of marriage. Correlations are then computed as the ratio of the predicted values of $\left(\hat{\delta}_{i,t}^{\omega}\right) \left\langle_\times^{HW}\right\rangle$ to the square root of the product of $\left(\hat{\sigma}_{i,t}^{\omega,H}\right)^2 \left\langle_\times\right\rangle$ and $\left(\hat{\sigma}_{i,t}^{\omega,W}\right)^2 \left\langle_\times\right\rangle$ predicted vales. The two standard error confidence intervals are computed using the delta method.

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

This effect reflects within-couple variation over time, as Table 14 shows that this pattern is robust to the inclusion of couple fixed-effects (columns 3 and 4). The pattern is robust to inclusion of controls for the presence of children (columns 5, 6 and 7). However, there is (statistically weak) evidence that the permanent covariance is lower in the presence of children (particularly young ones), and this seems to account for some but not all of the life-cycle pattern we observe. This life-cycle pattern is also present when looking at the co-movement of husbands' log incomes with wives' hours.

The most obvious candidate interpretation of this life-cycle pattern is that it reflects life-cycle changes in the relative importance of various economic benefits of marriage. Early in marriage, it may be relatively important that one spouse's production is a substitute for the production of the other; increases in income by one spouse will then be positively correlated with changes in market work and positively correlated with changes in home production by the other. This would imply the negative co-movement found early in marriage and in the presence of children. Later in marriage, complementarity of leisure become more important. Working less or retiring early is more appealing when you can spend the additional leisure time with your spouse, which would explain the increasingly positive co-movement of couples' incomes nearing retirement.

The degree of coordination is also increasing with calendar time, as shown in the first column of Table 14. Since this column also includes controls for years of marriage, the time effect reflects changing cohorts. Couples who married later tend to have more positively correlated innovations to permanent income. This may reflect couples increasingly meeting while at work or while studying in the same narrow field, in which case their incomes will naturally be more highly correlated.

While the results in Table 14 reveal systematic patterns in the permanenet covariance, they say nothing about its stochastic process Here, I compare the time-series properties of permanent covariance estimates, $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$, from actual data to those from wife-swap bootstrapped simulated data. The fourth row presents the average within-couple linear time-trend in $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$. This row shows that $\delta^{\omega}_{i,t}$ increases over time on average (significantly positive trend of 0.0017 per year on average), a pattern established already in Table 14 and in Figure 5. There is substantial between-couple variation in the within-couple time-trend (fifth row), though no more than is found in the "wife-swap bootstrapped" distributions of randomly paired couples. There is no evidence that couples differ in the life-cycle pattern of $\delta^{\omega}_{i,t}$. Also, I in the sixth row I present the coefficient of a regression of $\left(\hat{\delta}^{\omega}_{i,t}\right)\left\langle{}^{HW}_{\times}\right\rangle$ on $\left(\hat{\delta}^{\omega}_{i,t-1}\right)\left\langle{}^{HW}_{\times}\right\rangle$. While correlated estimation error makes the coefficient strongly negative, it is no different than for randomly paired couples. In short, there is no evidence of any interesting time-series patterns in $\delta^{\omega}_{i.t}$ except the life-cycle patterns that are common to all couples. If $\delta^{\omega}_{i.t}$ is viewed a measure of marriage quality that predicts divorce (as in Hess (2004)), the lack of heterogeneity in trends in $\delta^{\omega}_{i.t}$ suggest that changes over time in $\delta^{\omega}_{i.t}$ cannot explain why some couples get divorced and others do not.

5 Conclusion

This paper has developed a technique to estimate the covariance of couples' permanent innovations to income. In comparing the short-term changes in one spouses income to the long-term changes in the other's income that span them, the estimate avoids contamination from couples' correlated measurement errors. The product of these changes provides a measure of the permanent covariance that may prove useful in predicting a variety of outcomes. Since co-movement provides one measure of the economic benefits of marriage and their dynamics, it is natural to use a measure of co-movement to predict divorce, saving, fertility, and a variety of other household outcomes. I develop a "wife-swap bootstrap" test that rejects the hypothesis that couples' incomes evolve independently. More broadly, this technique provides a way to differentiate noise from heterogeneity, and therefore to correct for the attenuation bias induced by including a noisy measure of co-movement as an explanatory variable in a regression.

In addition to its methodological contribution, this paper documents several stylized facts about co-movement. First, the correlation of permanent innovations to income is roughly -10 percent on average. This relationship is apparent in responses to both large and small changes in income. When a wife increases her earnings in a given year, it implies a downward shift in the entire distribution of her husband's changes in income; he is more likely to have suffered a large long-term drop in income and less likely to have enjoyed a large long-term increase in income. This average affect masks dramatic heterogeneity, as there is too much variation in sample moments to be explained by measurement error alone. This heterogeneity is primarily explained by assortative mating in transitory income volatility, though there is also substantial life-cycle variation in co-movement. There is no evidence of latent heterogeneity in co-movement parameters.

There are more than one way to interpret results about co-movement. Co-movement could reflect the innate diversification benefits of marriage, the covariance that the two individuals labor income processes would have even if they were not married. Such correlated exogenous shocks could plausibly explain the positive covariance of couples who work in the same firm or industry, but they are unlikely to explain the substantial negative correlation for most couples. Instead, the negative correlation between couples changes in income likely reflects coordination, an endogenous choice to jointly change labor income. This is confirmed by the finding the permanent covariance of husbands' incomes is negative not only with wives' incomes but also with wives' hours worked and wives' labor force participation.

Changes in the motivation for coordination provides one possible explanation of the life-cycle variation in co-movement. Early in marriage, the substitution of spouses' production (both at home and at work) may be relatively important and drive negative co-movement, as positive shocks or choices for one spouse drive negative changes for the other. When one spouse decides they want to spend more time taking care of the children, their partner will have less need of doing the same and may choose to focus on market work. Later in marriage, the complementarity of spouses' leisure may be relatively important and drive positive co-movement. Couples will increase and decrease their effort or hours together if leisure time is pointless without someone to spend it with. Life-cycle variation in the relative importance of these effects may explain the life-cycle variation in co-movement. As such, life-cycle variation in co-movement may reflect life-cycle variation in the relative importance of various economic benefits of marriage.

References

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

Alvarez, J., M. Browning, and M. Ejrnaes (2001): "Modelling income processes with lots of heterogeneity," University of Copenhagen Working Paper.

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.

Fisher, R. A. (1935): *The Design of Experiments*. Oliver and Boyd, Edinburgh.

Hall, R. E., and F. S. Mishkin (1982): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 50(2), 461–482.

Hess, G. D. (2004): "Marriage and Consumption Insurance: What's Love Got to Do with It?," *Journal of Political Economy*, 112(2), 290–318.

Jepsen, L. K., and C. A. Jepsen (2002): "An Empirical Analysis of the Matching Patterns of Same-Sex and Opposite-Sex Couples," *Demography*, 39(3), 435–453.

Kennedy, P. E. (1995): "Randomization Tests in Econometrics," *Journal of Business and Economic Statistics*, 13(1), 85–9.

Lundberg, S. (1985): "The Added Worker Effect," *Journal of Labor Economics*, 3(1), 11–37.

Meghir, C., and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

Romano, J. P. (1989): "Bootstrap and Randomization Tests of some Nonparametric Hypotheses," *The Annals of Statistics*, 17(1), 141–159.

Shore, S. H. (2006): "For Better, For Worse: Marriage and the Business Cycle," University of Pennsylvania Working Paper.

Shore, S. H., and T. Sinai (2006): "Commitment, Risk, and Consumption: Do Birds of a Feather Build Bigger Nests?," earlier version is NBER Working Paper No. 11588.

Stephens, M. (2002): "Worker displacement and the added worker effect," *Journal of Labor Economics*, 20(3), 504–37.