

Georgia State University

ScholarWorks @ Georgia State University

UWRG Working Papers

Usery Workplace Research Group

5-1-2013

The Co-Movement of Couples' Incomes

Stephen H. Shore

Georgia State University, sshore@gsu.edu

Follow this and additional works at: https://scholarworks.gsu.edu/uwrg_workingpapers

Recommended Citation

Shore, Stephen H., "The Co-Movement of Couples' Incomes" (2013). *UWRG Working Papers*. 72.
https://scholarworks.gsu.edu/uwrg_workingpapers/72

This Article is brought to you for free and open access by the Usery Workplace Research Group at ScholarWorks @ Georgia State University. It has been accepted for inclusion in UWRG Working Papers by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

Working Paper 2013-5-1
May 2013

The Co-Movement of Couples' Incomes

Stephen H. Shore
Georgia State University

This paper can be downloaded at: <http://aysps.gsu.edu/uwrg/2013-0>

The Co-Movement of Couples' Incomes

Stephen H. Shore^{*†}

May 2013

Abstract

While there is a large literature on how individual incomes move over time, we know much less about couples' joint income dynamics. Current research on individual income dynamics has increasingly considered heterogeneity – do all individuals' incomes evolve in the same way, or does a particular individual's income evolve in the same way throughout their life? This paper considers the analogous questions for couples – do all couples' incomes move together in the same way, or does a particular couple's incomes move together in the same way throughout their marriage? In particular, I find evidence of correlated volatility; husbands with volatile incomes tend to have wives with volatile ones. I find weaker evidence for heterogeneity in the correlation of husbands' and wives' income changes, with some couples incomes moving together while others moving in opposite directions. Couples' income changes are negatively correlated early in marriage, particularly when young children are present, and become more positively correlated over time.

*Georgia State University; sshore@gsu.edu

†JEL Classification: D31 - Personal Income, Wealth, and Their Distributions; D13 - Household Production and Intrahousehold Allocation; C33 - Multiple Equation Models with Panel Data

‡keywords: couples, income dynamics, heterogeneity, wife-swap bootstrap

1 Introduction

There is a very large literature on individual income dynamics, on how individuals' incomes evolve over time. Much of this literature is focussed on income volatility, the variance of income changes.¹ Recent work in this area has focused on identifying latent heterogeneity in volatility; some people may face income changes with larger variances than others. (Meghir and Pistaferri, 2004; Alvarez, Browning, and Ejrnaes, 2001; Jensen and Shore, 2011, 2012)

The literature on couples' joint income dynamics – how couples' incomes move together – is much smaller. (Lundberg, 1985; Cullen and Gruber, 2000; Hyslop, 2001; Dynan, Elmendorf, and Sichel, 2007; Shore, 2010) Just as recent research has focussed on heterogeneity in individuals' income dynamics, this paper considers heterogeneity in couples' joint income dynamics; do all couples' incomes move together in the same way? Heterogeneity in couples' joint income dynamics could reflect assortative mating in volatility, so that individuals with volatile incomes tend to marry each other;² it could also reflect heterogeneity in co-movement, so that some couples' incomes move together while other couples' incomes move in opposite directions. Both of these phenomenon show up in the cross-section of couples' income changes as bivariate kurtosis (Mardia, 1970, 1974, 1980), the tendency of large (absolute) income changes for husbands and wives to coincide. In years in which a husband's earnings changes substantially (either rising or falling), his wife's income tends to change substantially (either rising or falling) as well. However, correlated volatility can be separated from heterogeneity in co-movement with panel data or other covariates given certain

¹Papers on this subject include Hall and Mishkin (1982); Gottschalk and Moffitt (1994); Moffitt and Gottschalk (1995); Daly and Duncan (1997); Carroll and Samwick (1997); Dynarski and Gruber (1997); Cameron and Tracy (1998); Geweke and Keane (2000); Haider (2001); Gottschalk and Moffitt (2002); Batchelder (2003); Hacker (2006); Comin and Rabin (2006); Gottschalk and Moffitt (2006); Hertz (2006); Winship (2007); Bollinger and Ziliak (2007); Bania and Leete (2007); Dahl, DeLeire, and Schwabish (2007); Shin and Solon (2008).

²Alternatively marriage could make income volatility for husbands and wives more similar than it would have been had they not wed.

assumptions.

These distinctions are important for understanding the economic effects of coupling. Positive assortative mating in volatility may be optimal given positive assortative mating in risk-aversion, as predicted by Chiappori and Reny (2006). Risk tolerant individuals may choose risky income streams for themselves, and also seek partners with risky income streams (leading to positive assortative mating on risk-aversion). Conversely, absent heterogeneity in risk-aversion, we would expect negative assortative mating in volatility, as the cost of marrying a high-risk spouse is lower for a low-risk person. Heterogeneity in the covariance of couples' income changes is important because it suggests differences across couples in the risk-sharing benefits of marriage. Nordblom (2004) shows that some of this variation in the diversification benefits of marriage may stem from differences in legal regimes that may affect the degree of commitment and cooperation while Chami and Hess (2005) shows that there is cross-state variation stemming from differences in states' levels of undiversifiable risk. Hess (2004) shows that such variation can predict divorce.

Changes over time in couples' joint income dynamics suggest changes in labor and leisure complementarities over the life cycle. This paper shows that early in marriage, particular when young children are present, couples' incomes are negatively correlated. Couples' income changes become more positively correlated as the number of years a couple has been married increases. One possible 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, one spouse's production may be a substitute for the production of the other; increases in income by one spouse will tend to coincide with increases in home production (and decreases in market work) for their partner. This suggests that the specialization in production described in Becker (1973) is particularly dominant early in marriage. Later in marriage, complementarity of leisure may become more important; this could explain the increasingly positive

co-movement of couples' incomes nearing retirement. This phenomenon is studied most frequently in the context of couples' joint retirement decisions, which frequently coincide. (Hurd, 1990; Burtless, 1990; Gustman and Steinmeier, 2000; Maestas, 2001; Michaud, 2003; Casanova, 2010) Simultaneous retirement is frequently motivated by leisure complementarities: leisure time in retirement is more enjoyable if you can share this leisure time with your spouse.

These ideas are applied to couples' income data from the Panel Study of Income Dynamics. In the data, wives' income changes are approximately uncorrelated with their husbands' income changes.³ However, they are not independent, as couples' squared income changes are positively correlated; there is bivariate kurtosis, so that husbands' large income changes (increases or decreases) tend to coincide with wives' large income changes (increases or decreases). A "wife-swap bootstrap" test strongly rejects the independence of couples income streams, finding substantial bivariate kurtosis. This procedure is appropriate when the pair of random variables (here, husbands' and wives' income changes) are unconditionally uncorrelated but each spouse's income changes may be autocorrelated (as in this case). This test is designed to measure the amount of matching that can be seen in couples' joint income dynamics, relative to a null hypothesis of random pairing; this paper strongly rejects the hypothesis that couples' joint income dynamics resemble what would be expected from random pairing. By comparing results for various measures of income and hours worked, much of this stems from large changes in wives' hours (and not wages per hour) coinciding with large changes in their husbands' incomes.

Correlated volatility can explain much of the observed bivariate kurtosis; wives whose income shocks have large variances tend to be married to husbands whose

³See sample moments from Table 3. For example, the unconditional sample correlation of simultaneous one-year changes in husbands' and wives' incomes is approximately -0.2 percent; the hypothesis that the covariance is zero cannot be rejected. Results are very similar for the transitory measure of couples' income changes. For the permanent measure of couples' income changes, the correlation is approximately -2 percent; the hypothesis that these covariances are zero can separately be rejected at the 95 percent confidence level but not the 99 percent confidence level.

income shocks also have large variances. Correlated variance parameters explain more than 28 percent or 90 percent (depending on the measure of income changes) of the observed bivariate kurtosis. This looks like the positive assortative mating on income risk of interest to Chiappori and Reny (2006).

Heterogeneity in co-movement – with some couples’ incomes moving together while other couples’ incomes moving in opposite directions – is also present. This covariance heterogeneity explains 10 percent to 33 percent of bivariate kurtosis.

2 Data

Data are drawn from the Panel Study of Income Dynamics (PSID). The PSID is a nationally representative panel of U.S. households that has tracked families annually from 1968 to the present. Data are not collected in even-numbered years after 1997; this paper uses data collected through 2005. However, since most analyses use one-year income changes, only data through 1997 will be used in most circumstances. The PSID includes data on households, including household food consumption and the education, income, hours worked, employment status, and age of husbands and wives. I use annual labor income as a measure of income. I restrict the sample to married couples, to couples where the marriage is the husband’s first, 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 for no more than 35 years.

I remove the predictable (to the econometrician) component of income and examine the time series properties of the unpredictable component, excess log income. As is common in the literature, this excess log income is the residual from a least-squares regression of the natural log of labor income (for either the husband or the wife) on the following 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.⁴ 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).

The residuals from this regression are Winsorized at the 5th and 95th percentiles, so that residuals below the 5th percentile are replaced by the 5th percentile value and those above the 95th percentile are replaced by the 95th percentile value. At the same time, values omitted from the initial regression because real annual income was below \$1,000 are given the 5th 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. One-year changes are demeaned.

Table 1 presents summary statistics on one-year changes in excess log income for husbands and wives. Note that most one-year excess log income changes are relatively small. The inter-quartile ranges for wives (x_{it} from -10 percent to 8 percent) and husbands (y_{it} from -8 percent to 10 percent) are modest. However, there are occasional very large changes in income, so that the standard deviations of one-year income changes (55 percent and 32 percent, respectively) are much larger than the inter-quartile ranges. These fat-tails could be the result of fat-tailed shocks (occasional large income changes) or heterogeneity (some observations are expected to have larger variances while others are expected to have smaller variances, though conditional on these variances tails are not fat).

⁴This procedure ensures that predictable income changes such as the typical life-cycle pattern of income are not included in measuring couples' idiosyncratic joint income dynamics.

Table 1: Distribution of Spouses' One-Year Change in Excess Log Income

Spouse	Wives	Husbands
Mean	0	0
St. Dev.	0.5490	0.3184
Observations	20,762	20,762
Minimum	-2.8499	-1.8283
5 th Percentile	-0.8390	-0.5258
25 th Percentile	-0.0955	-0.0796
50 th Percentile	-0.0179	0.0064
75 th Percentile	0.0806	0.0987
95 th Percentile	0.9305	0.4708
Maximum	2.8141	1.8410
lag	Autocorrelation	
1 Year	-0.2133	-0.3193
2 Years	-0.0766	-0.0445
3 Years	-0.0251	-0.0217
4 Years	-0.0395	-0.0169

This table presents the distributions of one-year changes in Winsorized excess log income for wives and husbands, x_{it} and y_{it} , respectively. The construction of Winsorized excess log incomes is explained in the text. In brief, annual log labor incomes for husbands and wives are separately regressed on a host of covariates. The residuals from these regressions are Winsorized at the 5th and 95th percentiles. These changes are de-meanded, so means are zero by construction. The median one-year change would be exactly zero in the absence of de-meaning, so -1 times the median values gives the average annual change. The sample is limited to observations where data exists in the six years prior to the year in question.

The patterns of autocorrelation are also presented in Table 1. One-year increases in income tend to be followed by decreases in the following year for both husbands and wives, with very small decreases in subsequent years. While small, autocorrelations at lags greater than one year are larger here than in Abowd and Card (1989), primarily because income changes are Winsorized. Another noteworthy result is that one spouse's income changes are nearly uncorrelated with lagged changes in the other's income.⁵

⁵For example, the correlation between y_{it} and x_{it-1} is 0.0001 and the correlation between y_{it} and x_{it-2} is -0.0028.

3 Results

3.1 Income Dynamics

Here, I present a standard income process. Model parameters from this process may differ across couples and over time. While more complex income processes are possible, it is standard in the literature to assume that excess log income is composed of permanent (p) and transitory (ε) components:

$$\begin{aligned} z_{yit} &= p_{yit} + \varepsilon_{yit}; \\ p_{yit} &= p_{yiT_0^i} + \sum_{\tau=T_0^i+1}^t \omega_{yi\tau}. \end{aligned} \tag{1}$$

Here, z_{yit} refers to the excess log income of the husband in household i in year t . The same process could be applied to wives as well, with x s replacing y s. x_{it} and y_{it} will be defined as changes in excess log income over an interval, $x_{it} \equiv z_{xit} - z_{xit-k}$ and $y_{it} \equiv z_{yit} - z_{yit-k}$. In equation 1, transitory income, ε_{yit} , is assumed to be i.i.d. with variance $(\sigma_{y\varepsilon}^2|it)$; permanent income, p_{yit} , is assumed to have a unit root so that innovations to permanent income, $p_{yit} - p_{yit-1} = \omega_{yit}$, are i.i.d. with variance $(\sigma_{y\omega}^2|it)$. Subsequently, “transitory variance” refers to the variance of transitory income, $(\sigma_{y\varepsilon}^2|it)$; “permanent variance” refers to the variance of innovations to permanent income, $(\sigma_{y\omega}^2|it)$. These conditional variances may differ across individuals and over time.

If husbands’ and wives’ incomes individually evolve as in equations 1, it is natural to consider the joint income process where couples’ income shocks may be correlated. For couple i at time t , I consider $E[\omega_{xit}\omega_{yit}] \equiv (\sigma_{xy\omega}|it)$ and $E[\varepsilon_{xit}\varepsilon_{yit}] \equiv (\sigma_{xy\varepsilon}|it)$, which I subsequently refer to as the “permanent covariance” and the “transitory covariance.” While husbands’ transitory shocks may be correlated with wives’ permanent ones, and *vice versa*, these cross-covariances are assumed to be zero here.

In this setting, I consider three $\{x_{it}, y_{it}\}$ measures to identify the variance-covariance

structure of different types of shocks: raw, permanent, and transitory. Each measure is named by the type of covariance identified by the product of husbands' and wives' income changes, $x_{it}y_{it}$. Couples' income change moments for each measure are shown in Table 3

1. Raw: The simplest measures of the variance or covariance of income changes come from contemporaneous one-year changes: $x_{rit} \equiv z_{xit} - z_{xit-1}$ and $y_{rit} \equiv z_{yit} - z_{yit-1}$. These income changes include both permanent and transitory components, so their squares and products will as well. From Table 3, the unconditional sample mean of $x_{rit}y_{rit}$ is close to zero, with an implied correlation of -0.2 percent (statistically insignificant difference from zero).
2. Permanent: To isolate the permanent covariance without contamination from the transitory variance, I consider the short-term change in a wife's income and the long-term change in her husband's income that spans this short term change: $x_{\omega it} \equiv z_{xit} - z_{xit-1}$ and $y_{\omega it} \equiv z_{yit+2} - z_{yit-3}$. So long as permanent shocks enter in over at most 2 periods and transitory shocks damp out in at most 2 periods (consistent with evidence from Abowd and Card (1989)), this measure isolates the permanent covariance even when the income process is much more general than the one specified here (Meghir and Pistaferri, 2004). From Table 3, the unconditional sample mean of $x_{\omega it}y_{\omega it}$ is slightly negative but close to zero, with an implied correlation of -2.6 percent (statistically different from zero at the 95 percent, but not the 99 percent, significance level).
3. Transitory: Under the specified income process, the transitory covariance can be identified by looking at the product of income changes for one spouse and their lag for the other spouse: $x_{\varepsilon it} \equiv z_{xit+1} - z_{xit}$ and $y_{\varepsilon it} \equiv z_{yit-1} - z_{yit}$. From Table 3, the unconditional sample mean of $x_{\varepsilon it}y_{\varepsilon it}$ is slightly negative but close to zero, with an implied correlation of -0.2 percent (statistically insignificant

difference from zero).

3.2 Determinants of Co-Movement

Figure 1: Co-Movement of Couples' Incomes Over the Life-Cycle

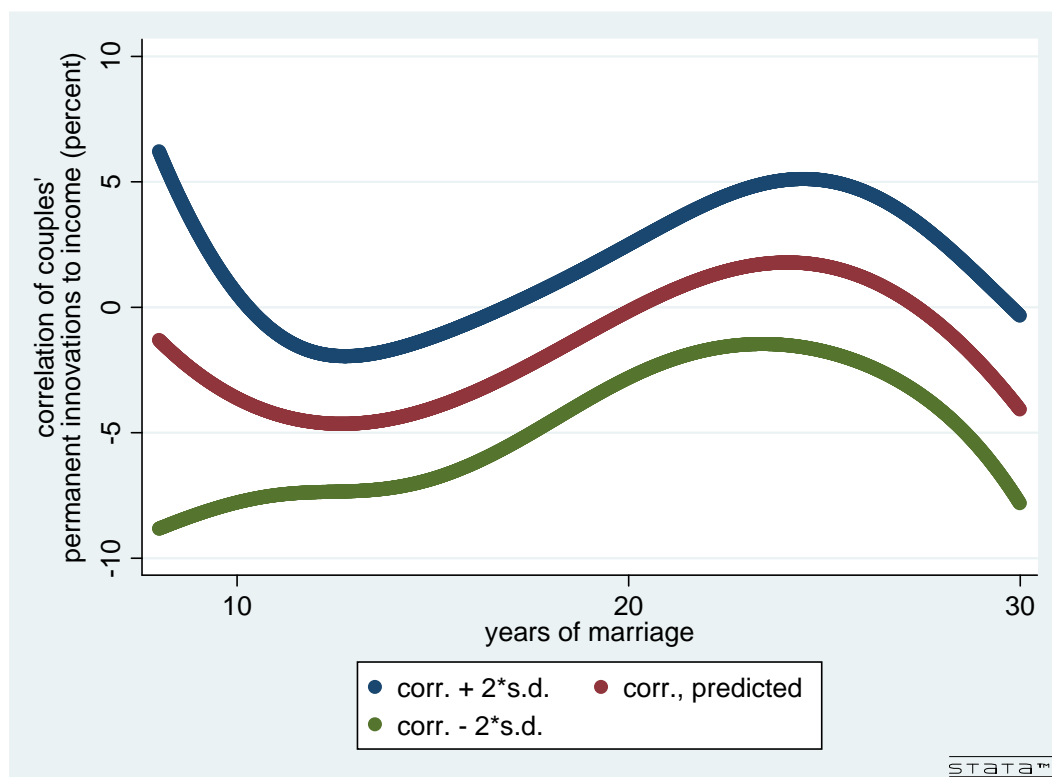


Figure plots the predicted correlation of permanent innovations to income as a function of the number of years of marriage. These are calculated as follows. First, the permanent covariance and permanent variances are calculated for each observation. These are each regressed on a three-degree polynomial 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. The two standard error confidence intervals are computed using the delta method.

While couples' co-movement is roughly zero on average (and insignificantly different from zero using the raw and transitory measures of co-movement), the correlation of husbands' and wives' income changes is not zero for every couple or zero at every

Table 2: Determinants of the Co-Movement of Couples' Incomes

Dependent Variable	Estimates of the One-Year Raw Covariance			Estimates of the Permanent Covariance		
	# of Years Married	0.0013*** (4.46)		0.0009*** (3.86)	0.0010** (1.99)	
# of Kids		-0.0035** (2.12)	-0.0015 (1.17)		-0.0000 (0.02)	0.0015 (0.79)
Husband's Years of Education			-0.0014** (2.02)			-0.0017* (1.69)
Wife's Years of Education			0.0016* (1.88)			0.0020* (1.66)
Fixed Effects?	yes	yes	no	yes	yes	no
Observations	20,762	20,762	20,762	15,478	15,478	15,478
R^2	0.0010	0.0002	0.0015	0.0003	0.0000	0.000

This table shows results from OLS regressions that predict permanent and one-year raw covariance estimates with covariates. t -statistics in parentheses. “*” “**” and “***” indicate significance at the 10%, 5%, and 1% levels, respectively.

point in the life cycle. In particular, there is strong life-cycle variation in co-movement. This is apparent in Figure 1, which is obtained by regressing permanent covariance estimates and variances separately on three-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 1 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. This correlation increases with the number of years of marriage. This finding is consistent with results from Shore (2010), which uses repeated observations on the cross-sectional covariance of couples' incomes to show that couples' incomes are negatively correlated early in marriage but positively correlated later in marriage. One possible 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 pro-

duction is a substitute for the production of the other; increasing in income by one spouse will tend to coincide by increasing home production and decreasing market work 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 may 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.

Table 2 presents results from regressions to predict co-movement with a host of covariates. The covariance of couples' income changes increases over the life-cycle of marriage. *Ceteris Paribus*, this increases the volatility of household income over time by reducing the diversification benefits of marriage. This will lead to increasing household income inequality over time for older couples (who have many years of compounded permanent shocks). While the presence of children reduces the covariance of couples' income changes, this can be explained fully by the number of years of marriage. There is weak evidence that that couples with high-education husbands and low-education wives have more negative covariances.

3.3 Heterogeneity in Couples' Joint Income Dynamics

The sample moments from Table 3 provides the moments needed to test for bivariate kurtosis, the tendency of couples large (absolute) income changes to coincide. The top panel of Table 4 presents the results of these tests, showing substantial and statistically significant bivariate kurtosis. The significance of the results is slightly higher using the "wife-swap bootstrap" test discussed in the Appendix. This test relaxes the assumption from the standard test that income changes are not autocorrelated; in the data, autocorrelations are negative for adjacent observations. The "wife-swap bootstrap" effectively provides a null hypothesis of how couples' incomes would jointly

Table 3: Sample Moments for Couples' Income Changes

Covariance Measure	Raw	Permanent	Transitory
x_{it} : wife's income change	One Year	One Year	One Year
y_{it} : husband's income change	Same One Year	Surrounding Five Years	Lagged One Year
x_{it}	0	0.0044	0.0005
x_{it}^2	0.3013	0.2970	0.3028
x_{it}^4	0.8560	0.8235	0.8540
y_{it}	0	0.0059	-0.0004
y_{it}^2	0.1014	0.1983	0.1022
y_{it}^4	0.1024	0.2173	0.1050
$x_{it}y_{it}$	-0.0004	-0.0055	0.0003
$x_{it}^2y_{it}^2$	0.0461	0.0692	0.0411
$x_{it}y_{it}x_{it-5}y_{it-5}$	0.0009	0.0017	0.0005
$x_{it}^2y_{it-5}^2$	0.0322	0.0593	0.0279
$y_{it}^2x_{it-5}^2$	0.0398	0.0765	0.0393
N	20,762	15,478	19,430

This table presents sample means over all i and t for which data on x_{it} and y_{it} are both available. z_{xit} is the excess log income of the wife from couple i in year t ; z_{yit} is the excess log income of the husband from couple i in year t . Sample sizes are smaller for the final three lead-lag moments. See text for details on variable construction.

evolve if husbands and wives were paired at random (but each spouse's income was free to evolve individually as it did in the data). The rejection of this null suggests that couples' large income changes tend to coincide far more than would be expected from random pairing.

Two possible sources of this pattern of bivariate kurtosis reflect heterogeneity in couples' joint income dynamics: correlated variances of husbands' and their wives' income changes, and heterogeneity in the covariance of husbands' and their wives' incomes. Appendix A shows how bivariate excess kurtosis can be decomposed into these components. Furthermore, that appendix shows how panel data can be used to bound the relative size of these components. The lower panel of Table 4 presents results that bound these potential sources of bivariate kurtosis.

Table 4: Sources of Bivariate Excess Kurtosis for Couples' Income Changes: Why Do Couples' Large (Absolute) Income Changes Coincide?

	Raw	Permanent	Transitory
excess unconditional bivariate kurtosis, $\hat{\kappa}_{xy}$	1.53	0.52	0.98
(z-stat)	(7.58)	(3.02)	(4.71)
("wife-swap bootstrap" z-stat)	(7.68)	(3.28)	(5.01)
	Fraction Explained by		
correlated variances, $cov_i((\sigma_x^2 i), (\sigma_y^2 i))$	> 38%	> 90%	> 28%
covariance heterogeneity, $var_i((\sigma_{xy} i))$	> 12%	> 33%	> 10%
excess conditional bivariate kurtosis, $\kappa_{xy} i$	< 40%	< -19%	< 55%

$x_{it} \equiv z_{xit} - z_{xit-1}$ and $y_{it} \equiv z_{yit} - z_{yit-1}$ if raw estimate; $x_{it} \equiv z_{xit} - z_{xit-1}$ and $y_{it} \equiv z_{yit+2} - z_{yit-3}$ if permanent estimate; $x_{it} \equiv z_{xit+1} - z_{xit}$ and $y_{it} \equiv z_{yit-1} - z_{yit}$ if transitory estimate. z-statistics are against the null hypothesis is that $\kappa_{xy} = 0$. The first z-statistic assumes that observations are independent over time and across individuals. The second z-statistic uses the "wife-swap bootstrap" explained in the text. This implicitly assumes that x_{it} and y_{it} are unconditionally uncorrelated but allows x_{it} (and also y_{it}) to be autocorrelated. The lower-bound on $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$ is calculated from the average of the sample covariance of x_{it}^2 and y_{it-5}^2 and the sample covariance of y_{it}^2 and x_{it-5}^2 . The lower-bound on $var_i((\sigma_{xy}|i))$ is calculated from the sample covariance of $x_{it}y_{it}$ and $x_{it-5}y_{it-5}$. The upper-bound on $\kappa_{xy}|i$ is calculated from these lower-bounds from equation 9. The percent of $\hat{\kappa}_{xy}$ explained by each of these components comes from equation 6 assuming that the other two components are zero.

Correlated variances of couples' income changes, $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$ explain much of the tendency of couples' large (absolute) income changes to coincide. Husbands whose incomes are volatile have wives whose incomes are volatile. The measure of this based on five-year leads and lags explains at least 38 percent, 90 percent, 28 percent of excess bivariate kurtosis for the raw, permanent and transitory measures of income changes, respectively. In the case of permanent variance, the large magnitude is particularly striking; husbands who receive large permanent shocks tend to have wives who receive large permanent shocks. This finding provides suggestive evidence of interest in models of assortative mating on risk (Chiappori and Reny, 2006)

While there is evidence of persistent covariances (and therefore covariance heterogeneity, $var_i((\sigma_{xy}|i))$), such heterogeneity is quantitatively smaller and accounts for far less of the observed excess bivariate kurtosis. If substantial heterogeneity in covariances exist in these data, they cannot be very persistent.

In the case of permanent income changes, observed excess bivariate kurtosis can be fully explained by correlated variances. In the case of transitory and raw income changes, substantial excess bivariate kurtosis remains unexplained. There is no way to know if this reflects parameter heterogeneity unexplained by the covariates used, reflects conditional excess bivariate kurtosis, or some combination.

Table 5: Raw Covariance of Husbands' Excess Log Incomes with Wives' Excess Hours, Excess Log Incomes, and Labor Force Participation

Husband's Variable: y		Log Income		
Wife's Variable: x	Log Income	Hours Worked	Log Income if in Labor Force	In Labor Force? 1 or 0
Implied Correl.	-0.2%	-3.3%*	2.6%*	-2.0%*
Excess Kurtosis	1.53*	0.91*	0.40	0.82*

Each column presents the estimates of the raw covariance, as discussed in the text. In each case, y refers to the Winsorized excess log income of the husband. The first row presents the implied sample correlation; the second row presents the implied excess kurtosis. "*" indicates significance at the 5% level.

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. This is significant because couples' incomes may covary either because of variation in wages, in hours worked, or labor force participation. Adjustment in hours worked (and relatedly in home production in leisure) have been shown to be an important source of benefit in marriage. (Vernon, 2010)

Table 5 presents estimates of raw covariance and excess bivariate kurtosis (ten-

dency of large absolute changes for the husband and wife to coincide) for several work-related variables for wives.⁶ The previous 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 hours (level of hours, not log hours, generated in the same way as excess log income) worked by wives, changes in excess log income for wives who remain working, and changes in labor force participation for wives.⁷ Note that all correlations are small and similar, between -4 and 3 percent. Excess bivariate kurtosis is greater in the “hours worked” and “in labor force” measures than for the “income if in labor force measure”; the hypothesis that there is no tendency of couples’ large income changes to coincide cannot be rejected conditioning on wives being in the labor force. This suggests that much of the variation of interest stems from changes in wives’ hours; these hours changes tend to be large at the same time that husbands’ incomes experience large changes.

4 Conclusion

This paper has decomposed observed bivariate kurtosis in couples’ income changes; absolute income changes of husbands and wives tend to coincide. There is some evidence of heterogeneity across couples in the covariance parameter governing their of income changes; there is strong evidence that husbands’ and wives’ have correlated parameters governing the variances of their income changes. In the case of permanent income changes, these two forms of heterogeneity explain all observed bivariate

⁶Examining long-term changes in these work-related variables for husbands would be less fruitful, since there is less adjustment in hours and labor force participation for men than for women.

⁷Excess 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 are 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.

kurtosis in couples' income changes.

The bounds on both forms of correlated heterogeneity identified here are useful for models of the household. The impact of intra-household risk-sharing (as proxied by the covariance parameter governing couples' income shocks) on savings, wealth or consumption will be attenuated – biased towards zero – in OLS regressions since couples' covariance parameters are measured with substantial error. For example, Hess (2004) uses couples' covariances to predict divorce as a test of competing theories of marriage. Since instruments for couples' covariances are weak (and of dubious exogeneity), it is more fruitful to exploit the full range of variation in covariances in the data. To correct for the attenuation bias caused by including noisy measures of covariance as right-hand-side variables, we need the fraction of variation in parameter estimates that stems from variation in parameters (as opposed to estimation error). This paper provides an upper bound on the extent of attenuation bias in such regressions.

Furthermore, this paper documents a high correlation between husbands' and wives' income change variances. This positive assortative mating is what would be expected in a model of couple formation in which risk-aversion varies across individuals. To the degree that preferences are uniform but the technologies that produce volatile incomes vary across individuals, negative assortative mating would be predicted.

A Appendix: Estimating Sources of Heterogeneity

A.1 Model

Consider two variables, x_i and y_i , that may not be independent of one another but are mutually independent across observations, i . In the case of couples' income changes studies in this paper, x_i is the one-year change in “excess” log income for a wife

in couple i and y_i is the one-year change in “excess” log income for her husband.⁸ The word “excess” (described in detail in Section 2) implies that any aggregate or predictable changes to income have been removed, so that x_i and y_i are residuals and therefore unconditionally mean zero by construction.⁹

Bivariate kurtosis has been used broadly to refer to the set of possible fourth moments coming from a pair of random variables: $E[x_i^4]$, $E[x_i^3 y_i]$, $E[x_i^2 y_i^2]$, $E[x_i y_i^3]$ and $E[y_i^4]$. Mardia (1970) proposes a summary statistic that combines these. Here, I focus on the symmetric moment, $E[x_i^2 y_i^2]$, because of the information that it encodes about correlated parameter heterogeneity, either the covariance between $\sigma_x^2|i$ and $\sigma_y^2|i$ across observations (denoted $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$) or heterogeneity across observations in $\sigma_{xy}|i$ (denoted $var_i((\sigma_{xy}|i))$).

If x_i and y_i have a conditionally bivariate normal distribution, then

$$E[x_i^2 y_i^2|i] = (\sigma_x^2|i)(\sigma_y^2|i) + 2(\sigma_{xy}|i)^2. \quad (2)$$

I follow Mardia’s convention of using this jointly normal baseline. I refer to the symmetric bivariate analog to excess kurtosis as excess bivariate kurtosis:

$$\kappa_{xy}|i \equiv 3 \left(\frac{E[x_i^2 y_i^2|i]}{(\sigma_x^2|i)(\sigma_y^2|i) + 2(\sigma_{xy}|i)^2} - 1 \right); \quad (3)$$

$$\kappa_{xy} \equiv 3 \left(\frac{E[x_i^2 y_i^2]}{\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2} - 1 \right). \quad (4)$$

$\kappa_{xy}|i$ measures bivariate kurtosis conditioning on observation-specific parameters such as the variances of x_i and y_i for a given i ; naturally, this is unobserved. κ_{xy} measures unconditional bivariate kurtosis and is straightforward to estimate from its constituent parts. Under conditional bivariate normality, $\kappa_{xy}|i = 0$. Note that if $x_i = y_i$, then measures of bivariate kurtosis collapse to the standard univariate definition of kurtosis.

To consider heterogeneity in lower (than fourth) order moments, I make the simplifying assumption that $\kappa_{xy}|i$ does not vary across observations. In this case, it is straightforward to rewrite equation 3 as:

$$E[x_i^2 y_i^2|i] = \left(\frac{\kappa_{xy}|i}{3} + 1 \right) ((\sigma_x^2|i)(\sigma_y^2|i) + 2(\sigma_{xy}|i)^2) \quad (5)$$

Subtracting $\left(\frac{\kappa|i}{3} + 1 \right) (\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2)$ from both sides, taking expectations (where $E[E[x_i^2 y_i^2|i]] =$

⁸The choice of who is x (wives) and who is y (husbands) corresponds to the second sex chromosome (XX for women and XY for men).

⁹While the mean-zero assumption is not central here, it allows the unconditional expectations $E[x_i^2]$ and $E[y_i^2]$ to be relabeled σ_x^2 and σ_y^2 and called variances and allows the unconditional expectation $E[x_i y_i]$ to be relabeled σ_{xy} and called a covariance. Since this paper considers latent heterogeneity, it admits the possibility that the variance-covariance matrix of couples’ income changes may differ *ex-ante* (but unobservably) across observations, i ; $\sigma_x^2|i$, $\sigma_y^2|i$, and $\sigma_{xy}|i$ denote the elements of this matrix.

$E[x_i^2 y_i^2]$ by the law of iterated expectations), dividing by $(\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2)$, and rearranging, equation 5 can be rewritten as:

$$\kappa_{xy} = \kappa_{xy|i} + (\kappa_{xy|i} + 3) \frac{\text{cov}_i((\sigma_x^2|i), (\sigma_y^2|i)) + 2\text{var}_i((\sigma_{xy}|i))}{\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2}. \quad (6)$$

In other words, unconditional bivariate kurtosis (κ_{xy} , which can be estimated from the data) reflects three (unobserved) factors:

1. conditional bivariate kurtosis, $\kappa_{xy|i}$;
2. covarying variances, $\text{cov}_i((\sigma_x^2|i), (\sigma_y^2|i))$; and,
3. heterogeneous covariances, $\text{var}_i((\sigma_{xy}|i))$.

In the first case, large income changes for husbands and wives tend to coincide (conditional on husbands' and wives' income variances and covariances); in the second case, husbands with high-variance income changes tend to have wives with the same; in the third case, some couples' incomes move together while others move in opposite directions. All three imply the tendency of large absolute income changes for husbands and wives to coincide. The i subscript on the variance and covariance operators refer to the cross-section of conditional moments over observations i . For example, $\text{var}_i((\sigma_{xy}|i)) > 0$ indicates that observations differ from one another in their *ex-ante* covariance, $\sigma_{xy|i}$. In the univariate case (setting $x_i = y_i$ so that $\kappa_x|i \equiv E[x_i^4|i] / (\sigma_x^2|i)^2 - 3$ and $\kappa_x \equiv E[x_i^4] / (\sigma_x^2)^2 - 3$), this reduces to:


$$\kappa_x = \kappa_x|i + (\kappa_x|i + 3) \frac{\text{var}_i((\sigma_x^2|i))}{(\sigma_x^2)^2} \quad (7)$$

Covariance heterogeneity and correlated variances appear identically in observed bivariate kurtosis. This is shown in the two panels of Figure 2. The two panels present the same data, eight hypothetical observations (shown as circles, which are in the same locations in each panel) for x_i and y_i . In particular, x_i and y_i both take on values of -1 , 0 , and 1 with probabilities $(1/4, 1/2, 1/4)$ and therefore $E[x_i] = E[y_i] = 0$ and $\sigma_x^2 = \sigma_y^2 = 1/2$. Were x_i and y_i to be independent, $E[x_i^2 y_i^2] = 1/4$. x_i and y_i are not independent (though they are unconditionally uncorrelated, $\sigma_{xy} = 0$) but the marginal distributions of x_i and y_i are unchanged. The key feature of this distribution is its excess bivariate kurtosis, the absence (compared with the distribution under independence) of mass where exactly one variable (x_i or y_i , but not both) is zero. Since non-zero values of x_i and y_i always coincide, the mean of $E[x_i^2 y_i^2] = 1/2$ compared to $1/4$ in the case of independence.

The two panels present different possible explanations for the bivariate kurtosis found in this hypothetical data: correlated variances ($\text{cov}_i((\sigma_x^2|i), (\sigma_y^2|i)) > 0$, right panel) or covariance heterogeneity ($\text{var}_i((\sigma_{xy}|i)) > 0$, left panel).

In the left panel, observations are either in a negative covariance state or a positive covariance state. Covariances, $(\sigma_{xy}|i)$, are either -1 (observations identified with a negative sign and running from top-left to bottom-right) or 1 (observations identified

Figure 2: Stylized Joint Distribution of Couples' Changes in Income



Stylized depiction of the data. See text for details.

with a positive sign and running from bottom-left to top-right) with equal probability. Conditional on the covariance, the distribution is trinomial (values of -1 , 0 , and 1 are possible).

In the right panel, some observations are in a high variance state while others are in a low variance state. Variances, $(\sigma_x^2|i)$ and $(\sigma_y^2|i)$, are either both 0 (marked with a “low” and clustered at zero) or both 1 (marked with a “high” and found at the corners) for both x_i and y_i and the variances for x_i and y_i are perfectly correlated. Conditional on the variances, the distribution is binomial (values of -1 and 1 are possible in the high variance state while only values of 0 are possible in the low variance state).

If we observe the unconditional distribution depicted in these panels, where large absolute values of x_i and y_i tend to coincide, this could reflect either correlated variances or covariance heterogeneity. A third extreme possibility is that there is no *ex-ante* heterogeneity; unconditional bivariate kurtosis reflects conditional bivariate kurtosis and not correlated heterogeneity. In other words, all observations are drawn from the same distribution which has the feature that large absolute changes of x_i and y_i happen to coincide. Of course, any combination of conditional bivariate kurtosis, correlated variances, and covariance heterogeneity will be consistent with the unconditional joint distribution described here.

A.2 Testing for Correlated Heterogeneity

Here, I present distributions for a test statistic for unconditional bivariate kurtosis. The aim is to test the null that there is no excess unconditional bivariate kurtosis, the joint normal baseline.

Under the null hypothesis of no bivariate kurtosis when $\sigma_{xy} = 0$ (a strong but testable assumption appropriate for the application to follow), for a randomly chosen i from the population, $x_i^2 y_i^2$ will have mean $\sigma_x^2 \sigma_y^2$ and variance $(\text{var}_i((\sigma_x^2|i)) + \sigma_x^4) \kappa_x (\text{var}_i((\sigma_y^2|i)) + \sigma_y^4)$

$\sigma_x^4 \sigma_y^4$. This is merely the product of $E[x_i^4]$ and $E[y_i^4]$ less the square of the mean. Note that under the null hypothesis and assuming moments are finite, $(\text{var}_i((\sigma_x^2|i)) + \sigma_x^4) \kappa_x$ can be estimated with $\frac{1}{N} \sum_i x_i^4$ and $(\text{var}_i((\sigma_y^2|i)) + \sigma_y^4) \kappa_y$ can be estimated with $\frac{1}{N} \sum_i y_i^4$. Since observations are assumed to be *iid*, under the null hypothesis with $\sigma_{xy} = 0$ the sample variance of $x_i y_i$, $\frac{1}{N} \sum_i x_i^2 y_i^2 - (\frac{1}{N} \sum_i x_i y_i)^2$ will have mean $\bar{\sigma}_x^2 \bar{\sigma}_y^2$ and variance $\frac{1}{N} ((\text{var}_i(\sigma_x^2|i) + \sigma_x^4) \kappa_x (\text{var}_i(\sigma_y^2|i) + \sigma_y^4) \kappa_y - \sigma_x^4 \sigma_y^4)$. Since we have the distribution of the sample variance it is straightforward to test that null.

Formally, the sample moment $\frac{1}{N} \sum_i x_i^2 y_i^2$ just allows for a test of the independence of shocks, $E[f(x_i) f(y_i)] = E[f(x_i)] E[f(y_i)]$. Independence requires that this be true for all $f()$ and $g()$ and here we look only at second moments, $f(x_i) = x_i^2$ and $g(y_i) = y_i^2$. The novelty here is that equation 6 decomposes this particular rejection of independence into conditional bivariate kurtosis and two types of latent correlated heterogeneity. In the example that follows, such correlated heterogeneity is of economic interest. Do all couples' incomes jointly evolve in the same way?

A.3 “Wife-Swap Bootstrap”

So far, $\{x_i, y_i\}$ pairs have been assumed to be independent of other pairs. For a cross-section of randomly chosen individuals who face idiosyncratic shocks, this assumption may be relatively innocuous. When data comes from a panel, this is seldom true. I add time subscripts (e.g., x_{it} , $(\sigma_y^2|i, t)$) to accommodate autocorrelation. In this case, the sample variance, $\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it}^2 y_{it}^2 - (\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it} y_{it})^2$ will be drawn from a distribution with same mean as in the *i.i.d.* case, $\sigma_x^2 \sigma_y^2$, but not the same variance:

$$\begin{aligned} & \frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it}^2 y_{it}^2 - \left(\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it} y_{it} \right)^2 \\ & \sim \left(\sigma_x^2 \sigma_y^2, \frac{1}{NT} \left((\text{var}_i((\sigma_x^2|i, t)) + \sigma_x^4) \kappa_x (\text{var}_i((\sigma_y^2|i, t)) + \sigma_y^4) \kappa_y - \sigma_x^4 \sigma_y^4 \right) \right. \\ & \quad \left. + \frac{1}{T} \sum_i \sum_t \sum_{s \neq t} \text{cov}(x_{is}^2 y_{is}^2, x_{it}^2 y_{it}^2) \right) \end{aligned} \quad (8)$$

The first part of the variance (same as in the *i.i.d.* case) is trivial to estimate from sample data as

$$\frac{1}{NT} \left(\frac{1}{N^2 T^2} \sum_i \sum_t x_{it}^4 \sum_i \sum_t y_{it}^4 - \frac{1}{N^4 T^4} (\sum_i \sum_t x_{it}^2)^2 (\sum_i \sum_t y_{it}^2)^2 \right);$$

covariance terms (stemming from autocorrelation) are more difficult to estimate. The main challenge in a non-rectangular panel is that attrition may be related to the autocorrelation. Without attrition, $\text{cov}(x_{is}^2 y_{is}^2, x_{it}^2 y_{it}^2)$ can be estimated from data under the null as $\frac{1}{N} \sum_i x_{is}^2 x_{it}^2 \frac{1}{N} \sum_i y_{is}^2 y_{it}^2$. An alternative way to obtain the same variance can be obtained by noting that under the null, $\text{var}(\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it}^2 y_{it}^2) = \text{var}(\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it}^2 y_{jt}^2)$ for a randomly chosen $j \neq i$. (The non-rectangularity problem can be overcome if j is chosen so that i and j have the same number of observations.) As a result, it is straightforward to obtain the variance of the estimator by repeatedly sampling $\frac{1}{N} \frac{1}{T} \sum_i \sum_t x_{it}^2 y_{jt}^2$ for different choices of j and taking the variance of these. When x and y refer to the incomes of husbands and wives, this involves randomly pairing all

husbands and wives from the data, and calculating the estimator for this synthetic pair. Doing this repeatedly builds up a reference distribution under the null. I use the tongue-in-cheek name *wife-swap bootstrap* to refer to this procedure.

A.4 Bounding Correlated Heterogeneity

After rejecting the null of no excess bivariate kurtosis (because $\kappa_{xy} > 0$), we know that $cov_i((\sigma_x^2|i), (\sigma_y^2|i)) > 0$, $var_i((\sigma_{xy}|i)) > 0$, $\kappa_{xy|i} > 0$, or some combination of these. $var_i((\sigma_{xy}|i)) > 0$ indicates that the covariance differs across observations. $cov_i((\sigma_x^2|i), (\sigma_y^2|i)) > 0$ indicates that observations with high-variance x also tend to have high-variance y . While identified from the same moment in the data $\frac{1}{N}\sum_i x_i^2 y_i^2$, they reflect completely different phenomenon. Consider the application to couples, and x refers to the change in husbands' incomes and y refers to the change in wives' incomes. $var_i((\sigma_{xy}|i)) > 0$ could be interpreted as saying that the diversification benefits of marriage (proxied by $(\sigma_{xy}|i)$) vary across couples. By contrast, $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$ identifies assortative mating in risk, which could be a test of models of optimal partner selection.

Without additional information from covariates, conditional kurtosis, correlated variances and covariance heterogeneity are observationally equivalent. To separate them, we must use covariates Z_i to obtain observation-specific estimates of $(\sigma_x^2|i)$, $(\sigma_y^2|i)$, and $(\sigma_{xy}|i)$, and then identify the heterogeneity in $(\sigma_{xy}|i)$ or correlated heterogeneity in $(\sigma_x^2|i)$ and $(\sigma_y^2|i)$ that can be traced out by variation in Z_i . Consider the following set of regressions (where each element of x_i , μ_{xi} , etc, refers to a vector):

$$\begin{pmatrix} x_i^2 \\ y_i^2 \\ x_i y_i \end{pmatrix} = Z_i \times \begin{pmatrix} \beta_x \\ \beta_y \\ \beta_{xy} \end{pmatrix} + \begin{pmatrix} \mu_{xi} \\ \mu_{yi} \\ \mu_{xyi} \end{pmatrix}.$$

Variation in $(\sigma_{xy}|i)$ traced out by Z_i places a lower bound on $var_i((\sigma_{xy}|i)) \geq \beta_{xy} Z' Z \beta_{xy}$; correlated variation in $(\sigma_x^2|i)$ and $(\sigma_y^2|i)$ traced out by Z_i ($\beta_x Z' Z \beta_y$) provides one source of $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$.¹⁰ Since additional correlated variation in variances could be of either sign, the total magnitude of correlated variation in variances is not bounded by $\beta_x Z' Z \beta_y$. Having said that, the panel data approach outlined in Section A.5 provides a setting where this is likely to be a lower bound.

From equation 6, lower bounds on $var_i((\sigma_{xy}|i))$ and $cov_i((\sigma_x^2|i), (\sigma_y^2|i))$ imply an upper bound on the importance of conditional bivariate kurtosis in explaining unconditional bivariate kurtosis. These are the upper-bounds for the importance of conditional bivariate kurtosis under the assumption that $\kappa_{xy|i}$ (defined in equation 3) are the same across individuals:

$$\kappa_{xy|i} \leq \frac{\kappa_{xy} (\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2) - 3(\beta_x Z' Z \beta_y + 2\beta_{xy} Z' Z \beta_{xy})}{\sigma_x^2 \sigma_y^2 + 2\sigma_{xy}^2 + \beta_x Z' Z \beta_y + 2\beta_{xy} Z' Z \beta_{xy}} \quad (9)$$

¹⁰This bears some similarity to Mardia and Marshall (1984) who provide maximum likelihood estimates of covariance heterogeneity (traced out by parametric variation) in a conditionally normal setting.

Note that all of the objects on the right-hand side of these inequalities can be estimated.

A.5 Panel Data

While panel data complicates estimation of unconditional bivariate kurtosis (see Section A.3), it also provides additional information useful in decomposing it. Couples i may differ from one another in their covariance parameter, $(\sigma_{xy}|i)$; and husbands with high variance parameters $(\sigma_y^2|i)$ may have wives with high variance parameters $(\sigma_x^2|i)$. With multiple observations from each couple, couple-specific estimates become possible. I assume that there exist s and t sufficiently far apart (for example, a fixed distance k) that common shocks from the two periods are uncorrelated. In the example that follows, I use $s = t - 5$. These assumptions are strong, but are readily testable in the case of couples, whose non-overlapping income changes are nearly uncorrelated and where any changes in the distribution of parameters is slow. For example, Abowd and Card (1989) show that innovations to income are not autocorrelated at lags greater than two years.

Most obviously, note that $|cov_i((\sigma_{xy}|is), (\sigma_{xy}|it))| \leq \sqrt{var_i(\sigma_{xy}|is) var_i(\sigma_{xy}|it)}$. If the distribution of $\sigma_{xy}|i$ is stable, then this implies $|cov_i((\sigma_{xy}|is), (\sigma_{xy}|it))| \leq var_i(\sigma_{xy}|it)$ (the last equality by the stability assumption). $cov_i((\sigma_{xy}|is), (\sigma_{xy}|it))$ can be readily estimated from the data as $\frac{1}{NT} \sum_t \sum_i (x_{is} y_{is} x_{it} y_{it} - \hat{\sigma}_{xy}^2)$, and this provides a lower bound for $var_i(\sigma_{xy}|it)$.

While it is not strictly required by the assumptions above, all but the most pathological distributions will exhibit

$$\begin{aligned} & \frac{1}{2} |cov_i((\sigma_x^2|is), (\sigma_y^2|it)) + cov_i((\sigma_x^2|it), (\sigma_y^2|is))| \\ & < \frac{1}{2} (cov_i((\sigma_x^2|is), (\sigma_y^2|is)) + cov_i((\sigma_x^2|it), (\sigma_y^2|it))) \\ & = cov_i((\sigma_x^2|it), (\sigma_y^2|it)) \end{aligned}$$

where the last equality follows from stability. Contemporaneous shocks should be more highly correlated than lead or lagged shocks with a large enough time-gap. This need not be true when one variable predicts subsequent values for other, but when $((\sigma_x^2|is), (\sigma_y^2|it))$ and $cov_i((\sigma_x^2|it), (\sigma_y^2|is))$ are positive and similar in value, contemporaneous shocks are more likely to have similar magnitudes. $((\sigma_x^2|is), (\sigma_y^2|it))$ and $cov_i((\sigma_x^2|it), (\sigma_y^2|is))$ can be readily estimated from the data with $\frac{1}{NT} \sum_t \sum_i (x_{is}^2 y_{it}^2 - \hat{\sigma}_x^2 \hat{\sigma}_y^2 - 2\hat{\sigma}_{xy}^2)$ and $\frac{1}{NT} \sum_t \sum_i (x_{it}^2 y_{is}^2 - \hat{\sigma}_x^2 \hat{\sigma}_y^2 - 2\hat{\sigma}_{xy}^2)$, respectively. This estimates a lower bound on $cov_i((\sigma_x^2|it), (\sigma_y^2|iy))$ as

$$\frac{1}{NT} \sum_t \sum_i \left(\frac{1}{2} x_{is}^2 y_{it}^2 + \frac{1}{2} x_{it}^2 y_{is}^2 - \hat{\sigma}_x^2 \hat{\sigma}_y^2 - 2\hat{\sigma}_{xy}^2 \right).$$

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.
- BANIA, N., AND L. LEETE (2007): “Income Volatility and Food Insufficiency in U.S. Low-Income Households, 1992-2003,” Institute for Research on Poverty Discussion Paper No. 1325-07.
- BATCHELDER, L. L. (2003): “Can Families Smooth Variable Earnings?,” *Harvard Journal on Legislation*, 40(2), 395–452.
- BECKER, G. S. (1973): “A Theory of Marriage: Part I,” *Journal of Political Economy*, 81(4), 813–846.
- BOLLINGER, C., AND J. P. ZILIAK (2007): “Welfare Reform and the Level, Composition, and Volatility of Income,” Working Paper.
- BURTLESS, G. (1990): “Comment on M. Hurd’s ‘The Joint Retirement Decisions of Husbands and Wives’,” in *Issues in the Economics of Aging*, ed. by D. A. Wise. University of Chicago Press, Chicago.
- CAMERON, S., AND J. TRACY (1998): “Earnings Variability in the United States: An Examination Using Matched-CPS Data,” Working Paper.
- CARROLL, C. D., AND A. A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), 41–71.
- CASANOVA, M. (2010): “Happy Together: A Structural Model of Couples’ Joint Retirement Choices,” Working Paper.
- CHAMI, R., AND G. D. HESS (2005): “For Better or For Worse? State-Level Marital Formation and Risk Sharing,” *Review of Economics of the Household*, 3(4), 1569–5239.
- CHIAPPORI, P.-A., AND P. J. RENY (2006): “Matching to Share Risk,” Working Paper.
- COMIN, DIEGO, E. L. G., AND B. RABIN (2006): “Turbulent Firms, Turbulent Wages?,” NBER Working Paper No. 12032.
- CULLEN, J. B., AND J. GRUBER (2000): “Does Unemployment Insurance Crowd Out Spousal Labor Supply?,” *Journal of Labor Economics*, 18(3), 546–72.
- DAHL, M., T. DELEIRE, AND J. SCHWABISH (2007): “Trends in Earnings Variability Over the Past 20 Years,” Working Paper: Congressional Budget Office.

- DALY, M. C., AND G. J. DUNCAN (1997): "Earnings Mobility and Instability, 1969-1995," Working Paper.
- DYNAN, K. E., D. W. ELMENDORF, AND D. E. SICHEL (2007): "The Evolution of Earnings and Income Volatility," Working Paper, Federal Reserve Board.
- DYNARSKI, S., AND J. GRUBER (1997): "Can Families Smooth Variable Earnings?," *Brookings Papers on Economic Activity*, pp. 229–303.
- GEWEKE, J. F., AND M. KEANE (2000): "An Empirical Analysis of Income Dynamics among Men in the PSID: 1968-1989," *Journal of Econometrics*, 96, 293–356.
- GOTTSCHALK, P., AND R. MOFFITT (1994): "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, pp. 217–254.
- (2002): "Trends in the Transitory Variance of Earnings in the United States," *Economic Journal*, 112, C68–73.
- (2006): "Trends in Earnings Volatility in the U.S.: 1970-2002," Working Paper.
- GUSTMAN, A., AND T. STEINMEIER (2000): "Retirement in Dual-Career Families: A Structural Model," *Journal of Labor Economics*, 18(3), 503–545.
- HACKER, J. (2006): *The Great Risk Shift: The Assault on American Jobs, Families, Health Care, and Retirement; And How You Can Fight Back*. Oxford University Press, Oxford.
- HAIDER, S. J. (2001): "Earnings Instability and Earnings Inequality in the United States, 1967-1991," *Journal of Labor Economics*, 19(4), 799–836.
- HALL, R. E., AND F. S. MISHKIN (1982): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," *Econometrica*, 50, 461–481.
- HERTZ, T. (2006): "Understanding Mobility in America," Center for American Progress Discussion Paper.
- HESS, G. D. (2004): "Marriage and Consumption Insurance: What's Love Got to Do with It?," *Journal of Political Economy*, 112(2), 290–318.
- HURD, M. (1990): "The Joint Retirement Decisions of Husbands and Wives," in *Issues in the Economics of Aging*, ed. by D. A. Wise. University of Chicago Press, Chicago.
- HYSLOP, D. R. (2001): "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review*, 91(4), 755–777.

- JENSEN, S. T., AND S. H. SHORE (2011): “Semi-Parametric Bayesian Modelling of Income Volatility Heterogeneity,” *Journal of the American Statistical Association*, 106(496), 1280–1290.
- (2012): “Changes in the Distribution of Income Volatility,” Working Paper, Johns Hopkins University.
- LUNDBERG, S. (1985): “The Added Worker Effect,” *Journal of Labor Economics*, 3(1), 11–37.
- MAESTAS, N. (2001): “Labor, Love, and Leisure: Complementarity and the Timing of Retirement by Working Couples,” mimeo, University of California at Berkeley.
- MARDIA, K. (1970): “Measures of multivariate skewness and kurtosis with applications,” *Biometrika*, 36, 519–530.
- (1974): “Applications of some measures of multivariate skewness and kurtosis in testing normality and robustness studies,” *Sankhya, Ser B*, 36(2), 115–128.
- (1980): “Tests of univariate and multivariate normality,” in *Handbook of Statistics, Vol 1*, ed. by E. Kotz, pp. 279–320. John Wiley and Sons, New York.
- MARDIA, K., AND R. MARSHALL (1984): “Maximum Likelihood Estimation of Models for Residual Covariance in Spatial Regression,” *Biometrika*, 71(1), 135–146.
- MEGHIR, C., AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1), 1–32.
- MICHAUD, P. (2003): “Joint Labor Supply Dynamics of Older Couples,” IZA Discussion Paper 832.
- MOFFITT, R. A., AND P. GOTTSCHALK (1995): “Trends in the Covariance Structure of Earnings in the U.S.,” Working Paper.
- NORDBLOM, K. (2004): “Cohabitation and Marriage in a Risky World,” *Review of Economics of the Household*, 2(3), 325–340.
- SHIN, D., AND G. SOLON (2008): “Trends in Mens Earnings Volatility: What Does the Panel Study of Income Dynamics Show?,” Working Paper.
- SHORE, S. H. (2010): “For Better, For Worse: Intrahousehold Risk-Sharing Over the Business Cycle,” *Review of Economics and Statistics*, 92(3).
- VERNON, V. (2010): “Marriage for Love, for Money..and for Time?,” *Review of the Economics of the Household*, 8(4), 433–457.
- WINSHIP, S. (2007): “Income Volatility, The Great Risk Shift, and the Democratic Agenda,” Working Paper.