Queen's Economics Department Working Paper No. 1195

Consumption and Real Exchange Rates in Professional Forecasts

Michael B. Devereux
University of British Columbia, CEPR, and NBER

Gregor W. Smith
Queen's University

James Yetman
Bank for International Settlements

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

1-2009

Consumption and Real Exchange Rates in Professional Forecasts

Michael B. Devereux, Gregor W. Smith, and James Yetman†

January 2009

Abstract

Standard models of international risk sharing with complete asset markets predict a positive association between relative consumption growth and real exchange-rate depreciations across countries. The striking lack of evidence for this link — the consumption/real-exchange-rate anomaly or Backus-Smith puzzle — has prompted research on risk-sharing indicators with incomplete asset markets. That research generally implies that the association holds in forecasts, rather than realizations. Using professional forecasts for 28 countries for 1990-2008 we find no such association, thus deepening the puzzle. Independent evidence on the weak link between forecasts for consumption and real interest rates suggests that the presence of 'hand-to-mouth' consumers may help to explain the evidence.

JEL classification: F41, F47, F37

Keywords: international risk sharing, Backus-Smith puzzle

†Devereux: Department of Economics, University of British Columbia, CEPR, NBER, `devm@ interchange.ubc.ca`; Smith: Department of Economics, Queen's University, `smithgw@ econ.queensu.ca`; Yetman: Bank for International Settlements, `james.yetman@bis.org`. We thank the Social Sciences Research Council of Canada and the Bank of Canada research fellowship programme for support of this research. The opinions are the authors' alone and are not those of the Bank of Canada or the Bank for International Settlements. We thank Rasool Zandvakil for skilled research assistance and seminar participants at the University of Washington, McGill University, the BIS, the Bank of Korea, and Hong Kong University of Science and Technology for helpful comments.

1. Introduction

A wide range of international macroeconomic models imply a positive relationship between relative consumption growth and real exchange rate depreciations as a consequence of integrated asset markets and risk sharing. As is well known, though, most of the evidence is that there is no such link in historical data, or that the relationship may even be negative. This rejection of an implication of these models is sometimes called the consumption/real-exchange-rate anomaly or Backus-Smith puzzle.

Recently, a number of researchers have proposed theoretical models in which asset markets are incomplete in specific ways, and shown that they can produce low or even negative correlations over time between relative consumption growth and the real exchange rate. This property is a form of *indirect* inference in favor of these models. But these same models typically also imply that the original relationship holds in conditional expectations, rather than state-by-state, due to conditional risk sharing. This alternative property can be tested using instrumental-variables methods, but such tests are hampered by the challenges of forecasting consumption growth and real exchange rate depreciations, which often have very little statistical predictability.

We provide a new, *direct* test of an implication of the incomplete-markets models. We investigate the link between relative consumption growth and real exchange rate changes in panels of forecasts made by professional forecasters. Outsourcing the forecasting problem in this way does not require us to calibrate a specific economic model. We provide graphical evidence using these forecasts for 28 countries over 1990-2008. We also specialize the evidence by pairs of countries, by time, and by exchange-rate regime. The overall conclusion is that there is still very little evidence of a positive relationship. This deepens the puzzle.

We then provide an interpretation of the failure of conditional risk sharing in the forecast data. For conditional risk sharing to hold, agents must have access to some asset markets, though not necessarily a full set of Arrow-Debreu securities. A substantial body of evidence suggests that households either do not use asset markets at all, or else fail to actively adjust their asset holdings optimally over time. In that case, even the conditional link between relative consumption growth and real depreciations may not hold. Instead the predicted relationship may be of either sign, depending on the sources of shocks and preference parameters. We develop a model of hand-to-mouth consumers which illustrates these points. Interestingly the model also illustrates a pitfall in a conventional test of the consumption/real-exchange-rate relationship. Researchers using that test may fail to detect conditional risk sharing, even if it holds.

Section 2 outlines notation and the risk-sharing conditions. Section 3 describes the tests of necessary conditions used in previous work as well as the findings of those tests. Section 4 outlines the sources and nature of the forecast data we use. Section 5 presents the evidence, in the form of scatter plots. Section 6 provides a simple interpretation based on the presence of rule of thumb consumers, and section 7 concludes.

2. Risk-Sharing Indicators and Incomplete Markets

We start with some basic theory. Take an N-country, world economy where households within a country i, $i = 1, 2, ..., N$, have identical preferences and unrestricted access to intra-national markets for risk sharing. Hence there is a representative household for each country. Time is discrete, beginning at $t = 0$. In each time period, the aggregate state is labeled z_t , where z_t comes from a finite set of possible states of the world. At time t, the state history is labelled $z^t = \{z_0, z_1, ..., z_t\}$, and $\pi(z^t)$ is the probability of history z^t. Households in country i have preferences given by

$$E_0 \sum_{t=0}^{\infty} \sum_{z^t} \beta^t \pi(z^t) u(c^i(z^t)) \tag{1}$$

where $c^i(z^t)$ represents country i's consumption composite at history z^t. We are assuming that preferences over consumption composites are identical across countries, and that discount factors β are also identical, but we do not necessarily assume that the composition of the consumption aggregates are the same in each country. There may be a country-specific, non-traded good for example.

Associated with each country's consumption aggregate is a consumer price index $p^i(z^t)$, defined in terms of currency i. Nominal exchange rates are defined with respect to a numeraire currency. Letting currency 1 be the numeraire currency, we define $s^i(z^t)$ as the currency–1 price of currency i, with $s^1(z^t) = 1$. Hence the real exchange rate between country 1 and country i is defined in the usual way as $q^i(z^t) = s^i(z^t)p^i(z^t)/p^1(z^t)$. We make no special assumptions about real exchange rate determination; $q^i(z^t)$ may reflect the presence of trade frictions, non-traded goods, sticky prices, or other goods market imperfections. Table 1 summarizes the notation used so far and supplemented below.

Asset markets are frictionless in the sense that agents in all countries face the same set of asset prices and returns, when measured in any numeraire. This assumption does not imply that assets markets are complete. But it means that for the set of assets that are traded across countries, all agents face the same prices and payoffs.

Table 1: Notation

t	time	i, j	countries
u	period utility	β	discount factor
γ	EIS	z^t	state history
c	consumption	p	price
π	state probability	m	asset price
s	nominal exchange rate	q	real exchange rate
y^{ij}	$\dot{c}^i_t - \dot{c}^j_t$	x^{ij}	$\dot{q}^{ij}_t \equiv \dot{s}^{ij}_t + \dot{p}^j_t - \dot{p}^i_t$

We begin with risk-sharing under complete markets. Assume that there exists a set of assets which have payoff in currency 1 in each possible history. Assets paying off in other currencies are redundant because markets are complete. Define $m(z^{t+1}|z^t)$ as the price of an asset which pays off one unit of currency in history z^{t+1}, conditional on history z^t. Then this price is determined by the condition:

$$m(z^{t+1}|z^t)\frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \beta\frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})}\pi(z^{t+1}|z^t) \tag{2}$$

where $\pi(z^{t+1}|z^t) = \pi(z^{t+1})/\pi(z^t)$. Since this condition holds for each country i, and all countries face identical asset prices, we must have:

$$\frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})}\Big/\frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \frac{u'(c^j(z^{t+1}))}{s^j(z^{t+1})p^j(z^{t+1})}\Big/\frac{u'(c^j(z^t))}{s^j(z^t)p^j(z^t)}, \tag{3}$$

$\forall i, j = 1, ..N$. This is the complete markets risk-sharing condition of Backus and Smith (1993). We can rewrite this condition as:

$$\frac{u'(c^i(z^{t+1}))}{u'(c^i(z^t))}\Big/\frac{u'(c^j(z^{t+1}))}{u'(c^j(z^t))} = \frac{q^{ij}(z^t)}{q^{ij}(z^{t+1})} \tag{4}$$

$\forall i, j = 1, ..N$, which says that the ex-post ratio of intertemporal marginal rates of substitution should be equal to the ex-post growth rate of the real exchange rate. Using dots to denote growth rates, and taking a linear approximation of the risk-sharing condition (4) around a non-stochastic steady state, we have:

$$\dot{c}^i_{t+1} - \dot{c}^j_{t+1} = \gamma \dot{q}^{ij}_{t+1} \tag{5}$$

$\forall i, j = 1, ..N$, where γ is the common, intertemporal elasticity of substitution. (No approximation is involved when utility is isoelastic.) Thus, full risk sharing implies that the difference between ex-post growth rates in consumption across countries i and j should be a positive, linear function of the ex-post growth rate of the real exchange rate q^{ij}.

Condition (5) presents a simple and intuitive condition for testing risk sharing across countries. Cross-country deviations in consumption per capita should occur only to the extent that there are offsetting real exchange rate changes. A country's relative consumption should rise when the relative price of consumption falls. Note that the source of real exchange rate variation has no bearing whatever on the prediction of condition (5). Whether real exchange rates vary due to changes in the relative price of non-traded goods, deviations from the law of one price in traded goods, or compositional effects of terms of trade changes, does not alter the predicted relationship between relative consumption and the real exchange rate.

This risk-sharing condition must be amended if some assets markets are missing. Consider the case where only non-contingent, currency-1-denominated, nominal bonds are available. This case has been studied by Kollmann (1995) and Corsetti, Dedola, and Leduc

(2008) for example. The price of such a bond is $m(z^t) = \sum_{z^{t+1}} m(z^{t+1}|z^t)$. This bond price is determined by the condition:

$$m(z^t)\frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \beta \sum_{z^{t+1}} \pi(z^{t+1}|z^t)\frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})}. \tag{6}$$

Again, since $m(z^t)$ is common across countries, this pricing implies:

$$\begin{aligned} &\sum_{z^{t+1}} \pi(z^{t+1}|z^t)\frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})}\Big/\frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \\ &\qquad\qquad \sum_{z^{t+1}} \pi(z^{t+1}|z^t)\frac{u'(c^j(z^{t+1}))}{s^j(z^{t+1})p^j(z^{t+1})}\Big/\frac{u'(c^j(z^t))}{s^j(z^t)p^j(z^t)}, \end{aligned} \tag{7}$$

$\forall i, j = 1, ..N$. Taking a linear approximation leads to the risk-sharing condition for this incomplete-markets economy:

$$\mathrm{E}\big[(\dot{c}^i_{t+1} - \dot{c}^j_{t+1})|z^t\big] = \mathrm{E}\big[\gamma \dot{q}^{ij}_{t+1}|z^t\big]. \tag{8}$$

Under this asset market arrangement, the time-t conditional expectation of the differences in consumption growth rates should be proportional to the conditional expected growth rate in the real exchange rate.

More generally, we may define any restriction on asset payoff contingencies by constructing a combination of possible histories ϕ^t, so that an asset pays off a unit of currency 1 for all histories $z^{t+1} \in \phi^{t+1}$. Obstfeld (1994) discusses this general, intermediate case. The approximate version of the empirical implication he draws is condition (8) but with the expectation conditional on the information set $\{\phi^{t+1}, z^t\}$.

The models we test in this paper have incomplete international risk sharing, but complete risk sharing across households within a country, so that there is a representative agent in each economy. Kocherlakota and Pistaferri (2007) consider the opposite case, where there is limited risk sharing within each economy but complete insurance against economy-specific shocks. They derive predictions, under several contracting schemes, for the properties of real exchange rates conditional on the cross-sectional distribution of consumption within an economy, and test those predictions using disaggregated data. The forecasts for aggregates we study cannot be used to test the predictions of their models, which require data disaggregated by households.

3. Tests and Previous Evidence

The complete-markets theory predicts a positive relationship between the predicted consumption growth differential and the predicted real depreciation:

$$\dot{c}_t^i - \dot{c}_t^j = \gamma \dot{q}_t^{ij}. \tag{9}$$

We define the composite variables $y_t^{ij} \equiv \dot{c}_t^i - \dot{c}_t^j$ and $x_t^{ij} \equiv \dot{q}_t^{ij}$, so the relationship is:

$$y_t^{ij} = \gamma x_t^{ij}. \tag{10}$$

This relabelling keeps notation simple and serves as a reminder of the axis labels on diagrams. The scatter plot of the two composite variables is an upward-sloping line in $x-y$ space, with slope γ. The strongest test of the complete-markets, risk-sharing condition is to inspect the scatter plot of y versus x, for example for a given country or year or for all data pooled. We call this the state-by-state test.

Backus and Smith (1993) studied two predictions of the complete-markets relationship. The first of these was the monotonicity property. They calculated the mean, standard deviation, and autocorrelation of the two sides of the equation. They then noted that the two sides should be positively related in a cross-section of pairs of countries. For example, a country-pair with a relatively volatile relative consumption growth also should have a relatively volatile real exchange rate growth. (The use of growth rates ensured stationarity.) The slope should be γ in means and standard deviations, and 1 in autocorrelation coefficients.

Studying the monotonicity property in means involves averaging over time like this:

$$\overline{y}^{ij} = \frac{1}{T}\sum_{t=0}^{T} y_t^{ij} \tag{11a}$$

and

$$\overline{x}^{ij} = \frac{1}{T}\sum_{t=0}^{T} x_t^{ij} \tag{11b}$$

and then inspecting the cross-country cross-plot for this relationship:

$$\overline{y}^{ij} = \gamma \overline{x}^{ij}. \tag{12}$$

The second prediction was for time-series correlations. The two growth rates should be perfectly correlated over time for any pair of countries. The sample covariance is:

$$\frac{1}{T}\sum_{t=0}^{T}(y_t^{ij} - \overline{y}^{ij})(x_t^{ij} - \overline{x}^{ij}) = \frac{1}{T}(\sum_{t=0}^{T} y_t^{ij} x_t^{ij}) - \overline{y}^{ij}\overline{x}^{ij}. \tag{13}$$

When $y = \gamma x$, $\mathrm{cov}_t(y, x) = \gamma \mathrm{var}_t(x)$, so the covariance is positive, and $\mathrm{corr}_t(y, x) = 1$.

Two notable features of these tests are (a) they are weaker than the state-by-state test and (b) they may conflict. Figure 1 illustrates these features using some arbitrary numbers. Black, grey, and white refer to three countries, each relative to a common base country. The circles represent annual data points for the growth rates x and y. The red circles, one for each country, give the country averages.

In the top panel of figure 1 the data points lie along a downward-sloping line for each country, failing the correlation test. But the means lie on an upward-sloping line, passing the cross-country monotonicity test. In the bottom panel, the data pass the correlation test in each country but fail the monotonicity test. Both data sets would fail the stronger, state-by-state test. Removing the red circles and not distinguishing between shades of grey, the points lie in a cloud, not along an upward-sloping line.

But in fact *both* these weaker tests reject. Backus and Smith used quarterly data for 8 OECD countries for 1971-1990. Using the monotonicity test in means (as well as other moments) they found clouds of points, rejecting the risk-sharing condition. Obstfeld (2007) examines more recent data for a wider set of countries. He too finds a negative association between average consumption growth differentials and real exchange rate changes.

Using the correlation test, Backus and Smith found coefficients near zero; the values ranged over [-0.08,0.17], with an average value of 0.045. Kollmann (1995) ran regressions like condition (5) for country pairs and found they had very low R^2, which provides similarly negative information on the correlations. (Kollmann also found that consumption and real exchange rates were not cointegrated in levels, thus rejecting another necessary condition.) This is consistent with the finding of Chari, Kehoe, and McGrattan (2002), who point out that, in time series observations, the correlation between relative consumption and real exchange rates is negative for most OECD economies. They refer to this discrepancy between theory and data as the 'consumption real exchange rate anomaly'.

So the complete-markets condition (5) is rejected. Theoretical approaches to explaining the evidence have relaxed the assumption of complete markets. Corsetti, Dedola, and Leduc (2008), Benigno and Thoenissen (2008), Selaive and Tuesta (2003), and Opazo (2006) all construct theoretical models in which the only asset that is traded across countries is a non-contingent bond. This gives a risk sharing condition of the form (8). This condition is then set in an environment in which a shock that generates a real exchange rate depreciation also leads to a rise in relative consumption. It then is possible to produce a negative correlation between consumption and real exchange rate changes, like the correlation found in the time series data. In Benigno and Thoenissen's model, for example, an unanticipated positive shock to the output of traded goods generates a wealth effect which raises the demand for non-traded goods. This leads to a rise in relative consumption combined with a real exchange rate depreciation. Similar wealth effects feature in the studies by Corsetti, Dedola, and Leduc (2008), Selaive and Tuesta (2003), and Opazo (2006).

While these papers develop general-equilibrium models in which condition (5) fails to hold, due to incomplete markets, they still imply that condition (8) holds: expected consumption growth should be positively correlated with expected real exchange rate changes. Since in almost any theoretical setting, consumption and real exchange rates are determined simultaneously, condition (8) makes no clear prediction regarding the correlation between consumption growth rate differences and real exchange rates. Rather, the condition implies that, adjusted for real depreciations, consumption growth differentials should be uncorrelated with any variables in the date-t information set.

If actual outcomes, state-by-state, lie in a cloud then can expected outcomes ever lie on an upward-sloping line, if expectations are rational? The answer is that they can do so for two reasons. First, forecast errors may be large. Adding volatile, white-noise errors to a regression may make it impossible to detect a positive slope in outcomes. Second, consumption forecast errors may be negatively correlated with exchange rate forecast errors. Again, then, the slopes of the scatter plots would differ between actual outcomes and expectations.

We next assess the appropriateness of testing the incomplete-markets, risk-sharing condition (8) using either the monotonicity-in-means test or the covariance test. We do not know z^t. But if we have instruments u_t and are confident that $u_t \in z^t$ then we can test and estimate with the sample versions of:

$$\mathrm{E}(y_{t+1}^{ij} - \gamma x_{t+1}^{ij} | u_t) = 0. \tag{14}$$

This moment condition can be used for estimation either country-pair by country-pair or for a pooled panel of data with a common γ. With more than one instrument a test of over-identifying restrictions is possible.

Obstfeld (1994), Kollman (1995) and Head, Mattina and Smith (2004) test various versions of this condition. In general these studies find only weak evidence supporting the incomplete markets condition. Obstfeld tested condition (8) (or its extension to a wider information set ϕ^{t+1}) by regressing individual consumption on world consumption, with the addition of world income and other variables to capture factors over which cross-country insurance contracts may not be written. Conditional on uninsurable control variables, there should be a unit coefficient on world consumption. This is decisively rejected in his estimates.

Kollmann (1995) studied the isoelastic utility model for G7 countries using data from 1972-1988 and lagged values of y and x as instruments. He found that the J-test did not reject, and interpreted this as supportive of the incomplete markets model. But $\hat{\gamma}$ was negative in about a third of cases, and insignificantly different from zero in many cases. Overall, then, there often were unpredictable departures from a cloud of data points, not from an upward sloping line.

Head, Mattina, and Smith (2004) noted that the incomplete-markets condition (8) is also necessary for the complete-markets condition (5). They studied that condition by

GMM for 10 OECD countries from 1961 to 2001. They examined traditional, isoelastic utility as well as models in which the marginal utility of consumption depends on government expenditure, real money balances, or external habit. They also considered models with exogenously missing asset markets but an endogenous discount rate that anchors the distribution of wealth and with endogenous market segmentation. Statistical tests reject all models of marginal utility, with one conspicuous exception. The model with external habit (which thus involves a moving average of consumption growth) passes the J-test and also yields significant parameter estimates with signs consistent with theory.

Drawing economic lessons from GMM estimation of the risk-sharing condition may be hampered by the problem of weak instruments. For many countries both consumption and real exchange rates (at least under floating nominal rates) are near random walks, which means that their growth rates are very difficult to predict. But drawing inference from instrumental-variables estimation requires valid instruments with significant predictive power. Without such instruments, standard confidence intervals may have incorrect coverage and the J-test may not have its nominal size. Neely, Roy, and Whiteman (2001), Stock and Wright (2000), and Yogo (2004) have drawn attention to the weak-instrument problem in estimating preference parameters.

An instrument of particular interest is $u_t = \iota$, a vector of ones. In that case, the estimating equations for the incomplete-markets condition become:

$$\mathrm{E}y_{t+1}^{ij} - \mathrm{E}\gamma x_{t+1}^{ij} = 0. \tag{15}$$

This is simply the monotonicity condition in means used by Backus and Smith (1993) and Obstfeld (2007). Provided these moments exist, then, the rejections of the risk-sharing condition using unconditional means apply to both the complete-markets and incomplete-markets versions.

There are potential pitfalls with the monotonicity-in-means test, though. First, the unconditional means may not be constant over history, so that the sample mean may not converge to a well-defined, population mean. Second, the unconditional means may be identical across countries so that the scatter plot or regression (12) does not identify γ. For example, suppose that the unconditional means of differential consumption growth, y_t^{ij}, and of real exchange rate growth, x_t^{ij} were both zero for all pairs of countries. Then these means could not be used to assess the slope of the conditional risk-sharing relationship.

What about the covariance test? This test does not naturally extend to the condition for risk-sharing with incomplete markets, essentially because the law of iterated expectations does not apply to second moments. Thus,

$$\mathrm{E}(y_{t+1}^{ij} - \gamma x_{t+1}^{ij} | z^t) = 0 \nRightarrow \mathrm{cov}_t(y_{t+1}^{ij}, x_{t+1}^{ij} | u_t) > 0. \tag{16}$$

If the instrument set coincides with the information used by market participants, so $u_t \equiv z^t$, then the conditional covariance is simply the covariance between their two forecast errors, which is not restricted by the theory. At the other extreme, if $u_t \equiv \iota$ then the covariance is unconditional, which again is not restricted. Indeed, the theoretical models with incomplete markets are designed specifically to be consistent with a correlation between y_{t+1}^{ij} and x_{t+1}^{ij} that is not positive. For example, Bodenstein (2008) studies a model in which asset markets are endogenously incomplete due to limited enforcement. He shows that it can reproduce a negative correlation using numerical examples. However, the covariance or correlation between forecasts *is* restricted:

$$\mathrm{E}(y_{t+1}^{ij} - \gamma x_{t+1}^{ij} | z^t) = 0 \Rightarrow \mathrm{cov}_t\left[(\mathrm{E}(y_{t+1}^{ij} | u_t), (\mathrm{E} x_{t+1}^{ij} | u_t)\right] > 0. \tag{17}$$

Implementing this test requires instrumental variables, as in the GMM estimation described earlier.

To sum up, the risk-sharing condition with incomplete markets can be tested in three ways:

a. A complete model that includes this condition can be tested indirectly by seeing whether it can reproduce a negative time-series covariance or correlation in actual outcomes, treating that like any stylized fact. But this approach does not provide a direct test of the condition. And it requires simulation of a complete model.

b. The monotonicity-in-means test across countries is a direct test of the incomplete markets condition. Disadvantages of this test are that it requires a significant number of countries and that it does not allow us readily to study the evolution of the condition over time. The test may be misleading if there is a trend or break in the unconditional moments. And it may be uninformative if the unconditional means do not differ across countries (for example if they are zero). In section 6 below, we provide an analytical example where the monotonicity-in-means test fails.

c. The condition can be tested by GMM (or the covariance between conditional expectations). One practical difficulty with this approach is that instruments may be weak, so that it may be challenging to find time-series variation in the forecasts. Both consumption growth, $\dot{c}_t^i$ and nominal depreciation, $\dot{s}_t^{ij}$, are notoriously difficult to forecast econometrically. A second practical difficulty arises from the need to construct real-time instruments rather than using revised data to represent historical expectations.

Our method in this paper is to directly test the risk-sharing condition using professional forecasts. This approach allows us to avoid the weak-instrument problem and yet to use numerous time-series observations for specific pairs of countries. The forecasts automatically are on a real-time basis. And we thus can study the evolution of the risk-sharing condition over time, with multiple observations per country pair.

Forecast-based tests of necessary conditions also have been applied by Smith and Yetman (2007) and Engel and Rogers (2008). Smith and Yetman study the Euler equation linking forecasts of inflation, consumption growth, and nominal interest rates in the US *Survey of Professional Forecasters.* Engel and Rogers study the present-value model of the current account using long-horizon forecasts for the G7 from *Consensus Economics.*

4. Forecast Data

Each month, *Consensus Economics* (`www.consensuseconomics.com`) surveys professional forecasters in numerous countries. The forecasters make predictions for macroeconomic variables including real consumption growth and CPI inflation rates for the current and subsequent calendar years. They also make predictions for nominal exchange-rate depreciations, this time at fixed horizons of 3, 12, and 24 months. Our data consist of all available observations on 12-month forecasts for all three of $\dot{c}$, $\dot{p}$, and $\dot{s}$ from the December surveys. Only in that month are the horizons for the foreign exchange predictions aligned with those from the predictions for consumption growth and inflation. (We cannot use the 24-month foreign-exchange predictions made in December 2004, for example, because forecasters are not asked their predictions for $\dot{c}$ and $\dot{p}$ for 2006 until January 2005.) These forecasts are thus denoted E_{t-12}.

We use the mean forecast, which is the summary statistic reported by Consensus Economics. They report individual forecasters' predictions for consumption growth and inflation but not for depreciations. And the set of forecasters is specific to each country. For both these reasons we cannot use disaggregated, individual forecasts.

Table 2 lists the 28 economies we study, along with the date at which their forecasts began. All series run to December 2007. We used all possible data points from this source. Because the data are proprietary, the appendix contains further notes on the forecasts, so that our work can be replicated or updated.

Table 2: Countries and Starting Years

Argentina	1993	Indonesia	1994	South Korea	1994
Australia	1993	Italy	1999	Spain	1999
Brazil	1994	Japan	1989	Sweden	1996
Canada	1989	Malaysia	1994	Switzerland	1998
Chile	1993	Mexico	1993	Thailand	1994
Columbia	1997	Netherlands	1999	UK	1989
Euro area	2002	New Zealand	1994	USA	1989
France	1999	Norway	1998	Venezuela	1993
Germany	1996	Peru	1998		
Hong Kong	1994	Singapore	1994		

We use 389 country-year combinations. We generally construct country pairs relative to the US, treating it as country i. But some currency forecasts are for rates of depreciation against the DM or Euro, so for those economies — France, Italy, the Netherlands, Norway, Spain, Sweden, and Switzerland — we use Germany as country i.

Most country-pairs have floating exchange rates. The exceptions are (a) Hong Kong and the US; (b) Argentina and the US for 1993-2000; and (c) the Euro-area economies each relative to Germany beginning in 1999. Within the Euro area after 1999 we thus do not need to align forecasts of consumption growth and inflation (which are on a calendar-year basis) with currency forecasts (which instead have fixed horizons) and so can use the monthly data for pairs chosen from France, Italy, the Netherlands, and Spain, relative to Germany. The forecasts are thus made from 23 to 0 months in advance of the end of the calendar year to which they apply. For this time period and set of countries we have 2 forecasts made each month (for two calendar years), 4 economy pairs, and 115 months (from January 1999 to July 2008) for a total of 920 observations. We also consider a start date of January 2002, when the new currency was introduced, leaving 632 observations.

The operator E_{t-h} denotes the forecast for calendar year t made h months in advance. Thus the risk-sharing condition in forecasts is:

$$\mathrm{E}_{t-h} y_t^{ij} = \gamma \mathrm{E}_{t-h} x_t^{ij}. \tag{18}$$

The condition predicts that the forecasts lie on an upward-sloping line, with slope γ. Our reporting method is simple. For pairs ij and year or month t we plot predicted differential consumption growth, $\mathrm{E}_{\mathrm{t-h}} y_t^{ij}$ on the vertical axis and predicted real depreciation, $\mathrm{E}_{t-h} x_t^{ij}$ on the horizontal axis. We plot these average forecasts for various pairs of countries, years, and for floating and fixed exchange rates.

These statistics are transitive; the point for pair jk automatically lies on the line connecting pair ji and pair ki. An upward-sloping line is evidence that the risk-sharing condition holds. So as not to overstate the impression created by the scatter plots, then, when we pool country pairs we plot only pairs relative to a base country i (generally the US).

We also compare the findings to those using realized, historical data from the IMF's *International Financial Statistics* (*IFS*). The *IFS* data apply to exactly the same set of countries and years as do the forecast data (except that we have forecasts for 2008 but not yet realized data), so that we can compare the state-by-state and forecast versions of the risk-sharing condition.

5. Evidence

Figure 2 presents evidence for all countries and years. The rate of real depreciation x_t^{ij} is on the horizontal axis, while the rate of relative, real consumption growth, y_t^{ij} is on the vertical axis. The first panel shows the realized outcomes (in dark red) and the corresponding economy-pair means (in light red). The latter correspond to the statistics graphed by Backus and Smith (1993) and Obstfeld (2007). The second panel then shows the scatter plot of forecasts (in dark blue) and the corresponding economy-pair mean forecasts (in light blue). The third panel in figure 1 then collects the realization means (in light red) and forecast means (in light blue) for each economy pair.

The conclusions from figure 2 are straightforward. For this recent data and wide range of countries, there is no evidence of an upward-sloping relationship in actual outcomes or their means or in forecasts. (We also constructed a version of figure 2 that applies only to pairs of economies with floating nominal exchange rates. The findings are the same.)

Figure 3 provides the data for individual economy-pairs. Realized outcomes again are in red and forecast outcomes in blue. Recall that the second traditional way to examine the evidence is to report the country-specific covariance or correlation in outcomes. Figure 3 shows that, for this group of economies, frequency, and time span, a negative correlation is *not* a stylized fact. For Argentina, Brazil, Chile, and New Zealand, each relative to the US, the red, realized outcomes slope down. But for the remaining 22 pairs the scatter plots are better described as clouds of points, with little correlation apparent.

The second question we can ask with the economy-pairs is whether there is greater evidence of a positive relationship in forecasts than in outcomes. An informal, visual inspection suggests that for Australia, Canada, Columbia, Indonesia, Malaysia, Mexico, and Venezuela, each relative to the US, (7 of the 26 economy pairs) there is such evidence of an upward-sloping pattern in forecasts to a greater extent than in realizations. As shown by the theoretical studies discussed in section 2, this discrepancy between realizations and forecasts can be consistent with a general equilibrium model where real exchange rates and consumption are determined simultaneously, with country-specific shocks and incomplete markets.

It is natural to wonder whether the results of tests for international risk-sharing trend over time. Flood, Marion and Matsumoto (2008), for example, present a measure of international consumption risk sharing (although without a control for real exchange rate changes): the variance of the log of a country's share of world consumption. They show that this variance has been trending down for industrialized countries, indicating gradual increases in risk-sharing. In a similar vein, we also disaggregated the data (both forecasts and realizations) by year, instead of by economy-pair. The results (not shown) show no trend to an upward slope in the scatter plots in both the forecasts and the realized data.

Several scholars have argued that the puzzle arises principally with floating exchange rates. Hess and Shin (2006), for example, show the correlation puzzle is largely due to the behaviour of the nominal exchange rate in OECD countries. They also study *intra-national* data from US states, and find that the correlation there (where the nominal exchange rate is one) is positive. Hadzi-Vaskov (2008) studies quarterly data for the Euro area for 1999-2006 and applies the correlation test. He finds a positive correlation for countries with a common currency, while the correlation remains negative for pairs of countries with floating nominal exchange rates over the same period. (These papers estimate regressions in the actual data – and so use the correlation test – but also use country-by-country and pooled instrumental-variables estimation.)

Figure 4 shows our findings for economy pairs with a currency board or common currency: Argentina–US 1992-2001; Hong Kong–US; France–Germany 1999-2007; Italy–Germany 1999-2007; Netherlands–Germany 1999-2007; and Spain–Germany 1999-2007. Again, whether in actual data or in forecasts or in means there is no evidence of an upward slope. In fact, these scatter plots tend to slope down, at least in means.

Finally, recall that for the Euro-area economies we have monthly forecasts of annual consumption growth and inflation rates, at 24 monthly horizons. We combine these monthly data to provide further evidence on the risk-sharing tests in this currency union. Figure 5 contains these intra-Euro area results, for France, Italy, the Netherlands, and Spain relative to Germany since 1999. (The results were very similar when we instead compared each country to the remaining Euro members as a group instead of to Germany.)

For this time period, set of economies, and set of horizons there is now considerably more evidence of a positive slope. The same is true for some individual economy pairs, shown in figure 6. The scatter plots for the Netherlands-Germany and Spain-Germany in particular slope up.

Figure 7 then disaggregates these high-frequency forecasts according to the year for which predictions apply. Again the idea is to see whether there is evidence of a trend in the slope linking the two conditional expectations, taking advantage of the additional data provided by this approach. Figure 7 does show an upward slope for several years, but it is hard to argue that there is any pattern over time.

Overall, then, the evidence suggests that the consumption-real exchange rate anomaly remains unresolved. The puzzle lies not just in the lack of a positive correlation between the movement in consumption differences and real exchange rates, but also in the absence of a positive relationship between the conditional expectations of consumption growth differences and real depreciations, where conditional expectations can be represented by forecasts. While the first failure (a problem for the complete-markets model) can in principal be resolved by a combination of limited financial assets (a non-contingent bond economy, for instance) and a judicious choice of parameters and pattern of shocks (as outlined by Benigno and Theonissen 2008, or Corsetti, Dedola, and Leduc 2008), the second is more problematic. Most international macroeconomic models with any form of capital mobility imply a version of (8). But condition (8) seems to be rejected in most of the forecasts.

One perspective on this evidence is given by relating it to the findings from forecast-based tests of the consumption Euler equation. Defining i_t as the one-period, nominal interest rate, we may decompose the consumption/exchange-rate linkage as follows:

$$\mathrm{E}\big[\dot{c}^i_{t+1}|z^t\big]-\mathrm{E}\big[\dot{c}^j_{t+1}|z^t\big]=\gamma\big(i^i_{t+1}-\mathrm{E}\big[\pi^i_{t+1}|z^t\big]\big)-\gamma\big(i^j_{t+1}-\mathrm{E}\big[(\pi^j_{t+1}|z^t\big]\big)=\gamma\mathrm{E}\big[\dot{q}^{ij}_{t+1}|z^t\big]. \quad (19)$$

The first equality comes from the consumption Euler equation; expected consumption growth is determined by the real interest rate. The second equality follows with the addition of the uncovered interest rate parity equation. Thus, a failure of condition (8) may be attributed to the breakdown of one or both of these relationships. Using professional forecasts, Smith and Yetman (2007) and Engel and Rogers (2008) find no support for the first relationship. Engel and Rogers's estimates with G7 forecasts show a strong relationship between (forecast) consumption growth and GDP growth, but no significant link between consumption growth and long term real interest rates. From US forecasts, Smith and Yetman (2007) find a small but significant *negative* relationship between short term interest rates and expected consumption growth. Of course, the second relationship is also highly suspect. The empirical failure of UIRP is well known. Movement in short term interest rates tend to be negatively related to exchange rates, rather than positively, as the theory suggests.

6. Hand-to-Mouth Consumers

Next we outline a simple framework which may help to explain the failure of the consumption-real exchange rate relationship in forecasts. We take the evidence of Smith and Yetman (2007) and Engel and Rogers (2008) as a guide. The high correlation between forecasts of consumption and income growth, and the weak connection between forecasted consumption and interest rates suggest that forecasters anticipate that household spending will be dictated primarily by current income. A simple way to implement this is to allow for a portion of households to be 'hand-to-mouth' consumers, in the sense of Campbell and Mankiw (1990). For hand-to-mouth consumers, consumption is simply equal to current income. But we emphasize that this interpretation does not simply add exogenous, noisy consumption growth to the scatter plots. Instead it drives a new pattern linking consumption and real exchange rates endogenously.

The following framework illustrates the effect of these consumers. Suppose that there are only two countries, labeled i and j. To keep notation simple, we omit the time subscript where appropriate. Households in country i, for example, receive random, non-storable endowments of a traded good w^i and a non-traded good v^i. The countries are of equal size, with household preferences defined over consumption of non-traded goods and an aggregate of home and foreign traded goods. Again in country i, households consume d_N^i of non-traded goods and an aggregate d_T^i of traded goods, in turn composed of a^i of their domestically-produced traded good, and b^i of the traded good imported from country j. We differentiate between two, country-specific traded goods because movements in the terms of trade may play an important role in real exchange rate responses, given home bias in preferences over traded goods.

Endowments of traded and non-traded goods are labelled w and v respectively. Market clearing thus requires that:

$$\begin{aligned} a^i + a^j &= w^i \qquad d_N^i = v^i \\ b^i + b^j &= w^j \qquad d_N^j = v^j. \end{aligned} \tag{20}$$

We next outline the consumption and price indexes, omitting the country-i superscript for ease of reading. Utility depends on an aggregate of traded goods, denoted d_T, given by:

$$d_T = [\mu^{\frac{1}{\theta}} a^{(1-\frac{1}{\theta})} + (1-\mu)^{\frac{1}{\theta}} b^{(1-\frac{1}{\theta})}]^{\frac{\theta}{\theta-1}}, \tag{21}$$

where $\mu > 0.5$ indicates home bias in tastes over the local traded good, and $\theta > 0$ measures the elasticity of substitution across goods. Then aggregate consumption is given by a higher-level CES aggregator:

$$c = [\alpha^{\frac{1}{\epsilon}} d_N^{(1-\frac{1}{\epsilon})} + (1-\alpha) d_T^{(1-\frac{1}{\epsilon})}]^{\frac{\epsilon}{\epsilon-1}}, \tag{22}$$

where α is the expenditure share on non-traded goods, common across countries, and ϵ is the elasticity of substitution between the traded and non-traded goods.

The prices of the two traded goods are p_a and p_b. The price index p_T for the composite traded good is defined as:

$$p_T = \left[\mu p_a^{(1-\theta)} + (1-\mu) p_b^{(1-\theta)}\right]^{\frac{1}{(1-\theta)}}. \tag{23}$$

The price of the non-traded good is denoted p_N. Allowing for the consumption of both traded and non-traded goods, the overall CPI is defined as:

$$p = \left[\alpha p_N^{(1-\epsilon)} + (1-\alpha) p_T^{(1-\epsilon)}\right]^{\frac{1}{(1-\epsilon)}}. \tag{24}$$

In each country there is a measure 1 of households, but only $n < 1$ of these households have access to international capital markets. Even then, capital markets are incomplete, allowing trade only in the form of a non-contingent, one-period bond. Unconstrained households can trade in this bond, with their holdings denoted h_t^i. The other $1 - n$ (constrained) households consume the value of their income. Campbell and Mankiw (1990) estimated $1 - n$ to be between 0.4 and 0.5 for the US economy. To be conservative, we choose a value of $1 - n = 0.4$ in the numerical example below.

Let c_u and c_c denote consumption of the unconstrained and constrained households respectively. Total consumption is

$$c = nc_u + (1 - n)c_c. \tag{25}$$

The budget constraints for unconstrained and constrained households, respectively, are given as:

$$\begin{aligned} p_t c_{ut} + h_{t+1} &= p_{Nt} v_t + p_{at} w_t + (1 + r_t) h_t \\ p_t c_{ct} &= p_{Nt} v_t + p_{at} w_t, \end{aligned} \tag{26}$$

where we do not need separate notation for individual and aggregate measures because there is a measure 1 of households.

Table 3 summarizes the new notation in this section.

Table 3: Notation

w^i, w^j	traded-good endowments	v^i	non-traded good endowment
a^i	consumption of w^i in country i	b^i	consumption of w^j in country i
d_T^i	traded goods consumption	d_N^i	non-traded goods consumption
p_a	price of w^i	p_b	price of w^j
p_T^i	price of traded composite	p_N^i	price of non-traded good
μ	share of a^i in d_T^i	θ	elasticity between a^i and b^i
α	share of d_N^i in c^i	ϵ	elasticity between d_T^i and d_N^i
c_u	unconstrained consumption	c_c	constrained consumption
n	unconstrained share	h	bond holdings

All households consume the non-traded good, and both varieties of the traded good. In addition, unconstrained households choose an optimal pattern of inter-temporal spending, given that they have free access to international bond markets. Thus, unconstrained households satisfy the bonds-only risk-sharing condition we have studied so far in this paper:

$$\mathrm{E}\big[(\dot{c}^i_{ut+1} - \dot{c}^j_{ut+1})|z^t\big] = \mathrm{E}\big[\gamma \dot{q}^{ij}_{t+1}|z^t\big]. \tag{27}$$

Constrained households, by contrast, have consumption growth limited by income growth. Taking a linear approximation of the budget constraint for hand-to-mouth households, we can approximate the expected consumption growth difference between *constrained* households in economies i and j as:

$$\begin{aligned}\mathrm{E}\big[(\dot{c}^i_{ct+1} - \dot{c}^j_{ct+1})|z^t\big] = \gamma \mathrm{E}\big[&(1-\alpha)(\dot{w}^i_{t+1} - \dot{w}^j_{t+1}) + \alpha(\dot{v}^i_{t+1} - \dot{v}^j_{t+1}) \\ &- 2(1-\mu)(1-\alpha)(\dot{p}_{bt+1} - \dot{p}_{at+1})|z^t\big].\end{aligned} \tag{28}$$

Comparing the constrained households across countries, the relationship between relative consumption growth and real depreciations will depend on the scale of shocks to traded goods versus non-traded goods, and the movement in the terms of trade.

From condition (28), holding the terms of trade $p_b - p_a$ constant, a rise in the country's $i's$ expected growth rate of traded goods will directly increase expected consumption growth of $i's$ constrained households. The scale of this effect depends on the share of traded goods in GDP, $(1-\alpha)$. At the same time, expected growth in traded good endowments will generate an expected real exchange rate appreciation, since it raises expected demand for $i's$ non-traded goods. So expected consumption growth of constrained consumers will rise (relative to that of the foreign country), but the expected real exchange rate will fall. However, the terms of trade will also adjust. Expected growth in output of traded goods will be associated with an expected fall in the terms of trade, or a rise in $p_b - p_a$. This will reduce expected consumption growth for constrained consumers in country i. The impact of this on relative consumption will be less the higher is the degree of home bias in preferences of traded goods, μ. But for $\theta > 1$, the negative effect of a terms of trade deterioration will not be enough to offset the direct effect of expected growth in traded good endowments. Therefore expected consumption growth rises, while the expected real exchange rate falls.

On the other hand, if income variation is primarily driven by persistent shocks to non-traded endowments, expected consumption growth of constrained households will co-move positively with the real exchange rate. For example, a rise in the country i endowment of the non-traded good will lead to a rise in consumption and a real exchange rate depreciation.

Hence, in the aggregate, the relationship between expected relative consumption growth and expected real exchange rate changes for constrained households may be positive or negative, depending on the source of shocks. Obviously, then, as a theoretical matter, this illustrative model can allow for movements in aggregate expected consumption growth differentials that are either positively or negatively related to movements in expected real exchange rates, depending on both the source of shocks, and the measure of constrained versus unconstrained households.

We carry out a simple, quantitative exercise that yields scatter plots of forecasts that can be compared to the empirical scatter plots in section 5. We follow Benigno and Thoenissen's (2008) calibration closely. Take $\theta = 2$, $\epsilon = 0.44$, indicating a higher elasticity of substitution between home and foreign traded goods than between traded and non-traded goods. Let $\mu = 0.72$, so that there is home bias in the consumption of traded goods, and assume the elasticity of intertemporal substitution is $\gamma = 2$. Take the relative size of the non-traded goods sector to be $\alpha = 0.55$. Assume that the log endowments in the traded goods sector, w^i and w^j, follow AR(1) processes, with persistence 0.82, and an innovation standard deviation of 1.9 percent. In the non-traded goods sector, endowments v^i and v^j also follow AR(1) processes, but with persistence of 0.5, and an innovation standard deviation of 0.7 percent. Benigno and Thoenissen's shocks are represented by Solow residuals, as their model incorporates endogenous labour supply and capital accumulation. Qualitatively, this difference is unlikely to affect the nature of the illustration.

We simulate 100 draws of 100 time periods from these endowment processes. The consumption/real exchange rate correlations then reflect a mixture of n unconstrained households that satisfy the original forecasting restriction (27) and $1-n$ households that satisfy the constrained version (28). We calculate the economy-pair means over 100 time periods of forecasts, denoted $\{\overline{\mathrm{E}_{t-1}x^{ij}}, \overline{\mathrm{E}_{t-1}y^{ij}}\}$, where the forecasts use the correct, population values of the AR(1) parameters for the endowment processes. These calculations thus simulate the monotonicity-in-means tests in forecasts.

Figure 8 shows the results. The upper panel of figure 8 shows simulations for $n = 1$, so that all households in both countries can access international bond markets. The scatter plot shows the average, conditional, expected growth rates for the difference of consumption and the real exchange rate (in light blue, just as in the historical survey forecasts). The positive relationship reflects the actual risk-sharing condition (27) in this incomplete-markets economy. The model displays conditional risk sharing as facilitated by bond markets, despite the fact that the correlation between realized consumption growth and real exchange rates may be negative.

The lower panel of figure 8 introduces hand-to-mouth consumers into the model, and shows average, conditional, expected growth rates of consumption differences and real exchange rate changes. Recall that $n = 0.6$. Again the scatter shows the mean forecast from each of 100 simulations of 100 periods. For the calibration and shock processes chosen, there is no apparent relationship between expected consumption growth and expected real exchange rate changes. We see only a cloud, similar to the findings in the survey forecast data.

Of course, the results would be different with different parameters and shocks. If $n = 0$, so that all households are hand-to-mouth, then the forecast diagram would have a distinct downward sloping shape, since in this case persistent endowment shocks, mostly coming from traded goods, would lead to an anticipated negative relationship between relative consumption growth and real depreciations.

Studies of incomplete risk-sharing that address the Backus-Smith puzzle predict an upward-sloping plot in forecasts or in average forecasts, like the one in the top panel of figure 8. As we have seen, such a pattern is not observed in professional forecasts. Our simple interpretation in this section ratchets up the approach used in papers by Kollmann

(1995), Corsetti, Dedola, and Leduc (2008), and Benigno and Thoenissen (2008). By adding a further departure from the standard model, in the form of a 40 percent share of hand-to-mouth consumers, we construct an example that avoids an upward slope in forecasts, as shown in the lower panel of figure 8. But, consistent with our earlier argument, we emphasize that reproducing the cloud-like pattern in forecasts again is only indirect evidence in favor of this model economy, rather than providing a direct test.

In section 3, we argued that conventional measures of conditional risk-sharing — the monotonicity-in-means test — may not be able to detect a positive correlation between expectations of consumption growth and real exchange rates. That pitfall in the means test creates a need for more direct measures of forecasts like those in this paper. The simulations here provide quite stark evidence of this problem. Figure 9 illustrates the average of ex-post growth rates under the same simulations as figure 8. There is a distinct downward sloping relationship, for both $n = 1$ and $n = 0.6$. This illustrates the difficulty with the monotonicity-in-means test in outcomes. While the case with $n = 1$ (by construction) implies conditional risk sharing in the sense of condition (8), the average values of ex-post consumption growth in this sample simulation are negatively related to ex-post real exchange rate changes. To see why the means test gives a misleading answer here, recall the discussion in section 3. In this model, there is no difference across simulations in the unconditional means of consumption growth or real exchange rate growth. This corresponds to a situation where all countries have the same trend in y (x) (as before, this is short-hand notation for consumption growth differences, and real exchange rate growth). Moreover, the unconditional means of y and x are zero. These properties lead to a test where the sample averages are dominated by the unexpected shocks to y and x.

This point may be illustrated more precisely. In the model described in this section (with $n = 1$), we may write the equilibrium values of y and x as:

$$\begin{aligned} y_t &= \delta_w E_{t-1}\ln(w_t) + \delta_v E_{t-1}\ln(v_t) + \psi_w \epsilon_{wt} + \psi_v \epsilon_{vt} \\ x_t &= \delta_w E_{t-1}\ln(w_t) + \delta_v E_{t-1}\ln(v_t) + \omega_w \epsilon_{wt} + \omega_v \epsilon_{vt} \\ \ln(w_t) &= \mu_w \ln(w_{t-1}) + \epsilon_{wt} \\ \ln(v_t) &= \mu_v \ln(v_{t-1}) + \epsilon_{vt} \end{aligned} \tag{29}$$

where ϵ_{wt} (ϵ_{vt}) represents the mean-zero, *iid* innovation to the traded (non-traded) good endowment. For illustrative purposes here, we assume shocks only to the country i endowment, and omit the i label. Note that, as must be the case, $E_{t-1}y_t = E_{t-1}x_t$. But also, $Ey_t = Ex_t = 0$. Moreover, the structure of the model implies that $\psi_w > 0$, while $\omega_w < 0$, as an innovation in the country $i's$ w–endowment raises the value of home consumption, but leads to a real appreciation.

The monotonicity-in-means test compares the following statistics across simulations (countries):

$$\begin{aligned} E_s(y) &= \frac{1}{T}\sum_{t=0}^{T} y_{t,s} = \frac{1}{T}\Big(\sum_{t=0}^{T}[\delta_w\mu_w \ln(w_{t-1,s}) + \delta_v\mu_v \ln(v_{t-1,s})] + \sum_{t=0}^{T}[\psi_w\epsilon_{wt,s} + \psi_v\epsilon_{vt,s}]\Big) \\ E_s(x) &= \frac{1}{T}\sum_{t=0}^{T} x_{t,s} = \frac{1}{T}\Big(\sum_{t=0}^{T}[\delta_w\mu_w \ln(w_{t-1,s}) + \delta_v\mu_v \ln(v_{t-1,s})] + \sum_{t=0}^{T}[\omega_w\epsilon_{wt,s} + \omega_v\epsilon_{vt,s}]\Big), \end{aligned} \tag{30}$$

where $E_s(.)$ denotes the sample average over simulation s. The experiment is to compare $E_s(y)$ and $E_s(x)$ across S simulations. Imagine the extreme case where shocks are purely transitory, and there are only shocks to the traded goods sector. Then for each simulation s, we have:

$$E_s(y) = \psi_w \frac{1}{T}\sum_{t=0}^{T}\epsilon_{wt,s} = \frac{\psi_w}{\omega_w}E_s(x). \tag{31}$$

Since $\psi_w/\omega_w < 0$, the sample means for y and x are negatively correlated across simulations. For all $T < \infty$, the scatter plot of $E_s(y)$ and $E_s(x)$ across the S simulations will show a downward sloping line. Thus the means test for conditional risk-sharing is entirely uninformative here.

In a less extreme case, the presence of persistent shocks and shocks to the non-traded endowment as well as the traded good endowment will introduce a tendency towards a positive correlation between $E_s(y)$ and $E_s(x)$, both due to the the ϵ_{vt} shocks which raise both y and x simultaneously, and to persistent shocks which have identical effects on $E_s(y)$ and $E_s(x)$. Nevertheless, for the calibration underlying figure 9, which is dominated by shocks coming from the traded good sector, and for the given persistence parameters (as in Benigno and Thoenissen, 2008), the means test fails to provide any evidence of consumption risk sharing, despite the fact that it is built into the model (for case $n = 1$). This underscores the need to use an alternative approach to measuring conditional expectations.

In summary then, the absence of a clear relationship between consumption and real exchange rates in forecasts is consistent with a general equilibrium setting where international capital markets are restricted to a subset of households. This same setting reproduces a negative correlation in the means of actual outcomes *i.e.* the Backus-Smith puzzle. But we do not offer here a deeper analysis of the sources of these capital market constraints.

7. Conclusion

A range of international macroeconomic models with incomplete asset markets predict a positive relationship between expectations of relative consumption growth and real depreciation. We provide a direct test of this risk-sharing condition using professional forecasts. This method avoids (a) the syndrome of weak identification associated with instrumental-variables estimation and (b) the pitfalls of the monotonicity-in-means test. Compared to using unconditional means for economy pairs it also provides many more observations. This method thus can test for trends in risk-sharing over time.

Generally there is little evidence of a positive relationship between forecasts in the data for 1990-2008. (And a negative correlation for economy pairs is not really a stylized fact either.) The cloud-like scatter plots cast doubt on the incomplete asset-market models. However, when we take advantage of the high-frequency (monthly) forecasts within the Euro area we do find greater evidence of positive slope.

A simple endowment-economy example with a role for hand-to-mouth consumers shows some potential for reproducing the cloud-like patterns in actual forecasts. Overall, though, the empirical results deepen the risk-sharing puzzle.

Appendix. Forecast Data

The data come from the three newsletters *Consensus Forecasts*, *Latin American Consensus Forecasts*, and *Asia Pacific Consensus Forecasts* provided by Consensus Economics Inc. Forecast nominal depreciations are calculated from the levels forecasts. It is not clear whether the forecaster provides the forecast of the level or the growth rate, as when these two measures are generally compatible up to a rounding error. For forecasts of the Euro area, we use only measures relative to the US and not European countries, so as to avoid overlap in the jurisdictions which would bias results towards finding risk sharing.

We omitted Brazil in 1993, due to its being an extreme outlier, with an inflation forecast of 2241% and forecast depreciation of 55%. We omitted forecasts for Taiwan because its corresponding realized data are not in *IFS*. For December 2003, the Asia-Pacific spreadsheet indicates that forecasts are for 2004 and 2005 (instead of 2003 and 2004). Closer inspection — comparing the forecasts with those made one month earlier and one month later — indicates that this is a typographical error in the spreadsheet. Eastern European forecasts, provided in a separate newsletter, are bi-monthly until April 2007. These are excluded from our panel because they are made in November and January, and so cannot be aligned with the calendar-year forecasts of the exchange rates.

References

Backus, David K. and Gregor W. Smith (1993) Consumption and real exchange rates in dynamic economies with non-traded goods. *Journal of International Economics* 35, 297-316.

Benigno, Gianluca and Christoph Thoenissen (2008) Consumption and real exchange rates with incomplete markets and non-traded goods. *Journal of International Money and Finance* 27, 926-948.

Bodenstein, Martin (2008) International asset markets and real exchange rate volatility. *Review of Economic Dynamics* 11, 688-705.

Campbell, John Y. and N. Gregory Mankiw (1990) Permanent income, current income, and consumption. *Journal of Business and Economic Statistics* 8, 265-79.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan (2002) Can sticky price models generate volatile and persistent real exchange rates? *Review of Economic Studies* 69, 533-63.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc (2008) International risk sharing and the transmission of productivity shocks. *Review of Economic Studies* 75, 443-473.

Engel, Charles, and John H. Rogers (2008) Expected consumption growth from cross-country surveys: Implications for assessing international capital markets. mimeo, Department of Economics, University of Wisconsin Madison,*IMF Staff Papers* forthcoming.

Flood, Robert, Nancy Marion, and Akito Matsumoto (2008) International risk sharing during the globalization era. mimeo, International Monetary Fund.

Hadzi-Vaskov, Metodij (2008) Does the nominal exchange rate explain the Backus-Smith puzzle? Evidence from the Eurozone. Tjalling C. Koopmans Research Institute discussion paper 07-32.

Head, Allen C., Todd D. Mattina, and Gregor W. Smith (2004) Real exchange rates, preferences, and incomplete markets: evidence, 1961–2001. *Canadian Journal of Economics* 37, 782-801.

Hess, Gregory and Kwanho Shin (2006) Understanding the Backus-Smith puzzle: It's the (nominal) exchange rate, stupid. MPRA paper 696.

Kocherlakota, Narayana R. and Luigi Pistaferri (2007) Household heterogeneity and real exchange rates. *Economic Journal* 117, C1-C25.

Kollmann, Robert (1995) Consumption, real exchange rates and the structure of international asset markets. *Journal of International Money and Finance* 14, 191-211.

Neely, Christopher J., Amlan Roy, and Charles H. Whiteman (2001) Risk aversion versus intertemporal substitution: A case study of identification failure in the intertemporal consumption capital asset pricing model. *Journal of Business and Economic Statistics* 19, 395-403.

Obstfeld, Maurice (1994) Are industrial-country consumption risks globally diversified? in Leonardo Leiderman and Assaf Razin, eds., *Capital Mobility: The Impact on Consumption, Investment and Growth.* Cambridge: Cambridge University Press.

Obstfeld, Maurice (2007) International risk sharing and the costs of trade. Ohlin Lectures.

Opazo, Luis (2006) The Backus-Smith puzzle: The role of expectations. Central Bank of Chile working paper 395.

Selaive, Jorge D. and Vicente Tuesta (2003) Net foreign assets and imperfect pass-through: The consumption real exchange rate anomaly. Board of Governors of the Federal Reserve System International Finance Discussion Papers 764.

Smith, Gregor W. and James Yetman The curse of Irving Fisher (professional forecasters' version). Queen's Economics Department working paper 1144.

Stock, James H. and Jonathan H. Wright (2000) GMM with weak identification. *Econometrica* 68, 1055-1096.

Yogo, Motohiro (2004) Estimating the elasticity of intertemporal substitution when instruments are weak. *Review of Economics and Statistics* 86, 797-810.

Figure 1: Test Conflict

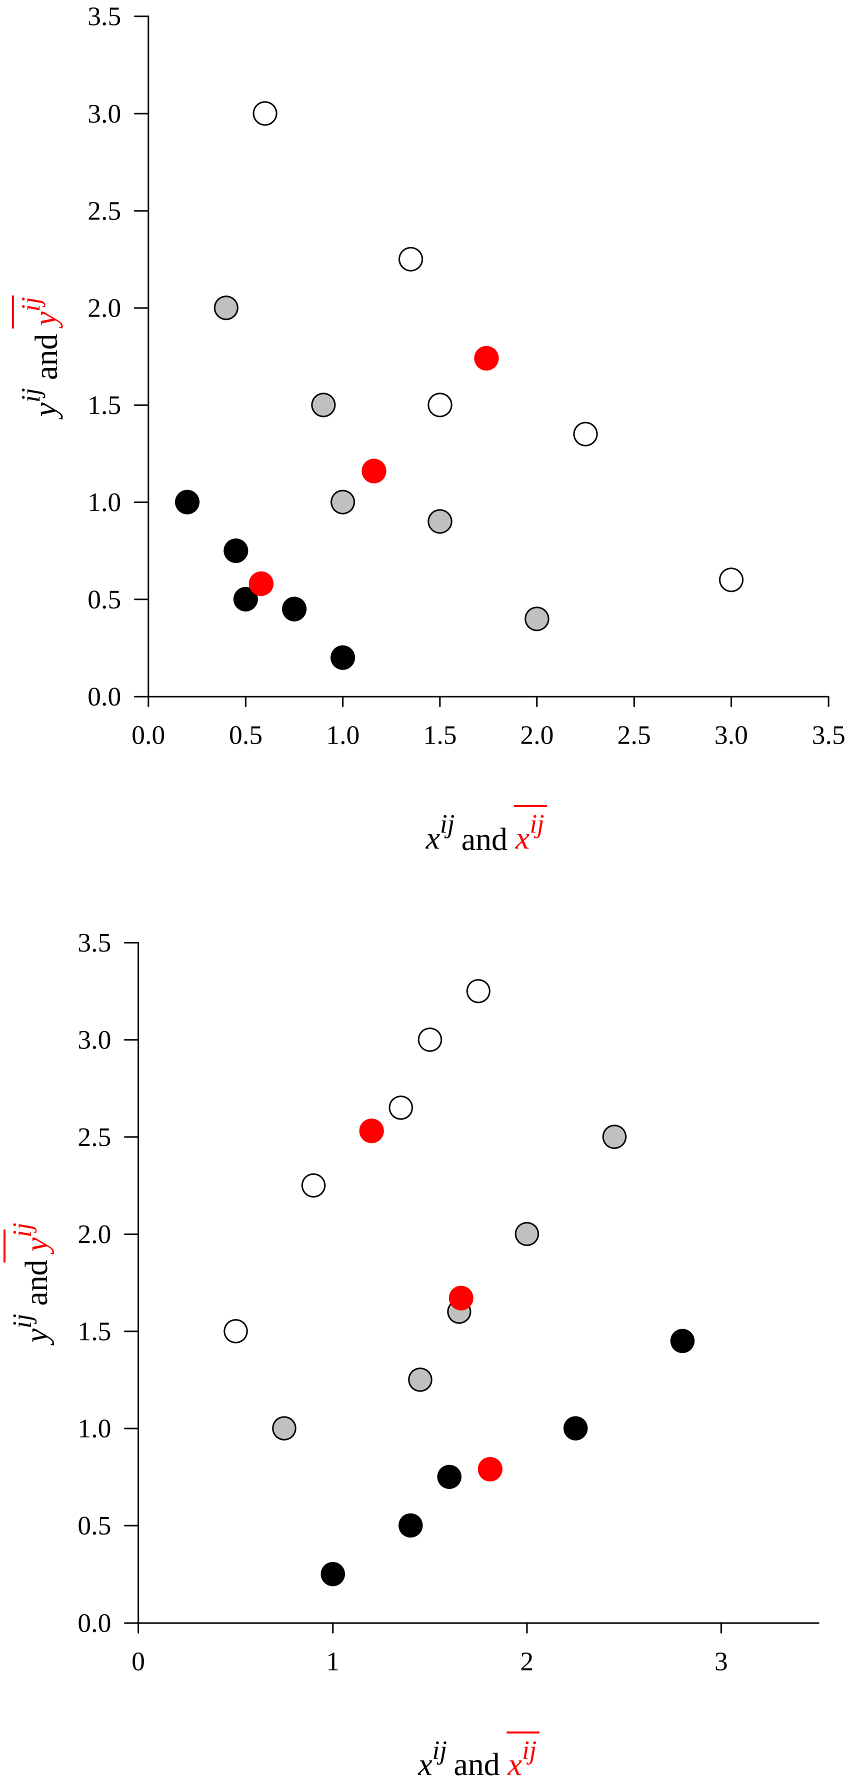

Figure 2: All Economies 1990-2008

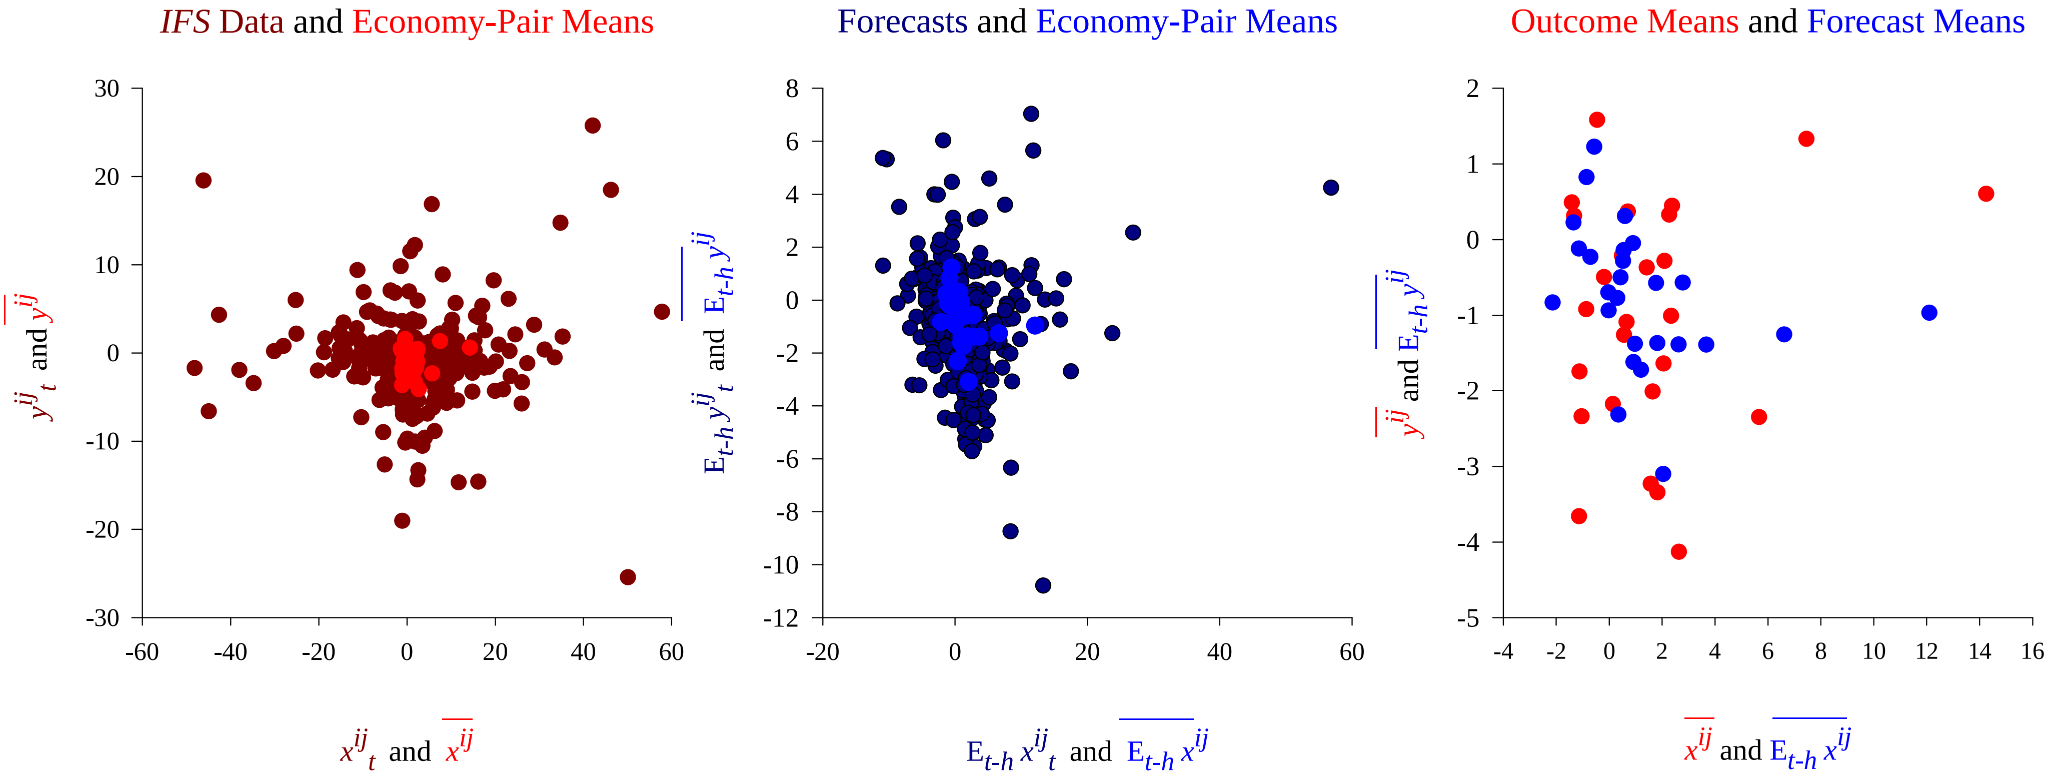

Figure 3: IFS and Forecast Data for Economy Pairs

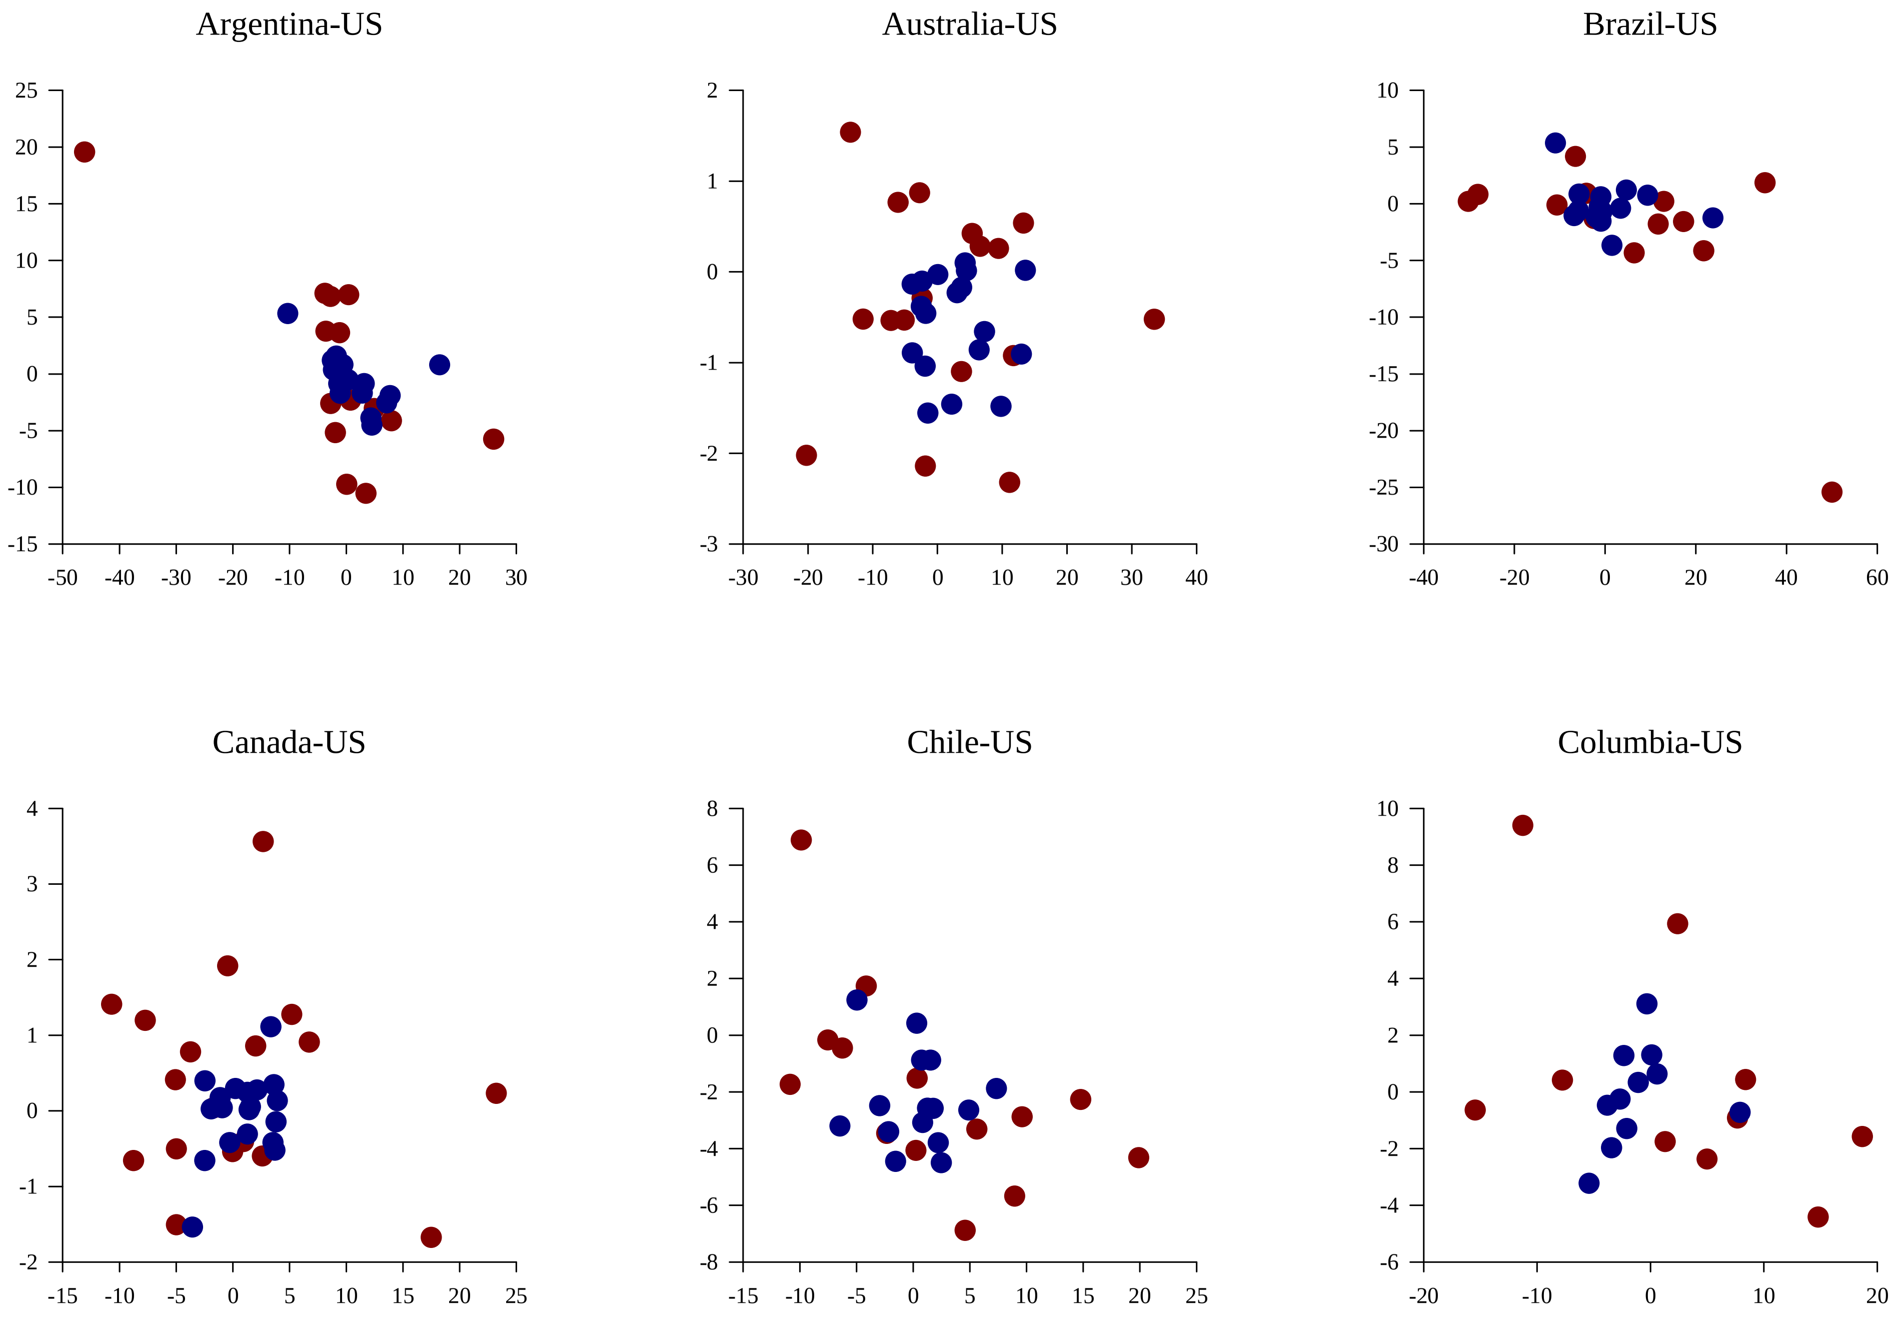

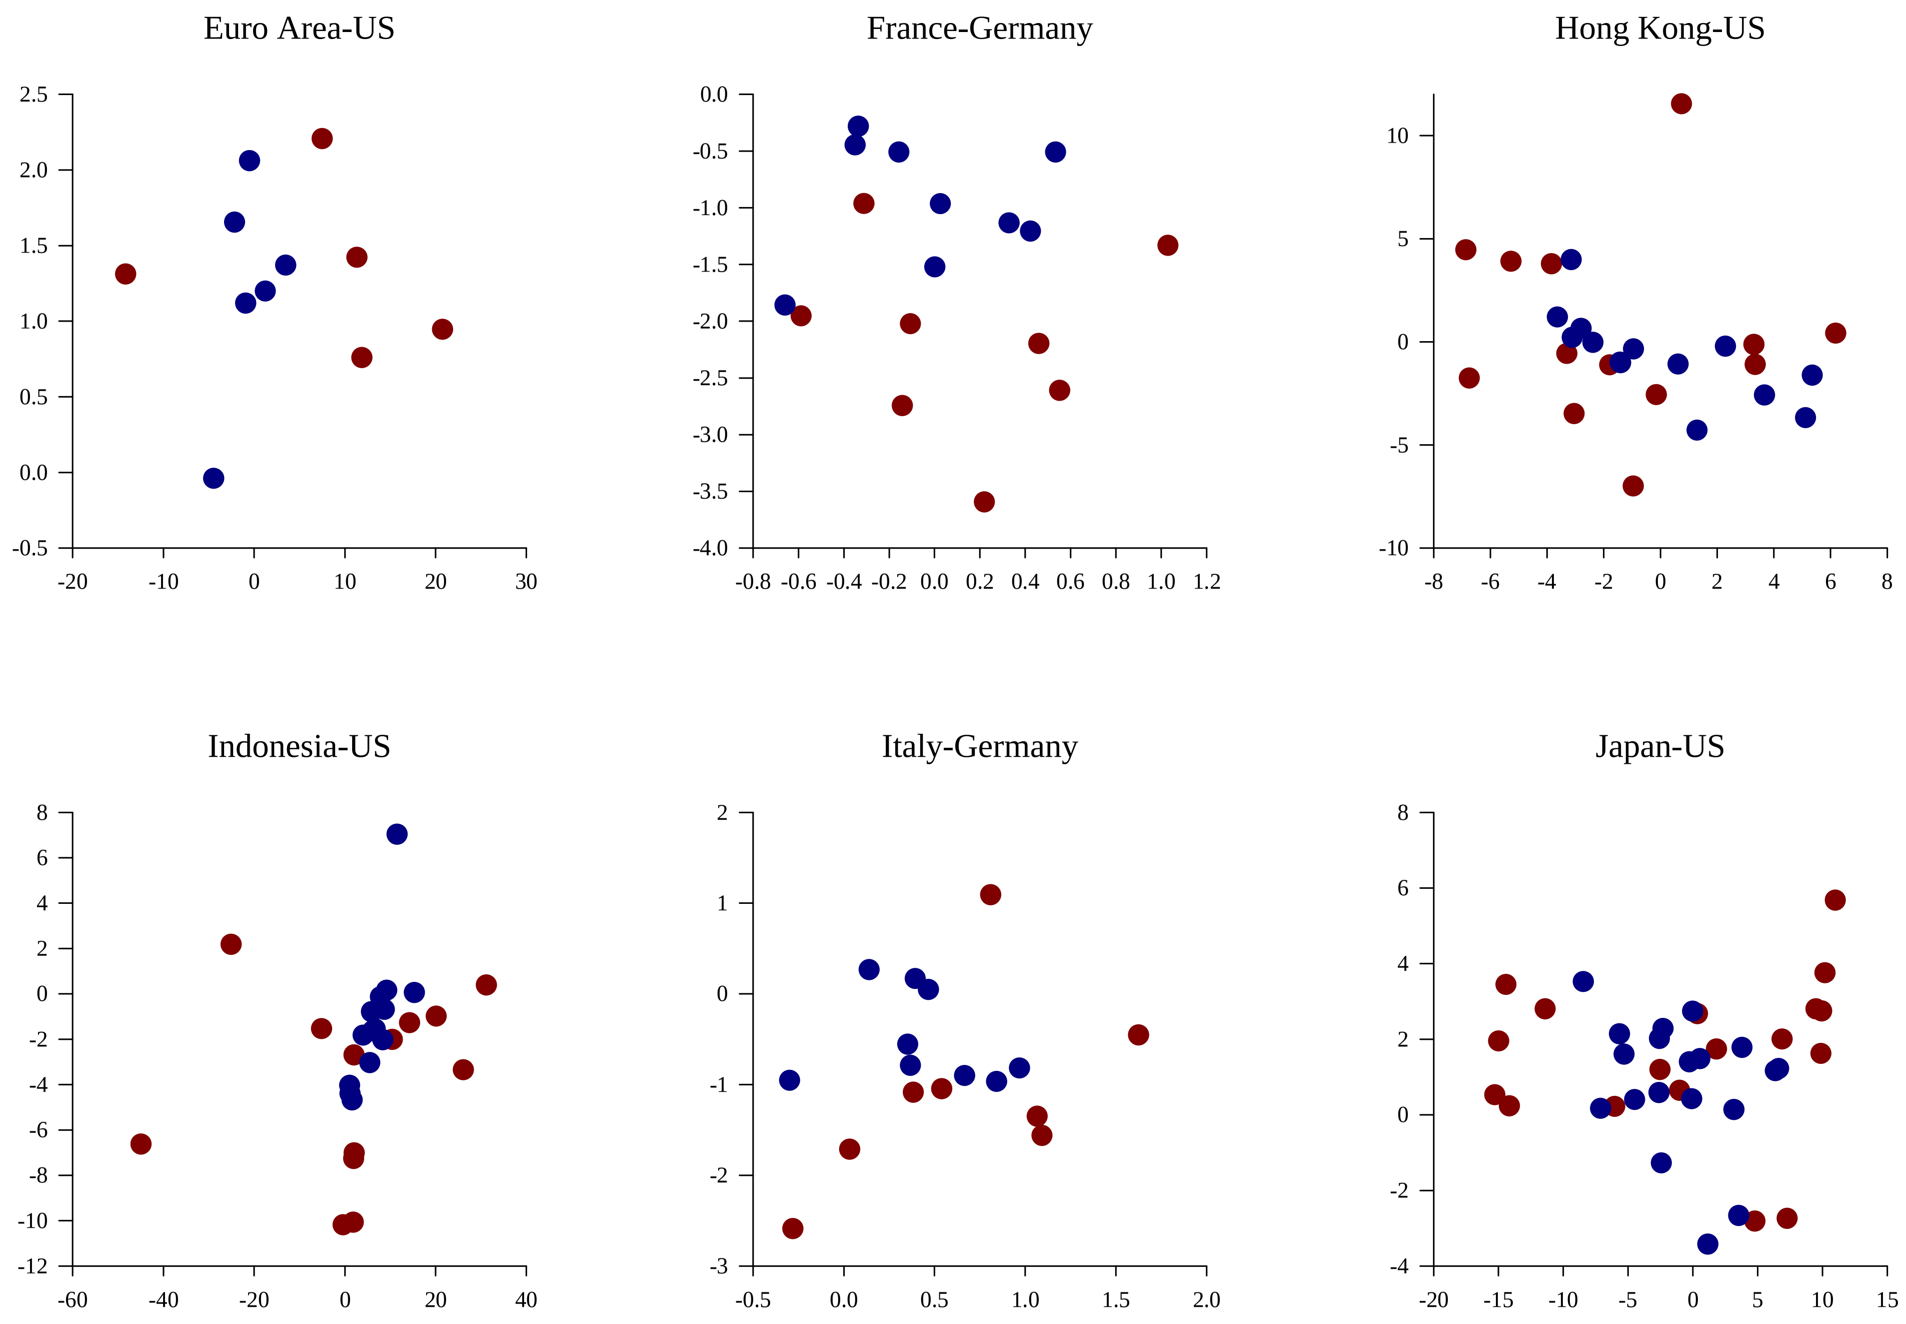
Euro Area-US
2.5
2.0
1.5
1.0
0.5
0.0
-0.5
-20
-10
0
10
20
30
France-Germany
0.0
-0.5
-1.0
-1.5
-2.0
-2.5
-3.0
-3.5
-4.0
-0.8
-0.6
-0.4
-0.2
0.0
0.2
0.4
0.6
0.8
1.0
1.2
Hong Kong-US
10
5
0
-5
-10
-8
-6
-4
-2
0
2
4
6
8
Indonesia-US
8
6
4
2
0
-2
-4
-6
-8
-10
-12
-60
-40
-20
0
20
40
Italy-Germany
2
1
0
-1
-2
-3
-0.5
0.0
0.5
1.0
1.5
2.0
Japan-US
8
6
4
2
0
-2
-4
-20
-15
-10
-5
0
5
10
15

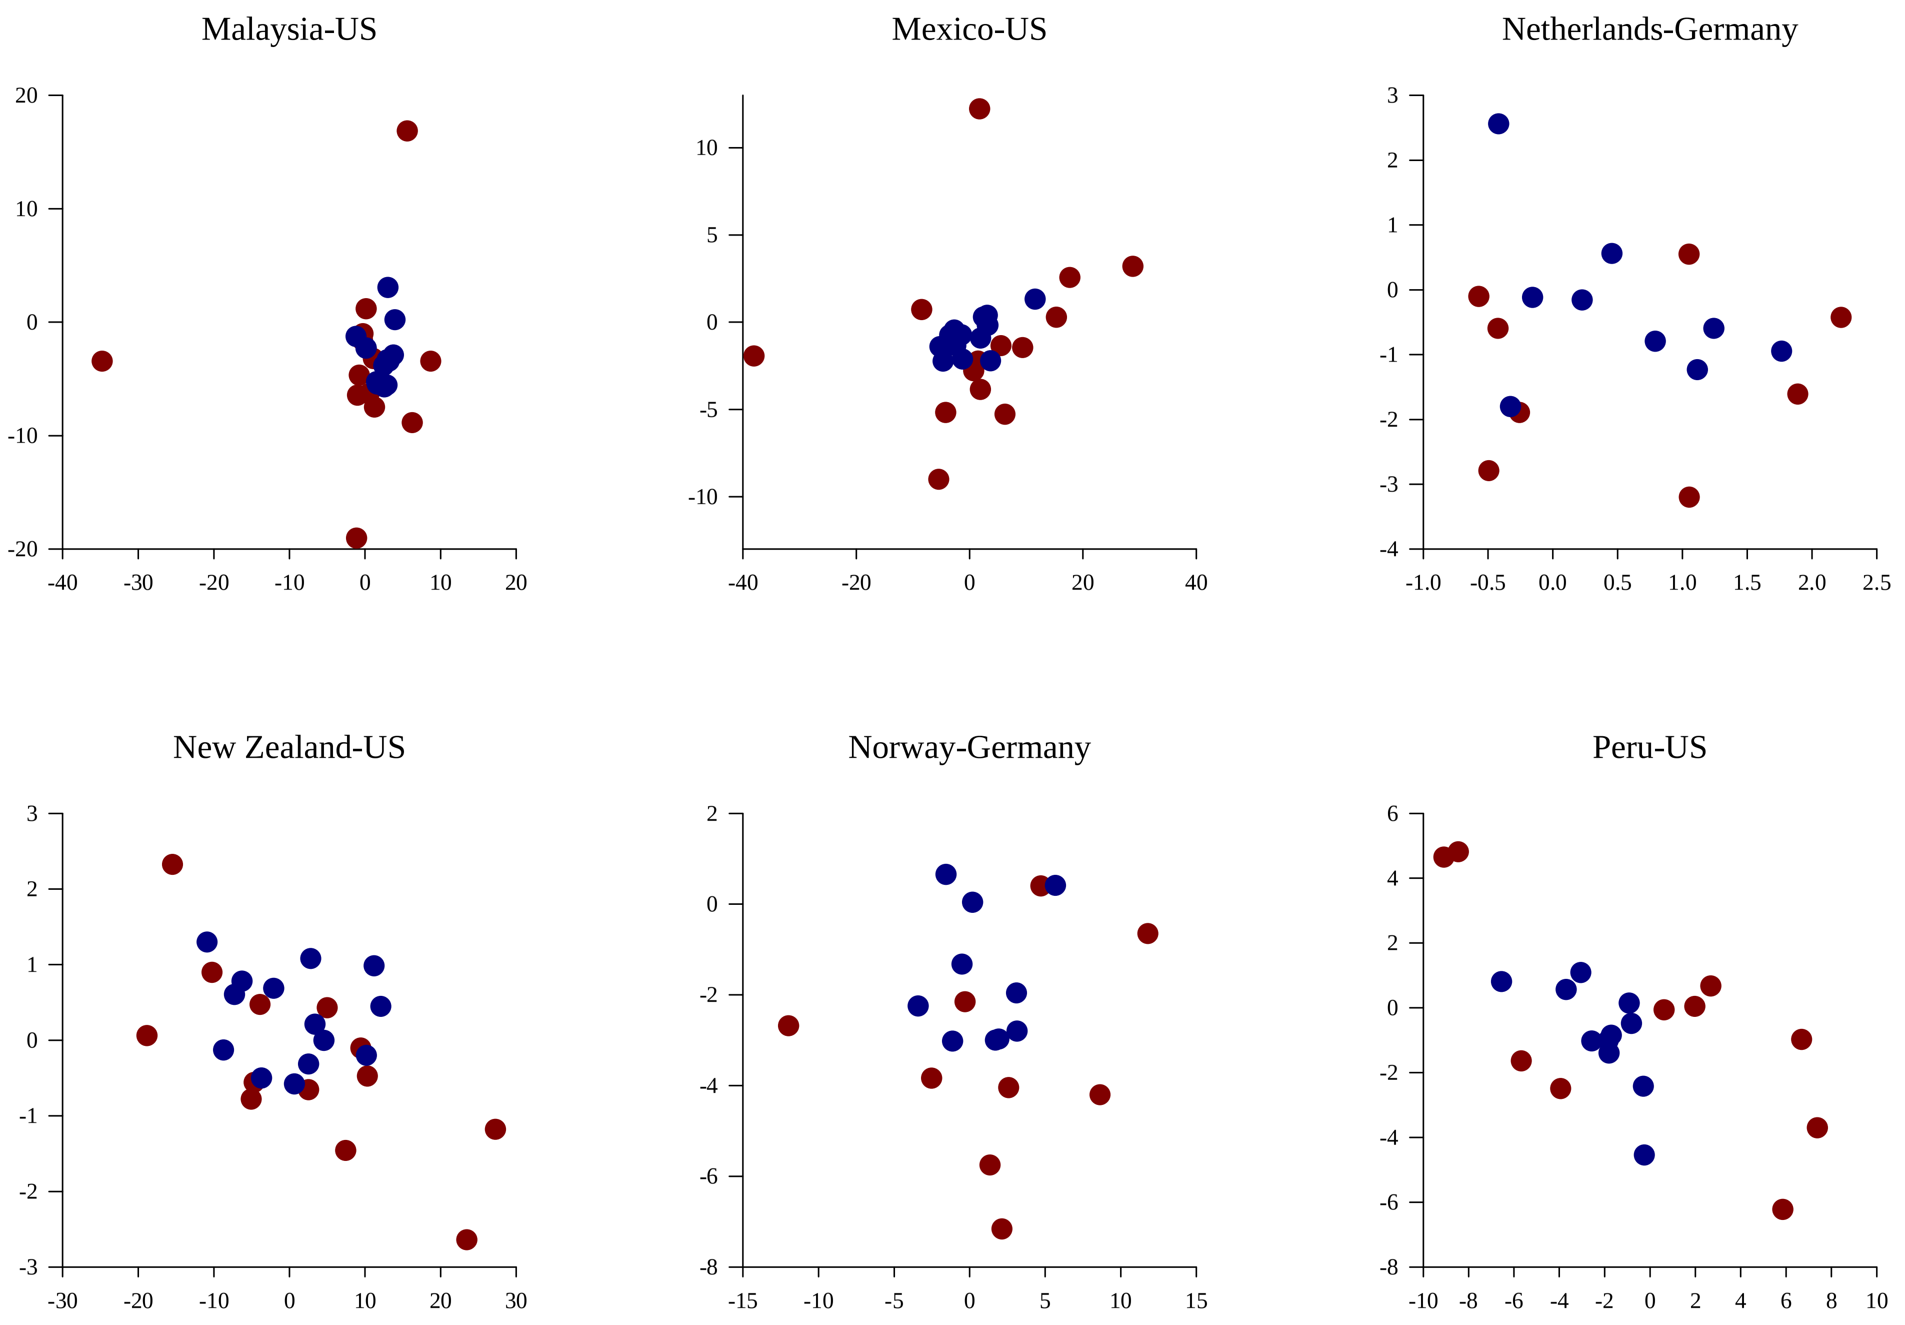
Malaysia-US
20
10
0
-10
-20
-40
-30
-20
-10
0
10
20
Mexico-US
10
5
0
-5
-10
-40
-20
0
20
40
Netherlands-Germany
3
2
1
0
-1
-2
-3
-4
-1.0
-0.5
0.0
0.5
1.0
1.5
2.0
2.5
New Zealand-US
3
2
1
0
-1
-2
-3
-30
-20
-10
0
10
20
30
Norway-Germany
2
0
-2
-4
-6
-8
-15
-10
-5
0
5
10
15
Peru-US
6
4
2
0
-2
-4
-6
-8
-10
-8
-6
-4
-2
0
2
4
6
8
10

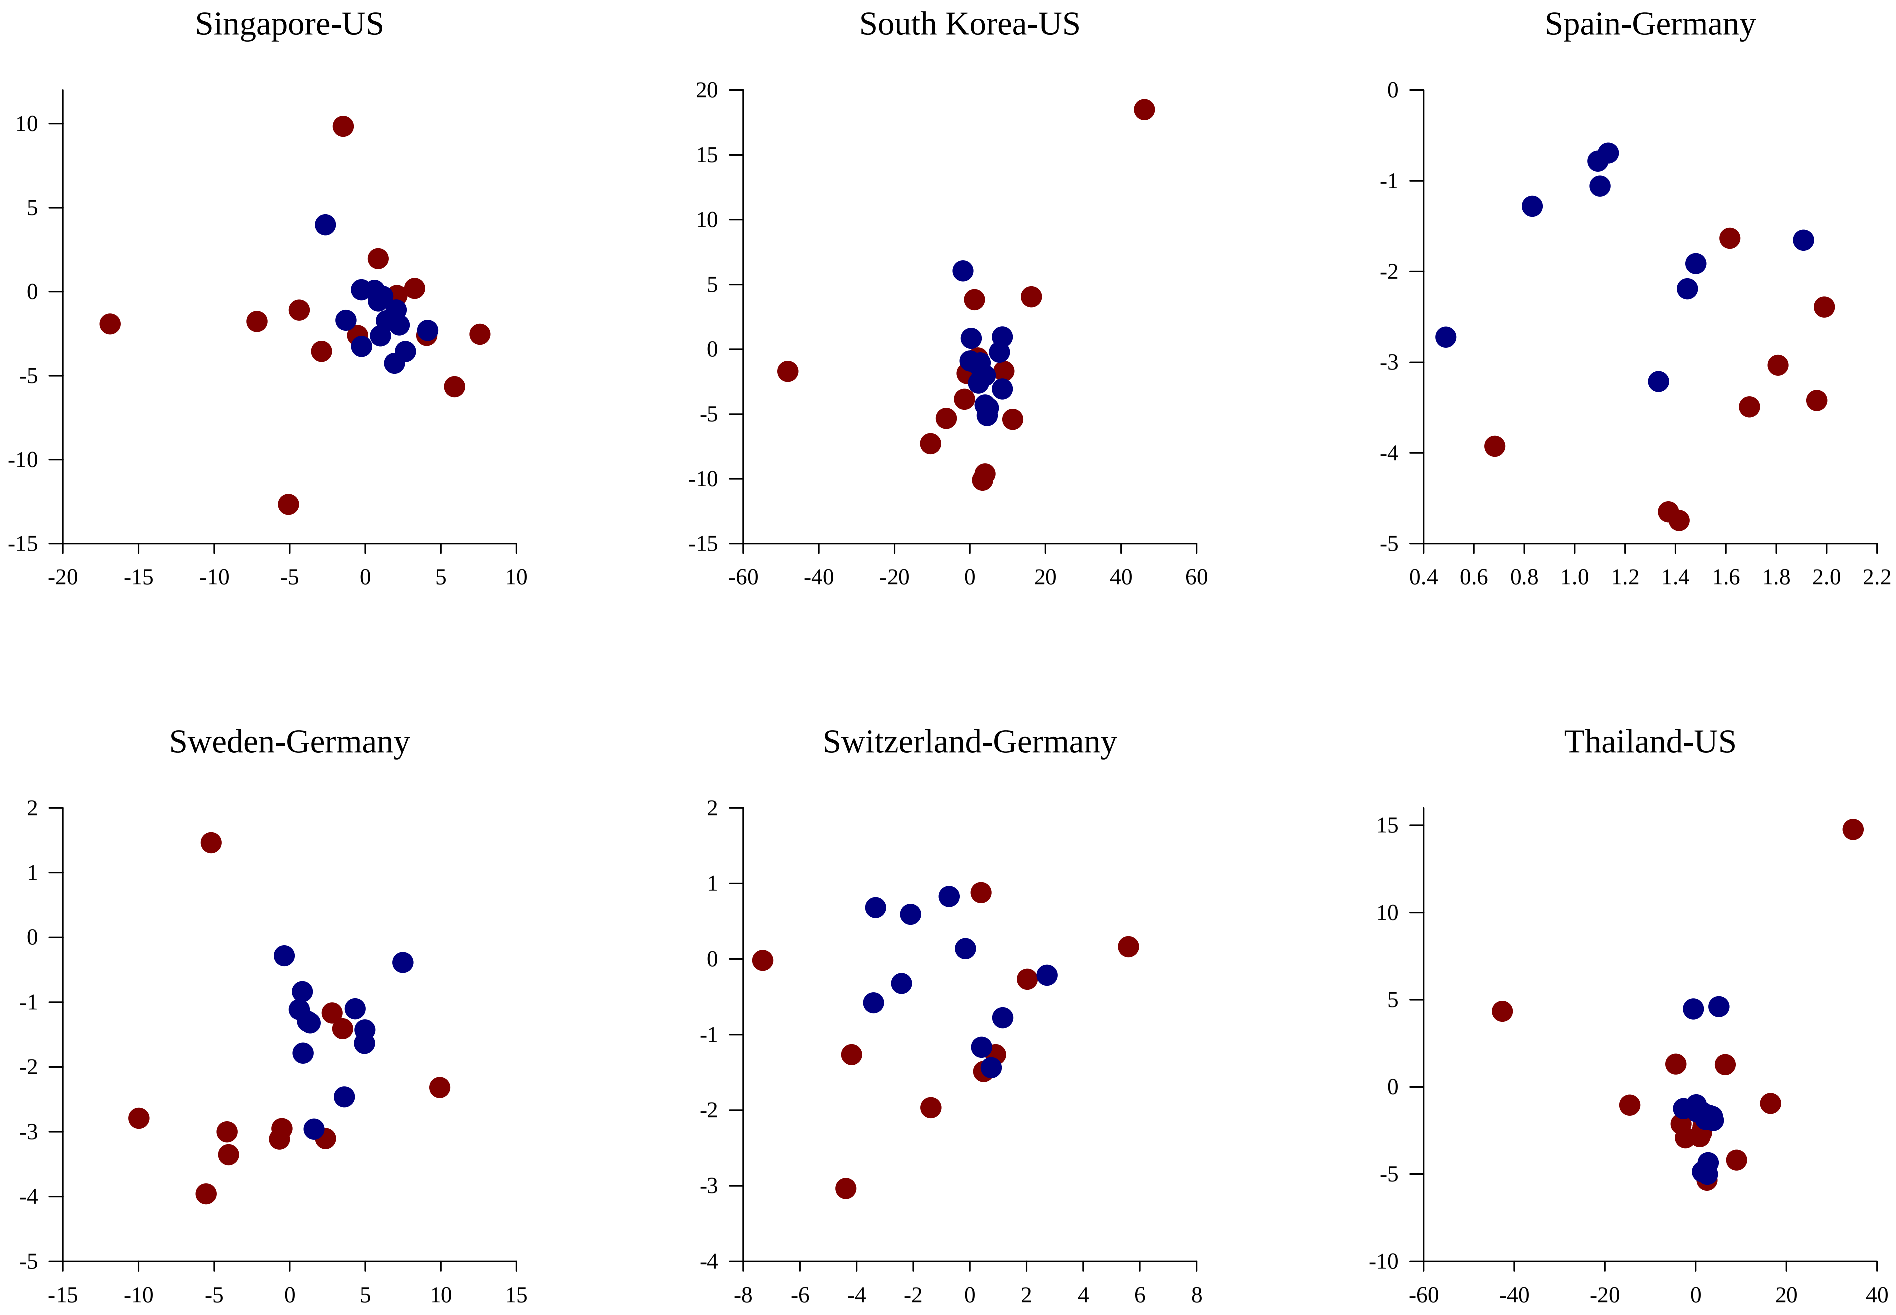
Singapore-US
South Korea-US
Spain-Germany
Sweden-Germany
Switzerland-Germany
Thailand-US
10
5
0
-5
-10
-15
-20
-15
-10
-5
0
5
10
20
15
10
5
0
-5
-10
-15
-60
-40
-20
0
20
40
60
0
-1
-2
-3
-4
-5
0.4
0.6
0.8
1.0
1.2
1.4
1.6
1.8
2.0
2.2
2
1
0
-1
-2
-3
-4
-5
-15
-10
-5
0
5
10
15
2
1
0
-1
-2
-3
-4
-8
-6
-4
-2
0
2
4
6
8
15
10
5
0
-5
-10
-60
-40
-20
0
20
40

UK-US

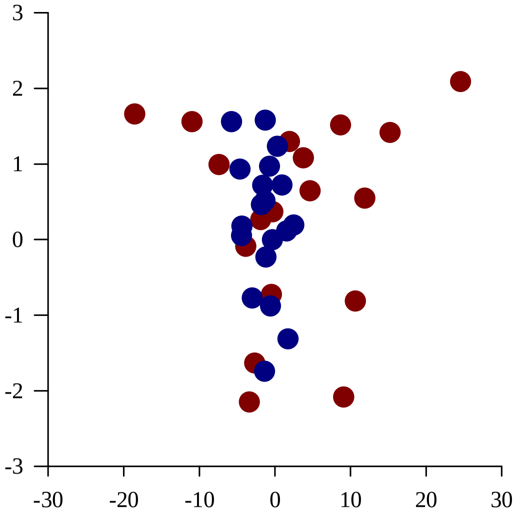

Venezuela-US

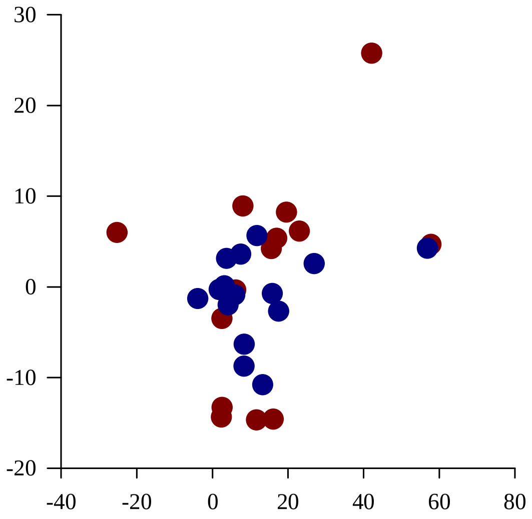

Figure 4: Currency-Union Economies

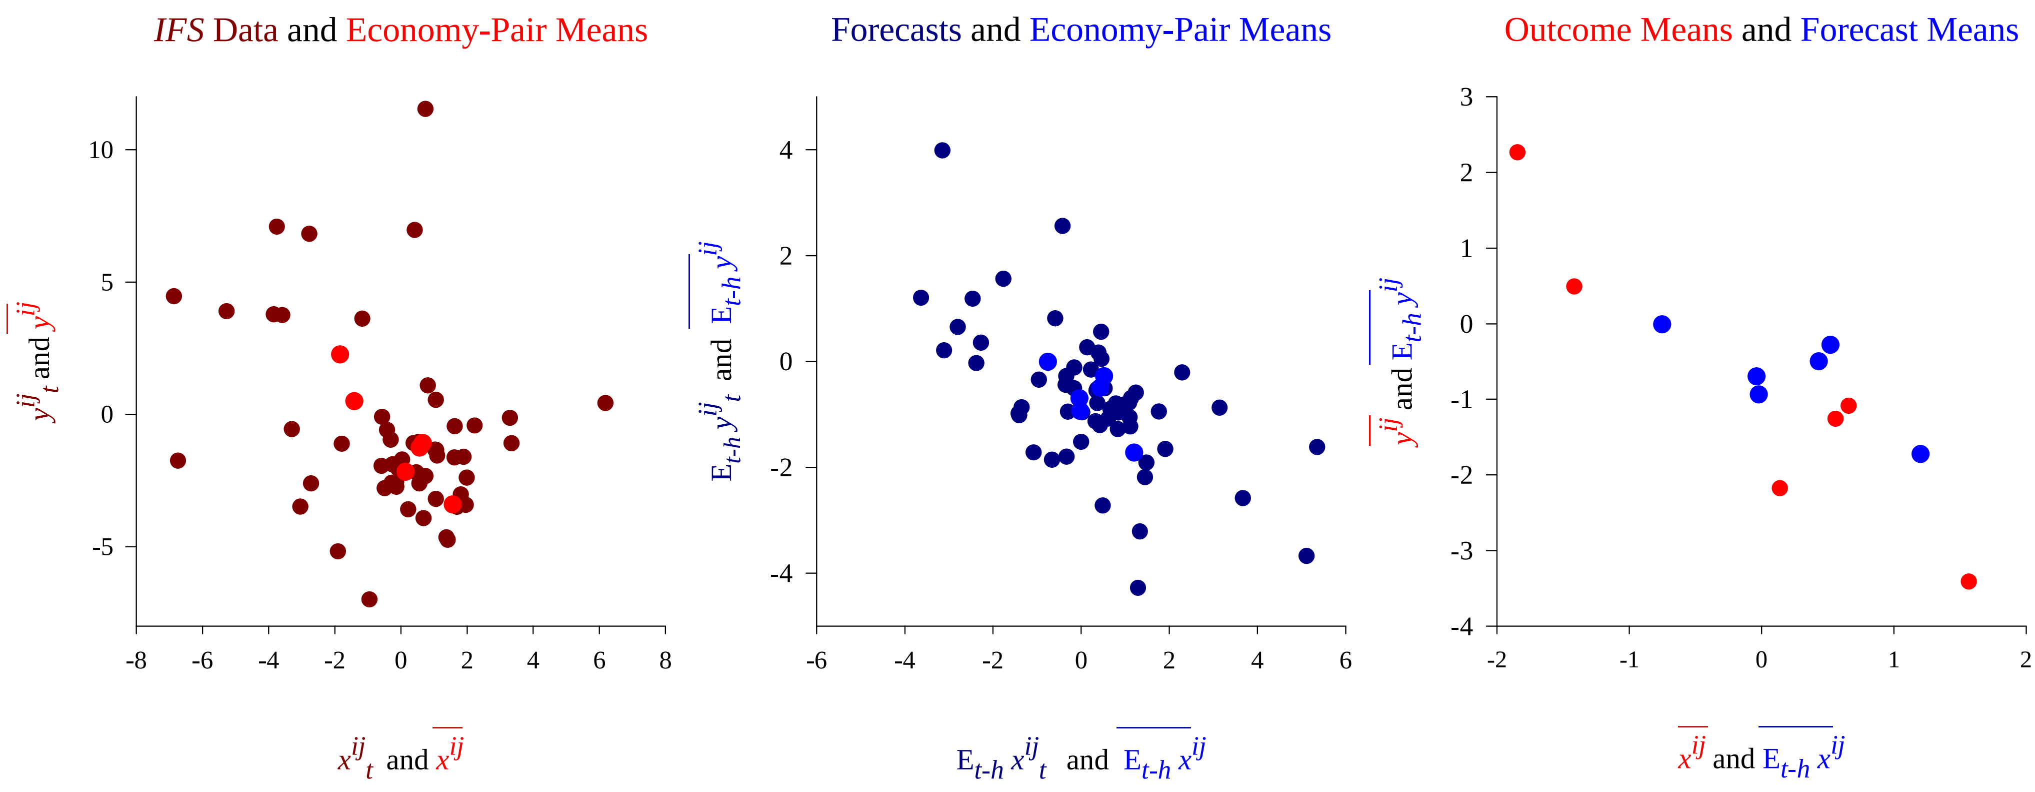

Figure 5: Euro Area Forecasts

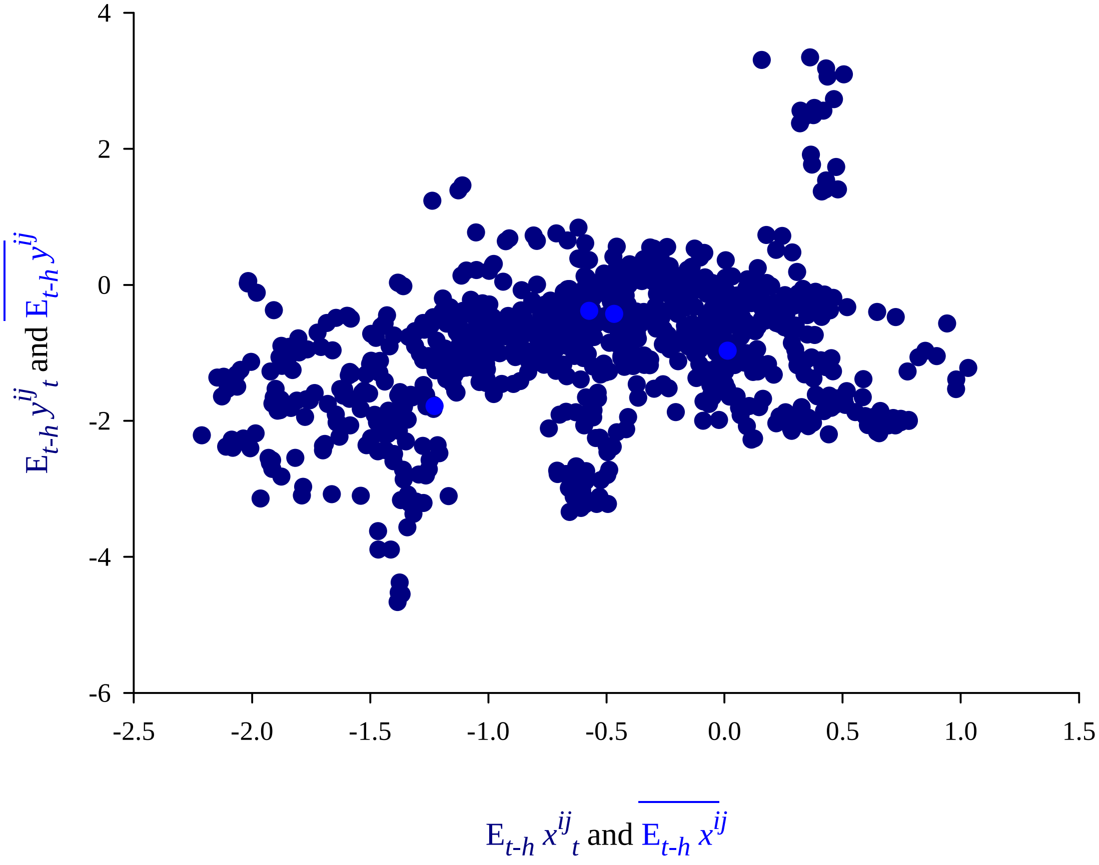

Figure 6: Euro Area Economy Pairs

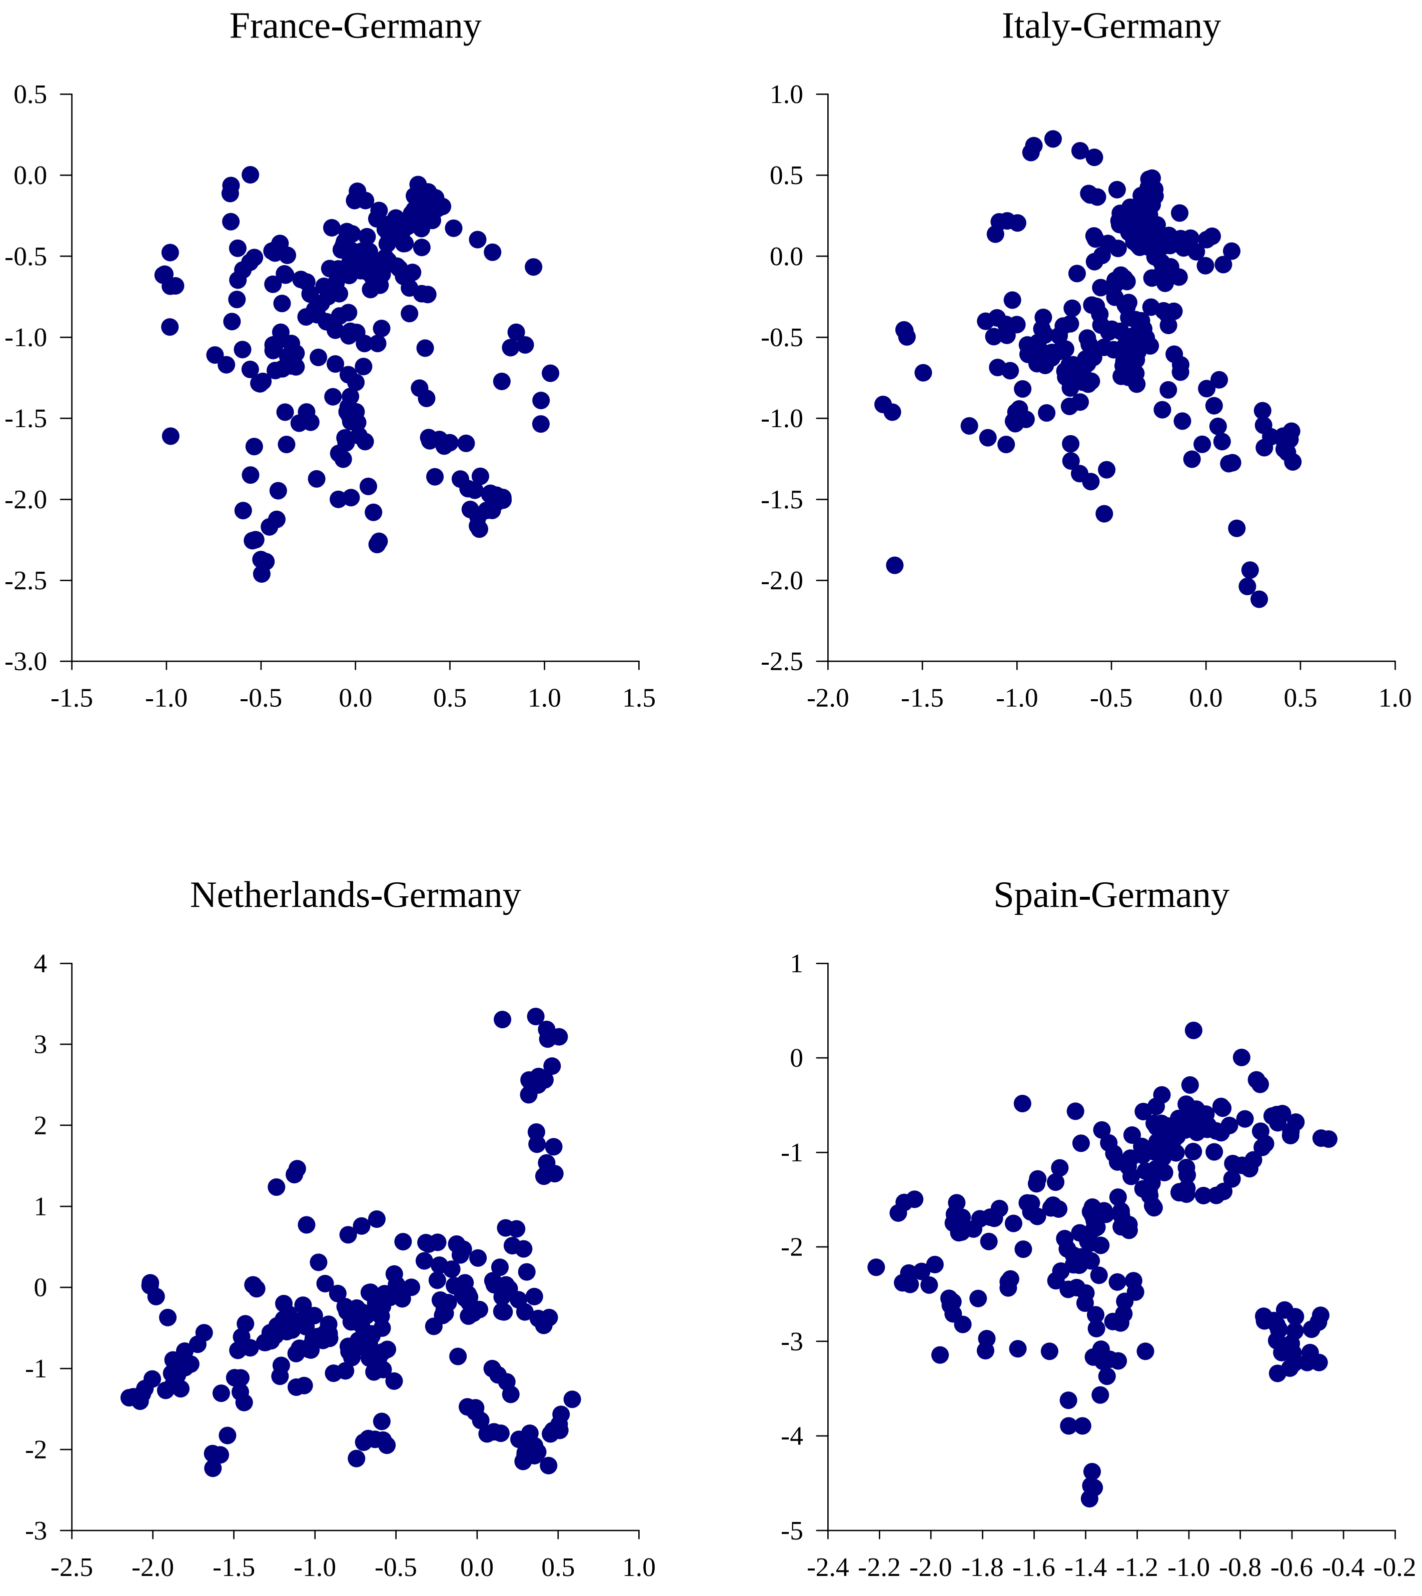

Figure 7: Euro Area Results by Year

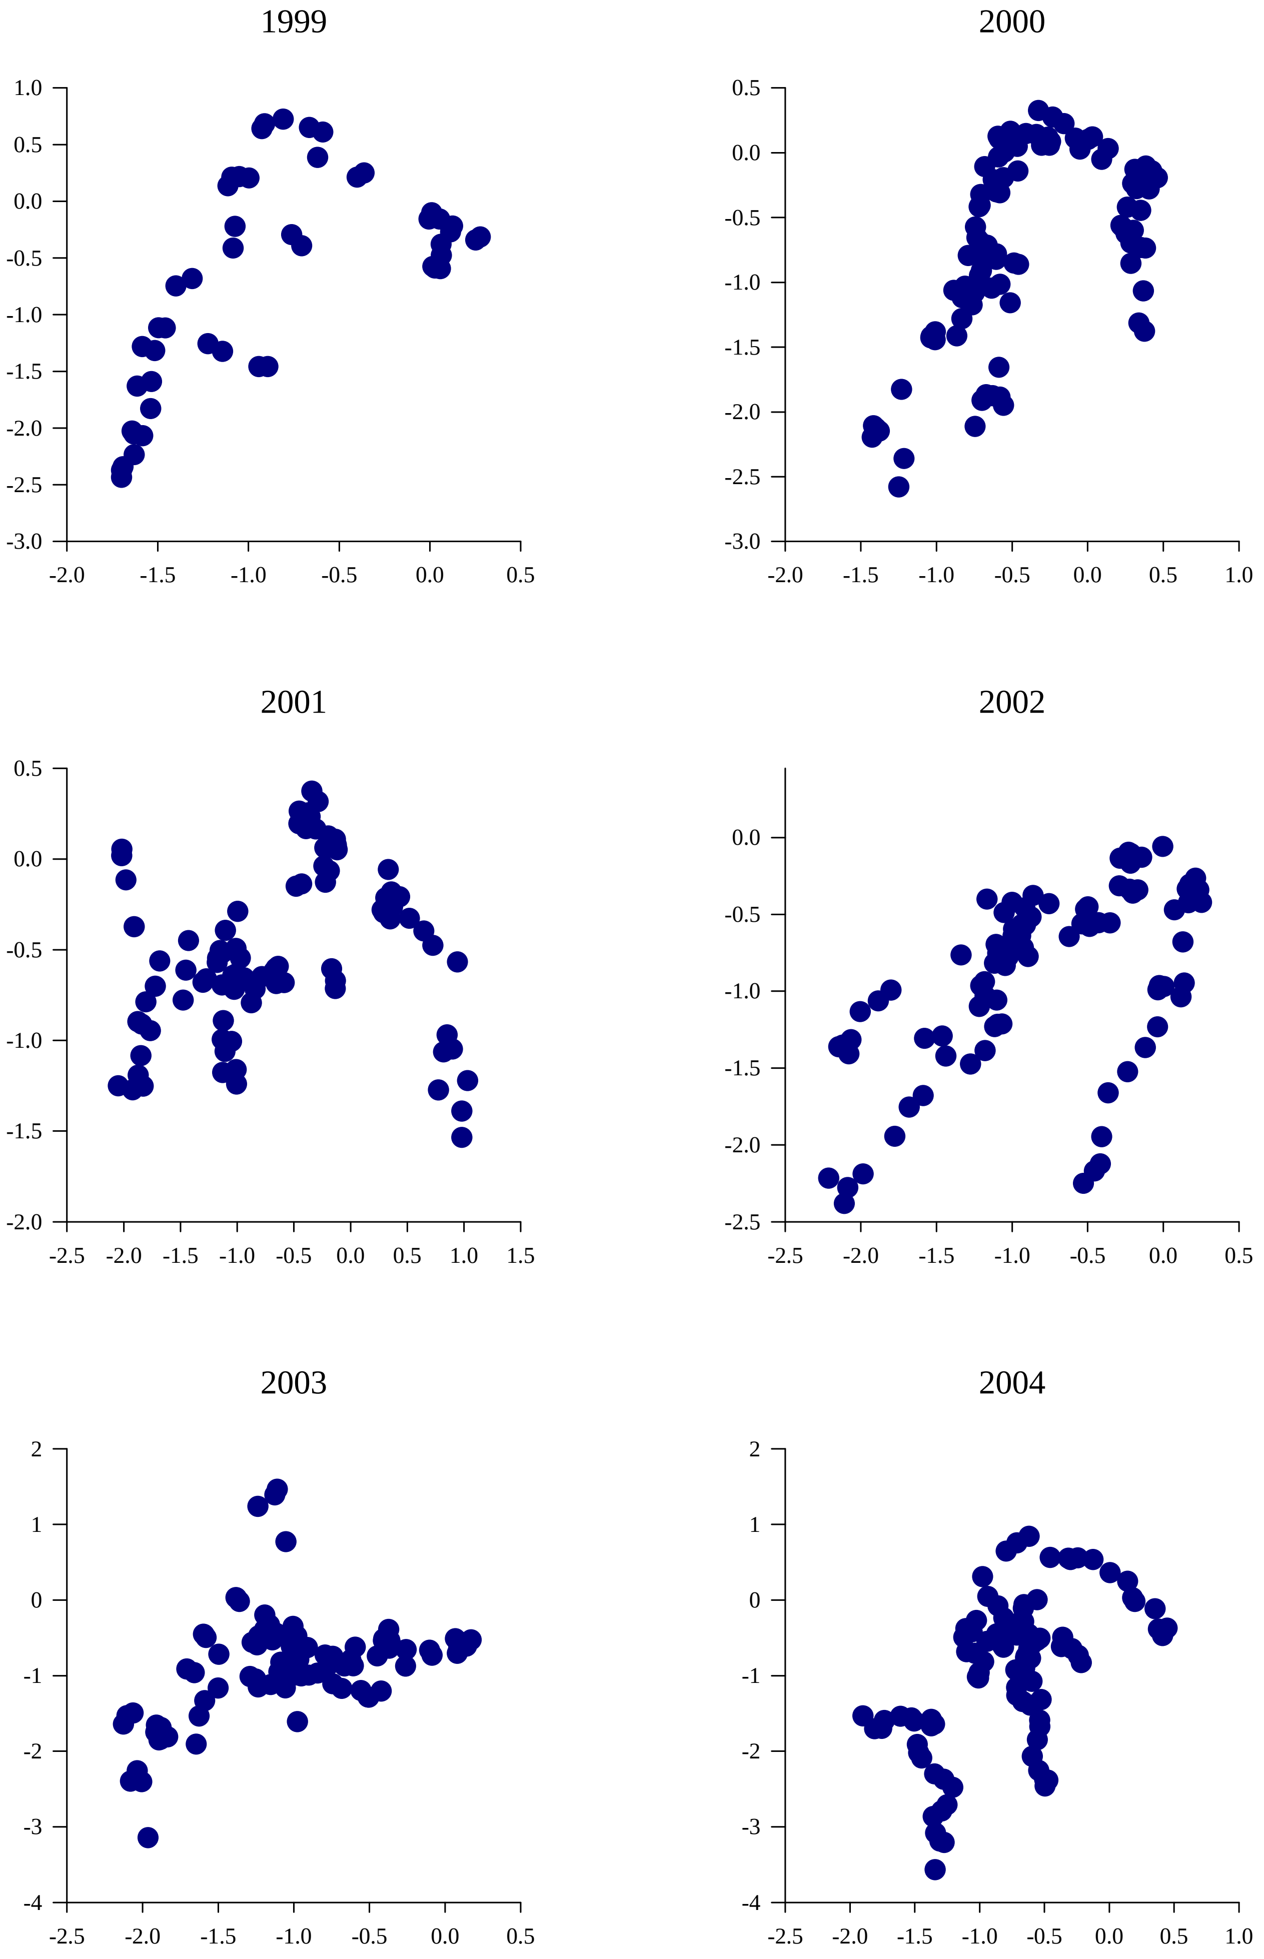

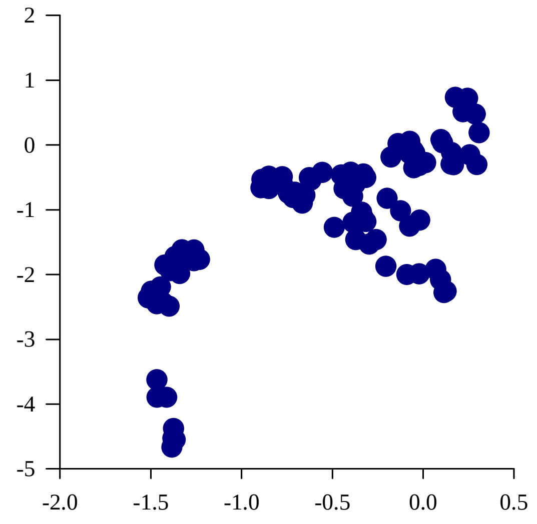
2005

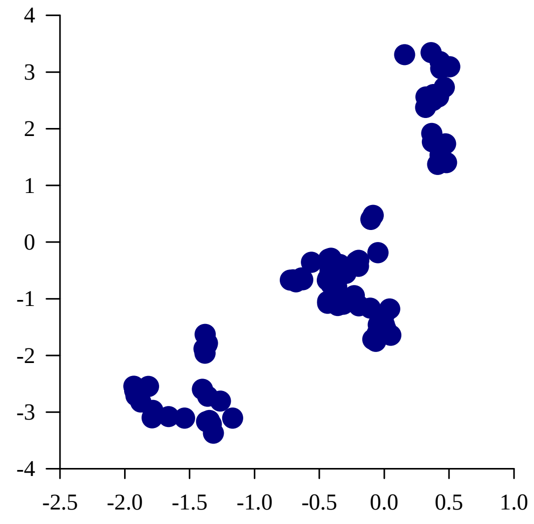
2006

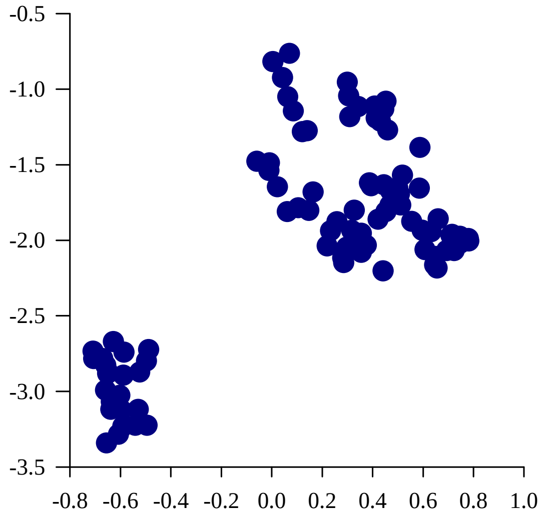
2007

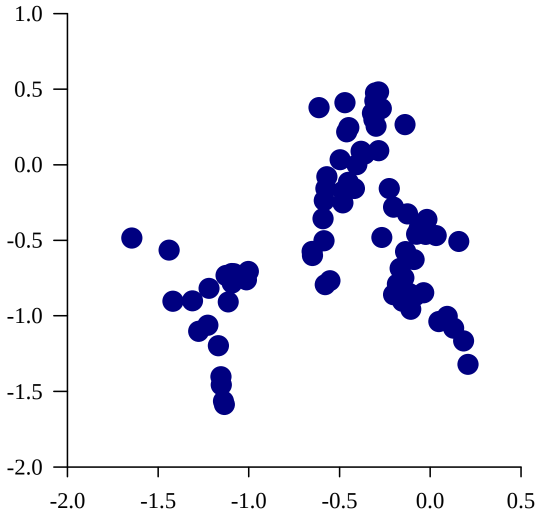
2008

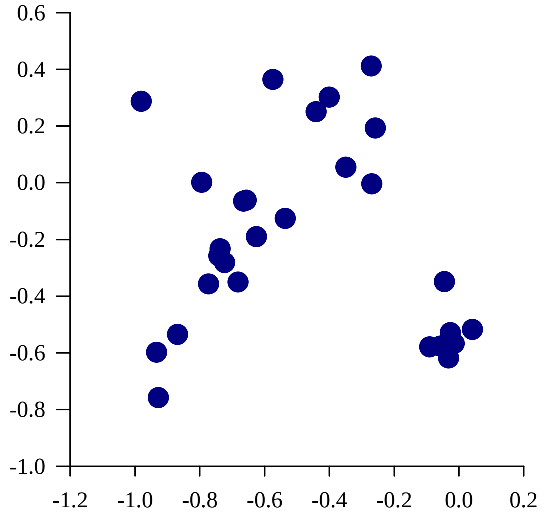
2009

Figure 8: Simulated Role of Hand-to-Mouth Households

Economy-Pair Forecast Means: *n=1*

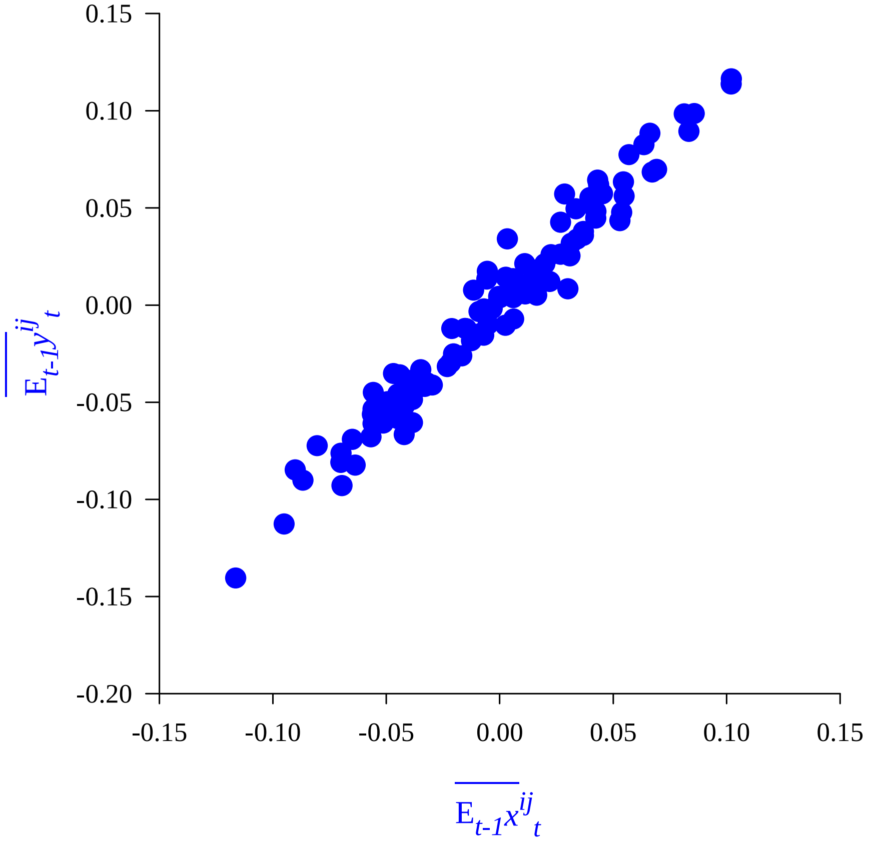

Economy-Pair Forecast Means: *n=0.6*

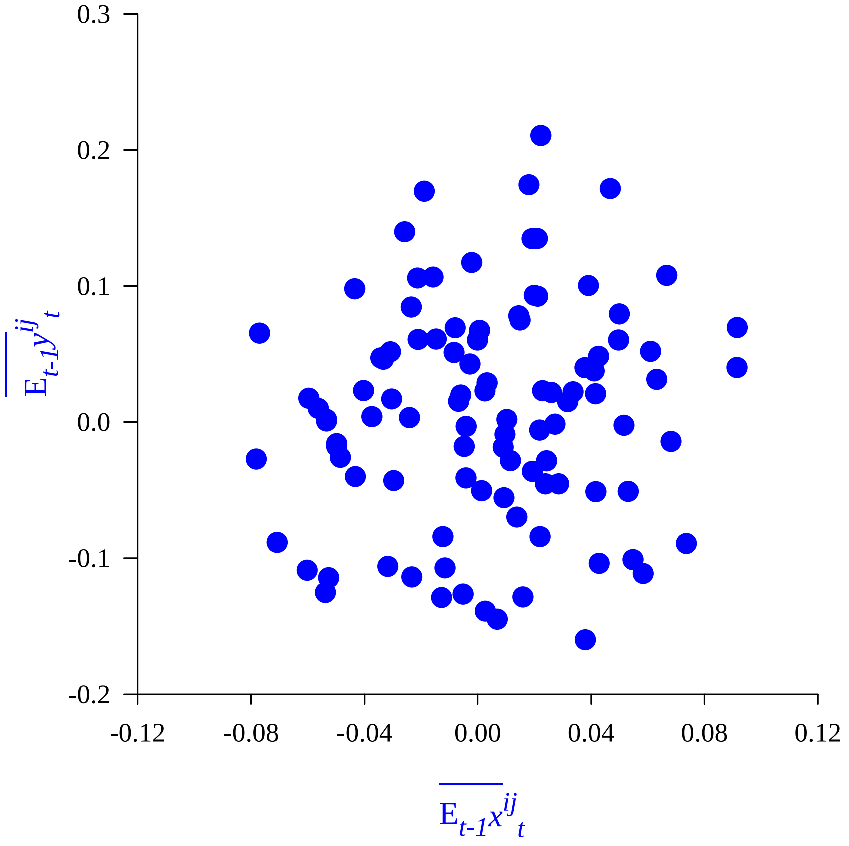

Figure 9: Simulated Role of Hand-to-Mouth Households

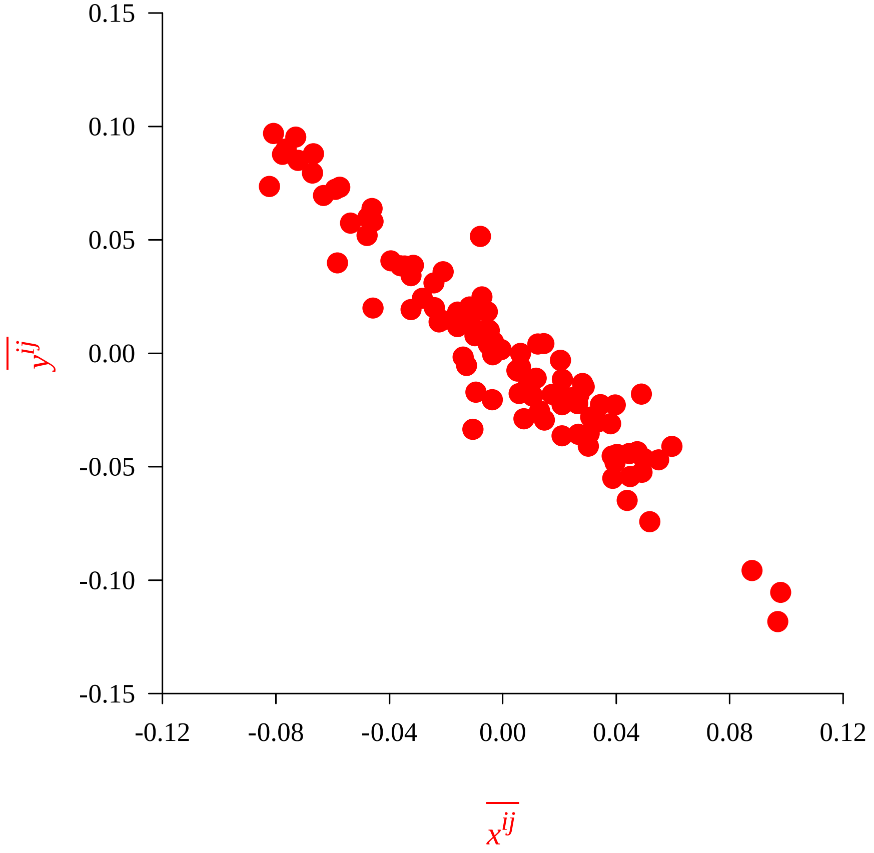

Economy-Pair Means: *n=1*

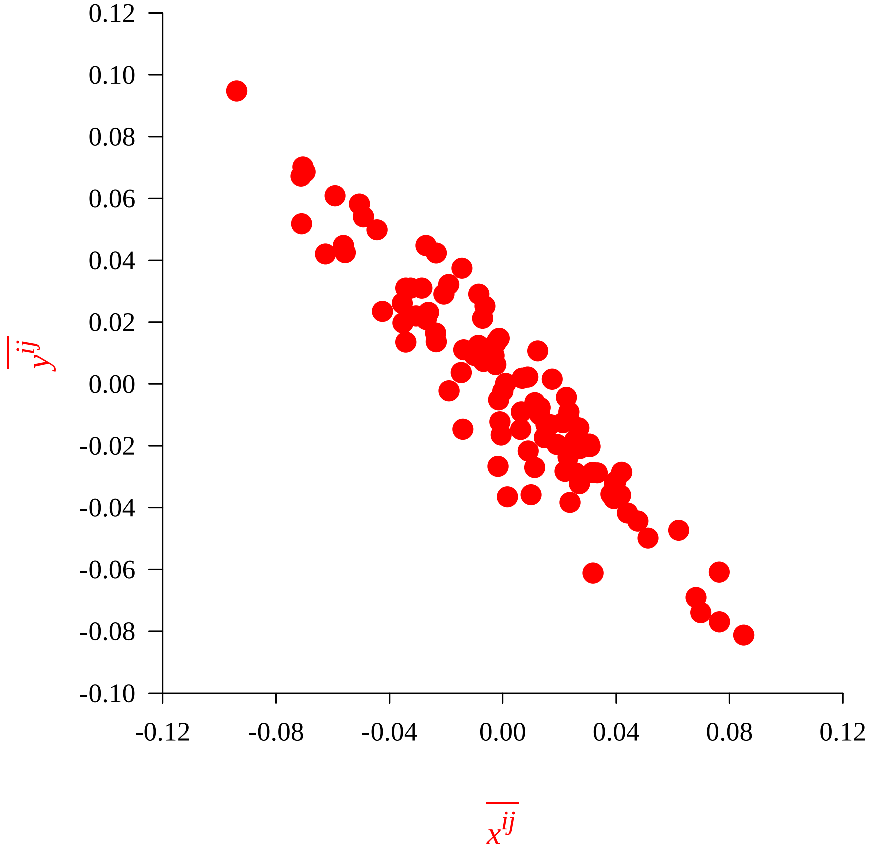

Economy-Pair Means: *n=0.6*