Consumption and Real Exchange Rates in Professional Forecasts

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

October 2008

Abstract

Standard models of international risk sharing with complete asset markets predict a positive association between relative consumption growth and real exchange-rate depreciation 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. Using separate findings from the application of forecasts for consumption and real interest rates, we suggest that the presence of 'hand-to-mouth' consumers may help to explain the anomaly.

JEL classification: F41, F47, F37

Keywords: international risk sharing, Backus-Smith puzzle

†Devereux: Department of Economics, University of British Columbia, `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 author's alone and are not those of the Bank of Canada or the Bank for International Settlements. We thank Rasool Zandvakil for skilled research assistance.

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. As is well known, though, most of the evidence is that there is no such link in historical data, or perhaps even a negative relationship. This rejection of an implication of these models sometimes is called the consumption/real-exchange-rate anomaly or Backus-Smith puzzle.

Recently, though, 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. This weaker 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 the link between expectations of relative consumption growth and those of real depreciations by using professional forecasts. The direct test does not require us to calibrate a specific economic model. And outsourcing the forecasting problem in this way provides us with a real-time series of forecasts that can be used to directly test an implication of the incomplete-markets models.

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.

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 brief summary.

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 let $\pi(z^t)$ be the probability of history z^t. Households in country i have preferences given by

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

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 $r^i(z^t) = s^i(z^t)p^i(z^t)/p^1(z^t)$. We make no special assumptions about real exchange rate determination; $r^i(z^t)$ 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.

Table 1: Notation

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

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.

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 $q(z^{t+1}|z^t)$ 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:

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

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:

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

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

$$\frac{u'(c^i(z^{t+1}))}{u'(c^i(z^t))}\Big/\frac{u'(c^j(z^{t+1}))}{u'(c^j(z^t))} = \frac{r^{ij}(z^t)}{r^{ij}(z^{t+1})} \tag{2}$$

$\forall i, j = 1, ..N$, 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. Use dots to denote growth rates. Taking a linear approximation of the risk-sharing condition (2) around a non-stochastic steady state, we have:

$$\dot{c}^i_{t+1} - \dot{c}^j_{t+1} = \gamma \dot{r}^{ij}_{t+1} \tag{3}$$

$\forall i, j = 1, ..N$, 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 of ex-post growth rates of consumption across countries i and j should be a positive, linear function of ex-post growth rate of the real exchange rate r^{ij}.

Condition (3) 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 (3). 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 (1995) for example. The price of such a bond is $q(z^t) = \sum_{z^{t+1}} q(z^{t+1}|z^t)$. This bond price is determined by the condition:

$$q(z^t) \frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \beta \sum_{z^{t+1}} \pi(z^{t+1}|z^t) \frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})}. \tag{4}$$

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

$$\sum_{z^{t+1}} \pi(z^{t+1}|z^t) \frac{u'(c^i(z^{t+1}))}{s^i(z^{t+1})p^i(z^{t+1})} \Big/ \frac{u'(c^i(z^t))}{s^i(z^t)p^i(z^t)} = \sum_{z^{t+1}} \pi(z^{t+1}|z^t) \frac{u'(c^j(z^{t+1}))}{s^j(z^{t+1})p^j(z^{t+1})} \Big/ \frac{u'(c^j(z^t))}{s^j(z^t)p^j(z^t)}, \tag{5}$$

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

$$\mathrm{E}\big[(\dot{c}^i_{t+1} - \dot{c}^j_{t+1})|z^t\big] = \mathrm{E}\big[\gamma \dot{r}^{ij}_{t+1}|z^t\big]. \tag{6}$$

Under this asset market arrangement, the time-t conditional expectation of the differences in consumption growth rates should equal the conditional expected growth 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 (6) 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.

3. Tests and Previous Evidence

The complete-markets theory predicts a positive relationship between the predicted consumption growth differential and the predicted real depreciation. Let dots denote annual growth rates.

$$(\dot{c}^i_t - \dot{c}^j_t) = \gamma(\dot{s}^{ij}_t + \dot{p}^j_t - \dot{p}^i_t)$$

To keep notation simple, we define the composite variables $y^{ij}_t \equiv \dot{c}^i_t - \dot{c}^j_t$ and $x^{ij}_t \equiv \dot{s}^{ij}_t + \dot{p}^j_t - \dot{p}^i_t$, so the relationship is:

$$y^{ij}_t = \gamma x^{ij}_t, \tag{7}$$

where γ is the elasticity of intertemporal substitution. The scatterplot of the two composite variables is an upward-sloping line, with slope γ. The strongest test of the complete-markets, risk-sharing condition is to inspect the scatterplot 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, 1 in autocorrelation coefficients.

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

$$\overline{y}^{ij} = \sum_{t=0}^{T} y_t^{ij}$$

and

$$\overline{x}^{ij} = \sum_{t=0}^{T} x_t^{ij}$$

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

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

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:

$$\sum_{t=0}^{T} (y_t^{ij} - \overline{y}^{ij})(x_t^{ij} - \overline{x}^{ij}) = \big(\sum_{t=0}^{T} y_t^{ij} x_t^{ij}\big) - \overline{y}^{ij}\overline{x}^{ij}$$

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

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 that ranged over [-0.08,0.17], with an average value of 0.045. Kollmann (1995) ran regressions like condition (3) for country pairs and found they had very low R^2, which provides similarly negative information on the correlations. (Kollmann also found that consumptions 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 discrepency between theory and data as the 'consumption real exchange rate anomaly'.

So the complete-markets condition (3) is rejected. Theoretical approaches to resolving this anomaly have relaxed the assumption of complete markets. Corsetti, Dedola, and Leduc (2008), Beningo and Thoenissen (2005), Selaive and Tuesta (2007), 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 (6). This condition is then combined with a shock process such that a shock that generates a real exchange rate depreciation leads to a rise in relative consumption. Then it is possible to account for the negative consumption real exchange rate correlation found in the time series data. For instance, in Beningo and Thoenissen's model, in an environment of incomplete markets, 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 lie behind the papers by Corsetti, Dedola, and Leduc (2007), Selaive and Tuesta (2007), and Opazo (2006).

While these papers develop structural DSGE models in which condition (3) fails to hold, due to incomplete markets, they still imply that condition (6) 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 (6) makes no clear predictions regarding the correlation between consumption growth rate difference and real exchange rates. Rather, the conditions imply 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 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, forecast errors may be correlated across countries. In a given year, for example, forecasts of consumption growth might be above actual outcomes for several different countries. Again, then, the slopes of the scatter plots would differ between actual outcomes and expectations.

We next assess the appropriateness for testing the incomplete-markets, risk-sharing condition (6) of the monotonicity-in-means test and the covariance test in turn. We do not know z^t. But if we have instruments w_t and are confident that $w_t \in z^t$ then we can test and estimate with the sample versions of:

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

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 (2003) test various versions of this condition. In general these studies find only weak evidence supporting the incomplete markets condition. Obstfeld tests condition (6) (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) note that the incomplete-markets condition (6) also is necessary for the complete-markets condition (3). They study that condition by GMM for 10 OECD countries from 1961 to 2001. They examine 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 consider models with exogenously missing asset markets markets but an endogenous discount rate that anchors the distribution of wealth and with endogenous market segmentation. They 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. [Note: Can we test this unconditionally or in forecasts?]

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 ι, a vector of ones. In that case, the estimating equations for the incomplete-markets condition become:

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

This is simply the monotonicity condition in means used by Backus and Smith (1993) and Obstfeld (2008). 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.

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,

$$\begin{aligned} \mathrm{E}(y_{t+1}^{ij} - \gamma x_{t+1}^{ij}|z^t) = 0 &\Rightarrow \mathrm{cov}((y_{t+1}^{ij} - \gamma x_{t+1}^{ij}|z^t) > 0 \\ &\nRightarrow \mathrm{cov}((y_{t+1}^{ij} - \gamma x_{t+1}^{ij}|w_t) > 0 \end{aligned}$$

[I'll get a formal citation for this, and explain a counter-example with $w_t = \iota$.]

So there is no direct, conditional covariance test for the incomplete-markets condition (6). But models with incomplete markets still can be tested indirectly by seeing whether they can reproduce a negative covariance or correlation in actual outcomes, treating that like any stylized fact. For example, Bodenstein (2008) studies a model in which asset markets are endogenously incomplete due to limited enforcement. He shows using numerical examples that it can reproduce a negative correlation.

Our method in this paper is to test the risk-sharing condition using professional forecasts. 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:

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

The forecasts lie on an upward-sloping line, with slope γ. But if the mean test continues to apply to the incomplete-markets condition, and rejects, then why are we using forecast data? There are four reasons:

1. using these conditional forecasts should add precision and so allow us to estimate γ and test its sign;
2. using these real-time forecasts may avoid both the syndrome of weak instruments and the pitfall of using revised data as instruments;
3. the test using unconditional means will not be valid if there is non-stationarity, in the form of a break in unconditional moments; [I'll compose an example or diagram where the test would be misleading.] and
4. these data allow us to see if there are trends over time in the evidence for risk sharing, because we can find conditional means in each year.

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 CCAPM links between 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 thus are 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 1 lists the 29 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
Eurozone	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 relative to one another 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, Germany, Italy, the Netherlands, and Spain. The forecasts thus are 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), 10 economy pairs, and 115 months (from January 1999 to July 2008) for a total of 2300 observations. We also consider a start date of January 2002, when the new currency was introduced, leaving 1580 observations.

Our reporting method is simple, as noted in section 3. 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. [Check notation since earlier sections used $t+1$.] We plot these 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 scatterplots, 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. [Need to relate this to discussion in section 3 of various tests and reason for using forecasts.]

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 [check this in forecasts].

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 *intranational* data from US states, and find the correlation there (where the nominal exchange rate is one) is positive. In a similar vein, Hadzi-Vaskov (2008) studies quarterly data for the Eurozone 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.

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 these 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.

6. A Simple Interpretation

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 expectation of consumption growth differences and real exchange rate changes, where conditional expectations can be represented by forecasts. While the first failure (complete markets) 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 (such as Benigno and Theonissen 2006, or Corsetti et al. 2007), the second is more problematic. Most DSGE international macro models with any form of capital mobility imply a version of condition (6). But condition (6) seems to be rejected in most of the above figures.

One perspective on this is given by Engel and Rogers (2008). We may decompose (6) as follows:

$$\begin{aligned} &\mathrm{E}\big[(\dot{c}^i_{t+1}|z^t\big] - \mathrm{E}\big[(\dot{c}^j_{t+1})|z^t\big] \\ &= i^i_{t+1} - \mathrm{E}\big[(\pi^i_{t+1}|z^t\big] - \big(i^j_{t+1} - \mathrm{E}\big[(\pi^j_{t+1}|z^t\big]\big) \\ &= \mathrm{E}\big[\gamma\dot{r}^{ij}_{t+1}|z^t\big]. \end{aligned} \tag{?}$$

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 (6) 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 on 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.

Here we suggest a simple framework which may throw some insight on the failure of the consumption-real exchange rate relationship. 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. Take the following illustrative framework (we spell this out more completely in the appendix). There are two countries, with stochastic endowments of both traded and non-traded goods in each country. The countries are of equal size, with household preferences defined over consumption of home and foreign goods traded goods. The home and foreign country price indices for the traded good are defined as:

$$P_{Tt} = \left[\mu P_{ht}^{(1-\theta)} + (1-\mu) P_{ft}^{*(1-\theta)}\right]^{\frac{1}{(1-\theta)}}$$

$$P_{Tt}^{*} = \left[(1-\mu) P_{ht}^{(1-\theta)} + \mu P_{ft}^{*(1-\theta)}\right]^{\frac{1}{(1-\theta)}},$$

where $\mu \geq 0.5$ indicates the potential for home bias in tastes over the local good, and $\theta > 0$ measures the elasticity of substitution across goods.

Allowing for the consumption of both traded and non-traded goods, the overall CPI for each country is defined as

$$P_t = \left[\alpha P_{Nt}^{(1-\epsilon)} + (1-\alpha) P_{Tt}^{(1-\epsilon)}\right]^{\frac{1}{(1-\epsilon)}}$$

$$P_t^* = \left[\alpha P_{Nt}^{*(1-\epsilon)} + (1-\alpha) P_{Tt}^{*(1-\epsilon)}\right]^{\frac{1}{(1-\epsilon)}}$$

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 good.

Each household receives a random endowment of its own country's traded good (e.g. Y_{ht} in the case of home) and the non-traded good. 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. The other $1-n$ households have to consume the value of their income (constrained households). Campbell and Mankiw (1990) estimate $1-n$ as high as 0.5 for the US economy.

Total consumption of home households is written as $C = nC_{ut} + (1-n)C_{ct}$, where C_{ut} (C_{ct}) represents consumption of unconstrained (constrained) households. The budget constraints for unconstrained and constrained households, respectively, are given as:

$$P_t C_{ut} + B_{t+1} = P_{Nt} Y_{Nt} + P_{ht} Y_{ht} + (1 + r_t) B_t \tag{8}$$

$$P_t C_{ct} = P_{Nt} Y_{Nt} + P_{ht} Y_{ht} \tag{9}$$

All home households consume the home 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 a condition equivalent to (6).

Constrained households, by contrast, have consumption growth limited by income growth. Taking a linear approximation of (9), and the analogous equation for the foreign country, using lower case letters as log deviations from a non-stochastic steady state, we can represent the expected consumption growth difference between constrained households in the home and foreign economy as:

$$\mathrm{E}\left[(\dot{c}_{ct+1} - \dot{c}^*_{ct+1})|z^t\right]$$

$$= \gamma \mathrm{E}\big[\alpha(\dot{y}_{ht+1} - \dot{y}_{ft+1}) + (1-\alpha)(\dot{y}_{Nt+1} - \dot{y}^*_{Nt+1}) - 2(1-\mu)(1-\alpha)(\dot{p}^*_{ft+1} - \dot{p}_{ht+1})|z^t\big].$$

Comparing the constrained households across countries, the relationship between relative consumption growth and real exchange rates changes will depend on the strength of shocks to traded goods versus non-traded goods, and the movement in the terms of trade. When real exchange rates are driven primarily by persistent shocks to traded goods, for given terms of trade, then *expected* consumption growth differences will tend to move in the opposite direction to expected real exchange rate changes. This is because expected growth in traded good endowments will generate an expected real exchange rate appreciation. But in fact, the terms of trade will not be constant. An expected growth in output of traded goods will be associated with an expected fall in the terms of trade. For $\theta > 1$, this will not be enough to offset the direct wealth effect of expected growth in traded good endowments, and in net, expected consumption growth rises, while the expected real exchange rate falls.

On other 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. 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.

Let us do a simple quantitative exercise that can be compared to the previous data findings. We follow Benigno and Theonissen's (2006) calibration closely. Take $\theta = 2$, $\epsilon = 0.44$, indicating a higher elasticity of substitution between home and foreign goods than between traded and non-traded goods. Let $\mu = .72$, and assume $\gamma = 2$. Take the size of the non-traded goods sector to be 0.55. Assume that log endowments in the traded goods sector follow persistent AR(1) processes, with persistence 0.82, and an innovation standard deviation of 1.9 percent. In the non-traded goods sector, endowments also follow an AR(1) process, 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 however.

Figure 8 shows the model results for the simulated series of consumption growth and real exchange rates. The first two panels of the top row of the figure show simulations for $n = 1$, so that all households in both countries can access international bond markets. The top left hand panel, Figure 8.1, shows the scatter plot of average conditional expected growth rate difference of consumption compared to conditional expected growth in the real exchange rate. As implied by (6), there is a clear positive relationship here. The top right hand panel, Figure 8.2 shows the average ex-post growth rate by simulation. Here we find a downward sloping relationship. Clearly, Figure 8.1 is the appropriate test. It shows that the model displays conditional risk sharing as facilitated by bond markets, despite the fact that the correlation between consumption growth and real exchange rates may be negative.

The bottom two panels of Figure 8 introduce hand-to-mouth consumers into the model. We set $n = 0.5$. 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 forecast data. Of course, for different parameters and shocks, the results would be different. If $n = 0$, so that all households are hand-to-mouth, the Figure 8.3 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 exchange rates.

Again, it is important to focus on the average of conditional expectations, rather than on correlations. Figure 8.4 shows that for the model with hand to mouth consumers, the correlation between consumption and real exchange rates is negative. But again, this is not the appropriate test.

In summary then, the absence of a clear relationship between consumption and real exchange rates in forecasts seems to be consistent with a general equilibrium setting where international capital markets are restricted to a subset of households. Here, we don't offer 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 the syndrome of weak identification associated with instrumental-variables estimation. Compared to using unconditional means for economy pairs it provides many more observations and allows us to draw inferences even with non-stationarity in unconditional moments. This method thus can test for trends 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.

Appendix. Forecast Data

The data come from 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, when these two measures are completely 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 it 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 nontraded goods. *Journal of International Economics* 35, 297-316.

Benigno, Gianluca and Christoph Thoenissen (2006) Consumption and real exchange rates with incomplete markets and nontraded goods. Manuscript, LSE and University of St. Andrews (February), forthcoming in the *Journal of International Money and Finance.*

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

Campbell, John and N. Gregory Mankiw (1990), Permanent Income, Current Income, and Consumption, *Journal of Business & Economic Statistics* 8, 265-79.

Chari, V. V., Patrick 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, forthcoming in *IMF Staff Papers.*

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. mimeo, Utrecht School of Economics.

Head, Allen, Todd 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 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 (1989) How integrated are world capital markets? Some new tests," in Guillermo A. Calvo et al., eds., *Debt, Stabilization and Development : Essays in Memory of Carlos Daz-Alejandro.* Oxford: Blackwell.

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. Bank of Chile working paper 395.

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

Smith, Gregor W. and Yetman, James (2007) The curse of Irving Fisher (professional forecasters' version). Working paper 1144, Department of Economics, Queen's University.

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

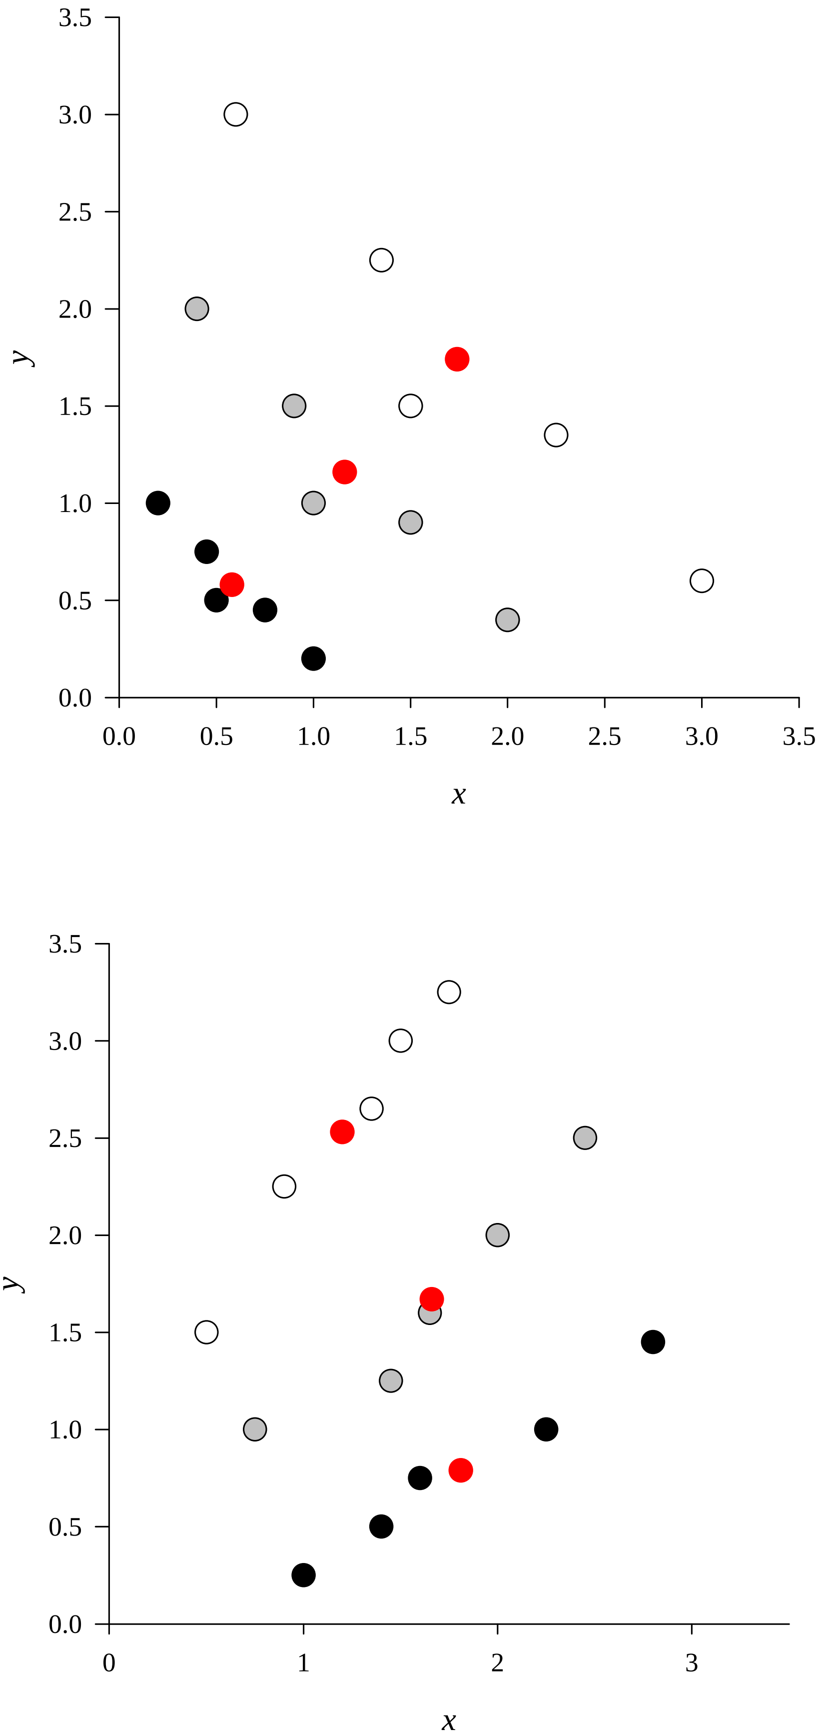

Figure 2: All Economies 1990-2008

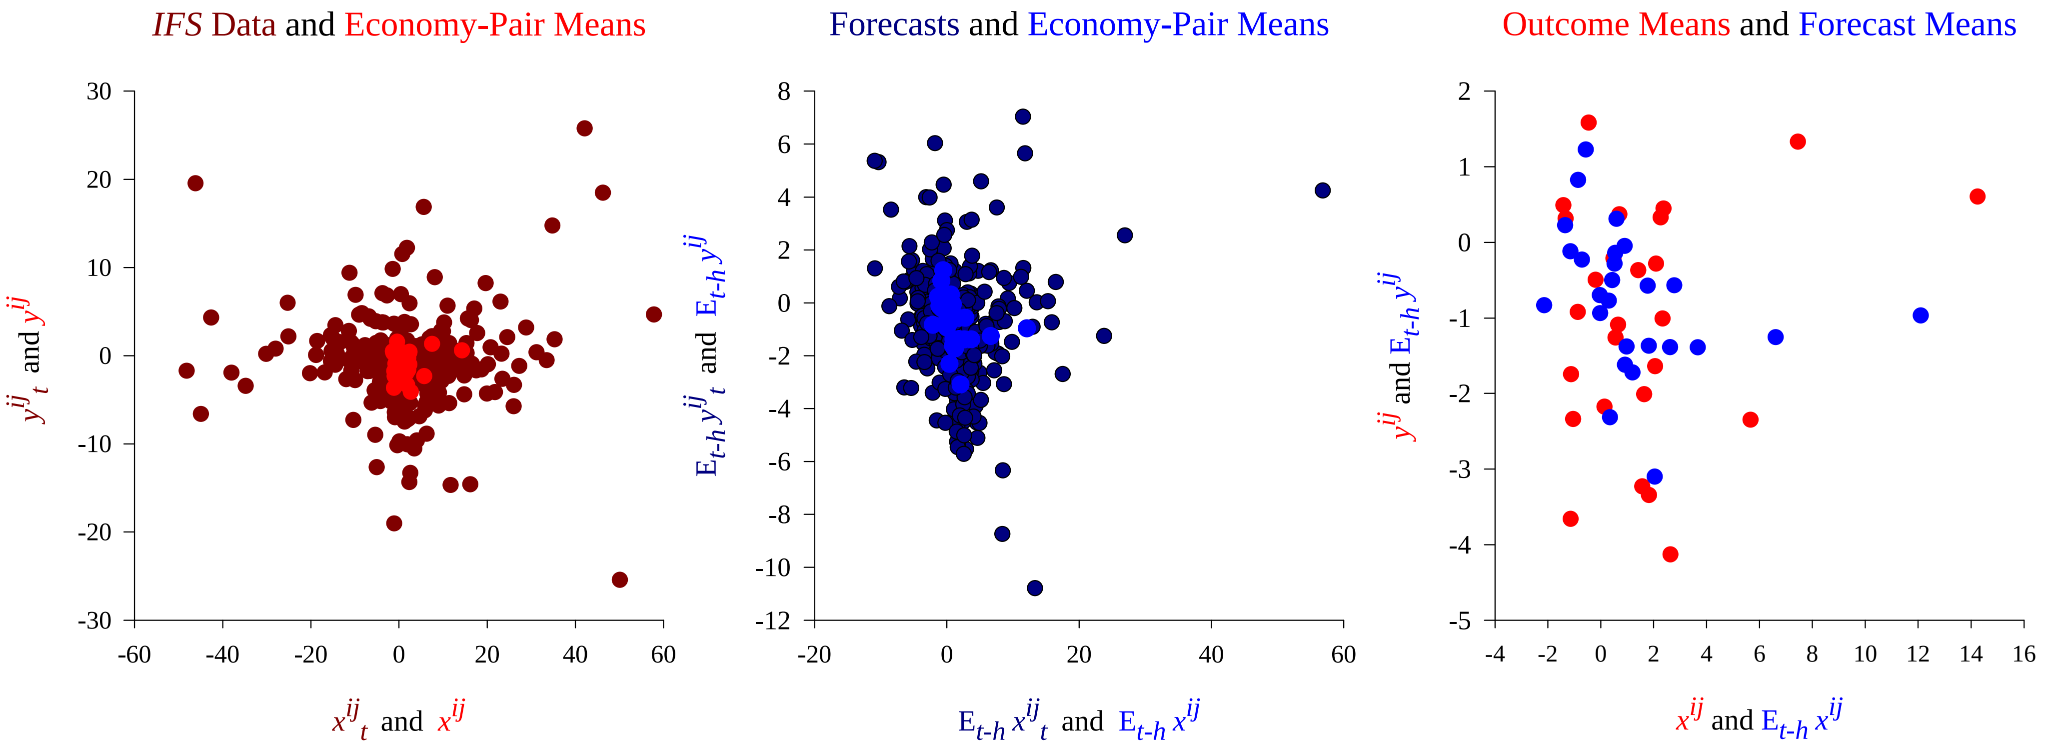

Figure 3: IFS and Forecast Data for Economy Pairs

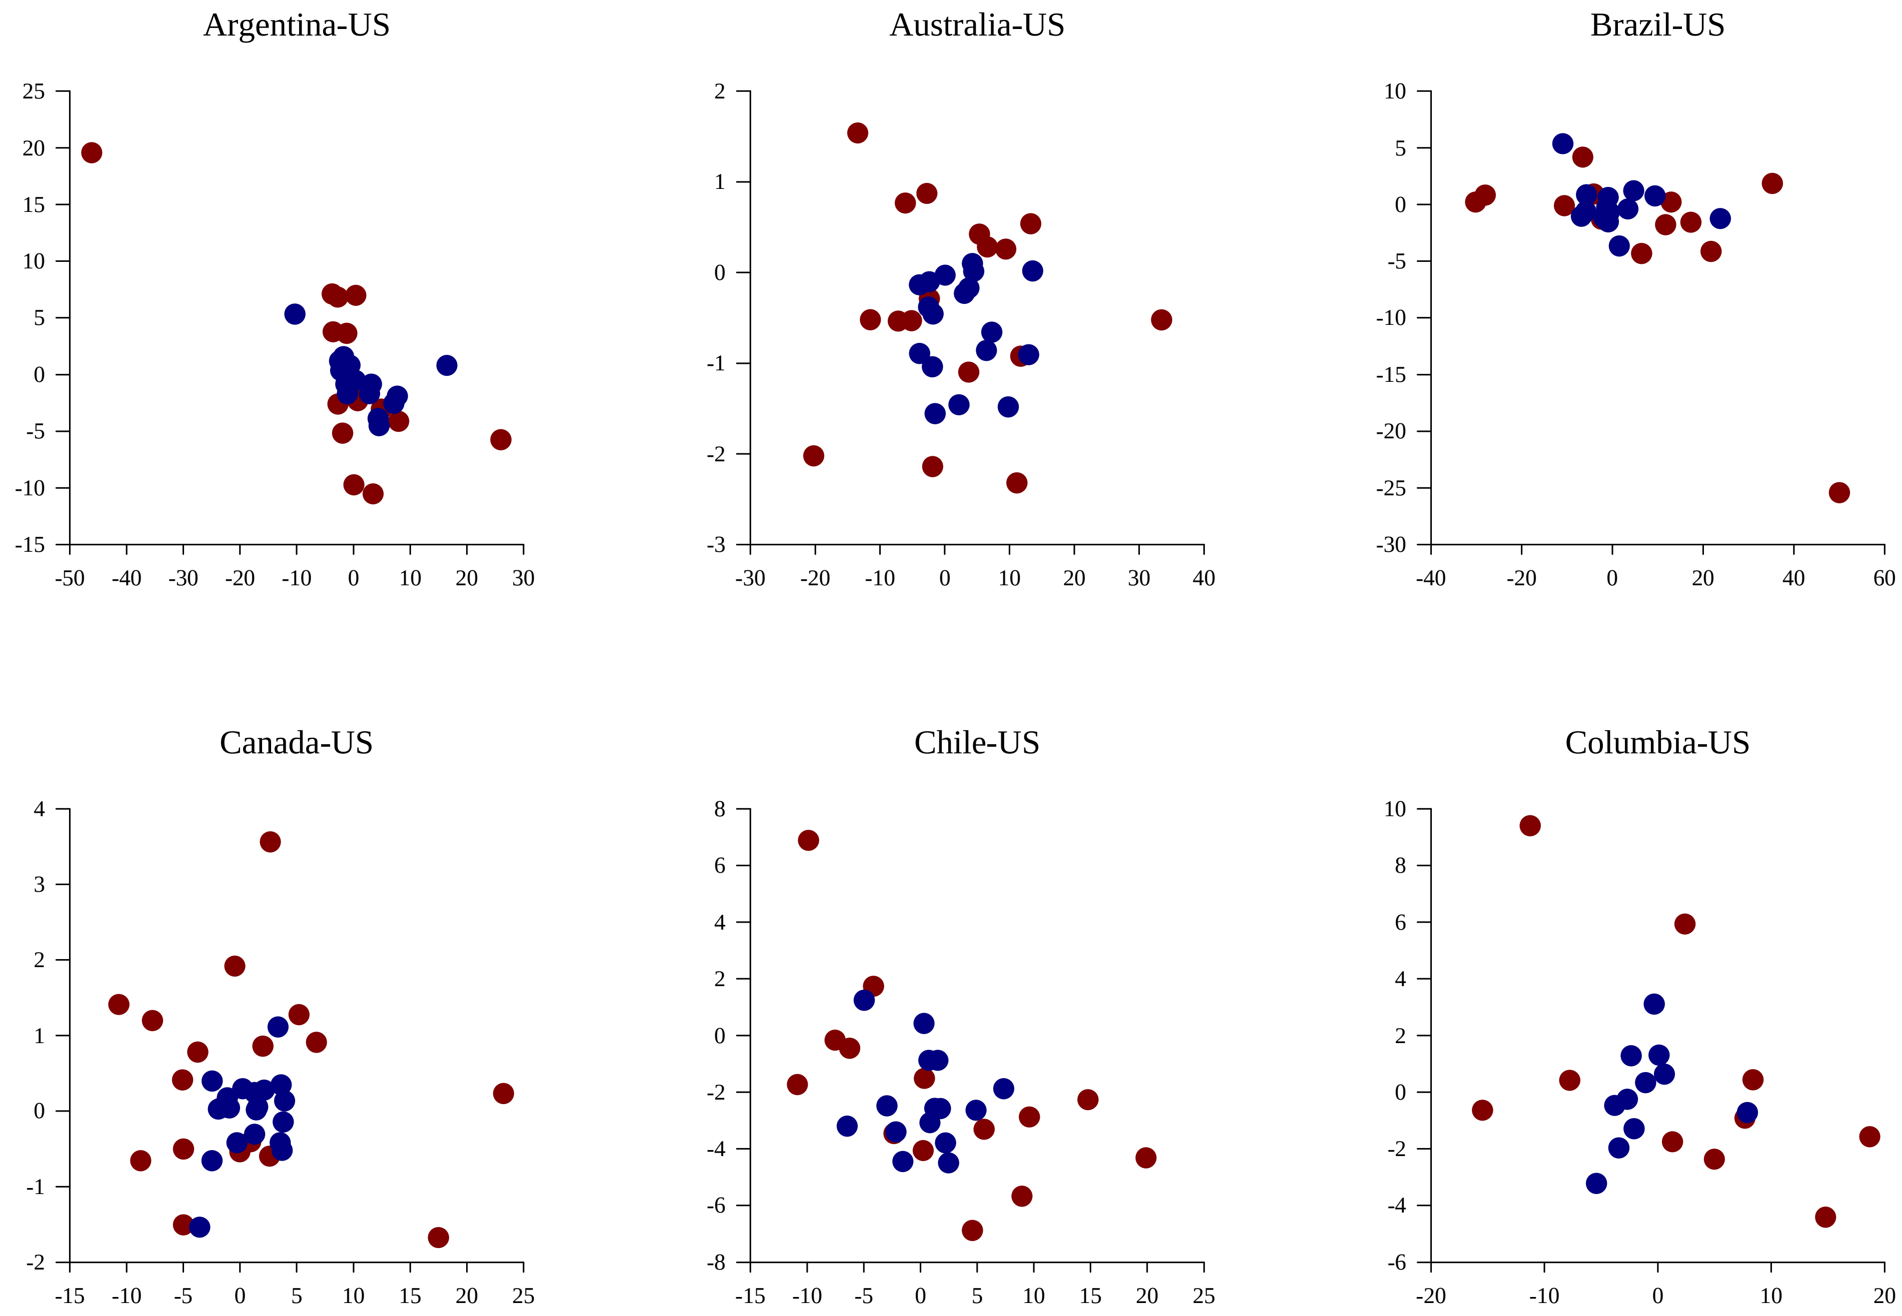

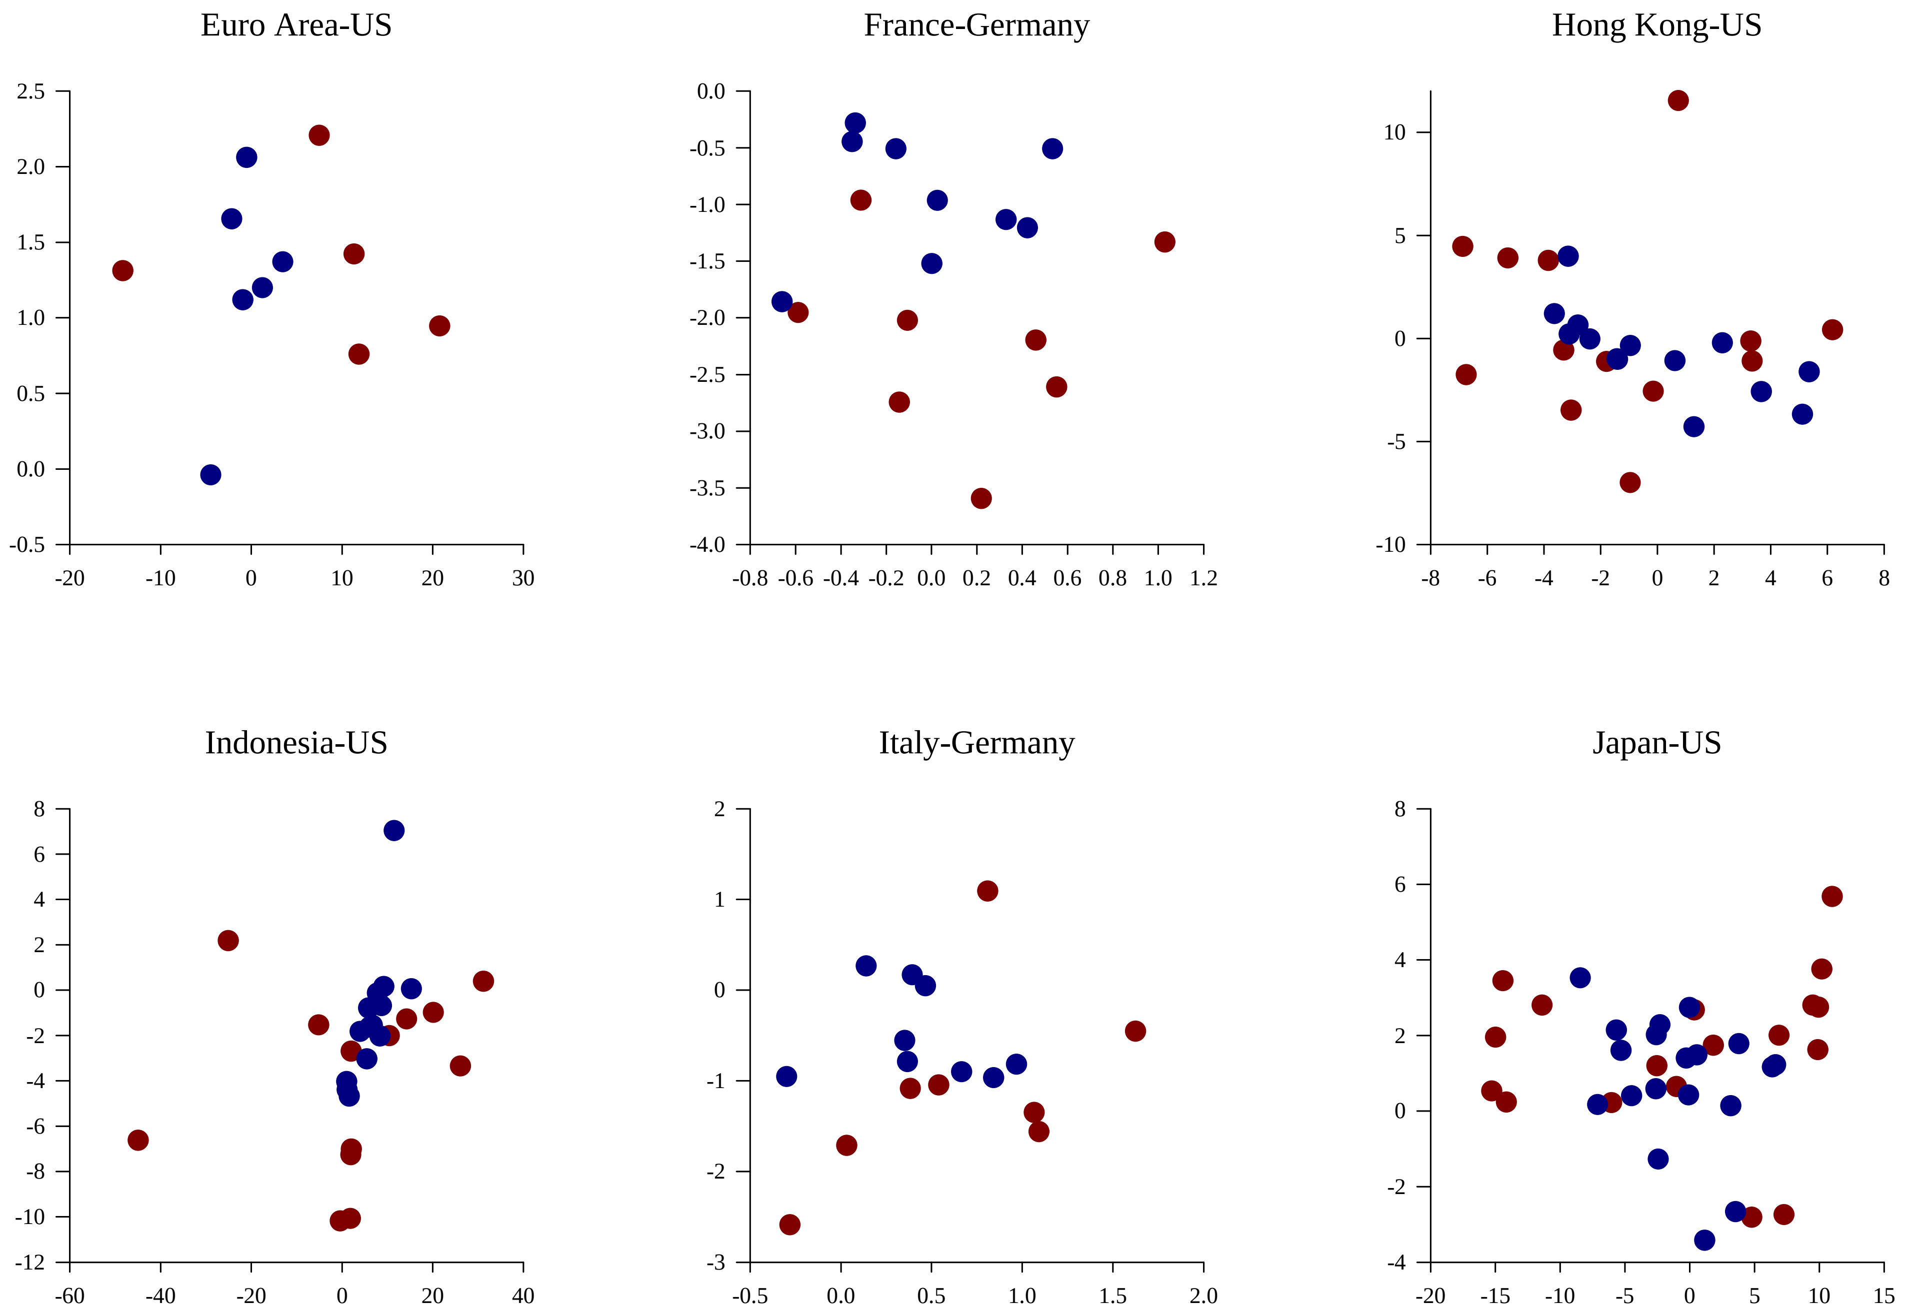
Euro Area-US
France-Germany
Hong Kong-US
Indonesia-US
Italy-Germany
Japan-US

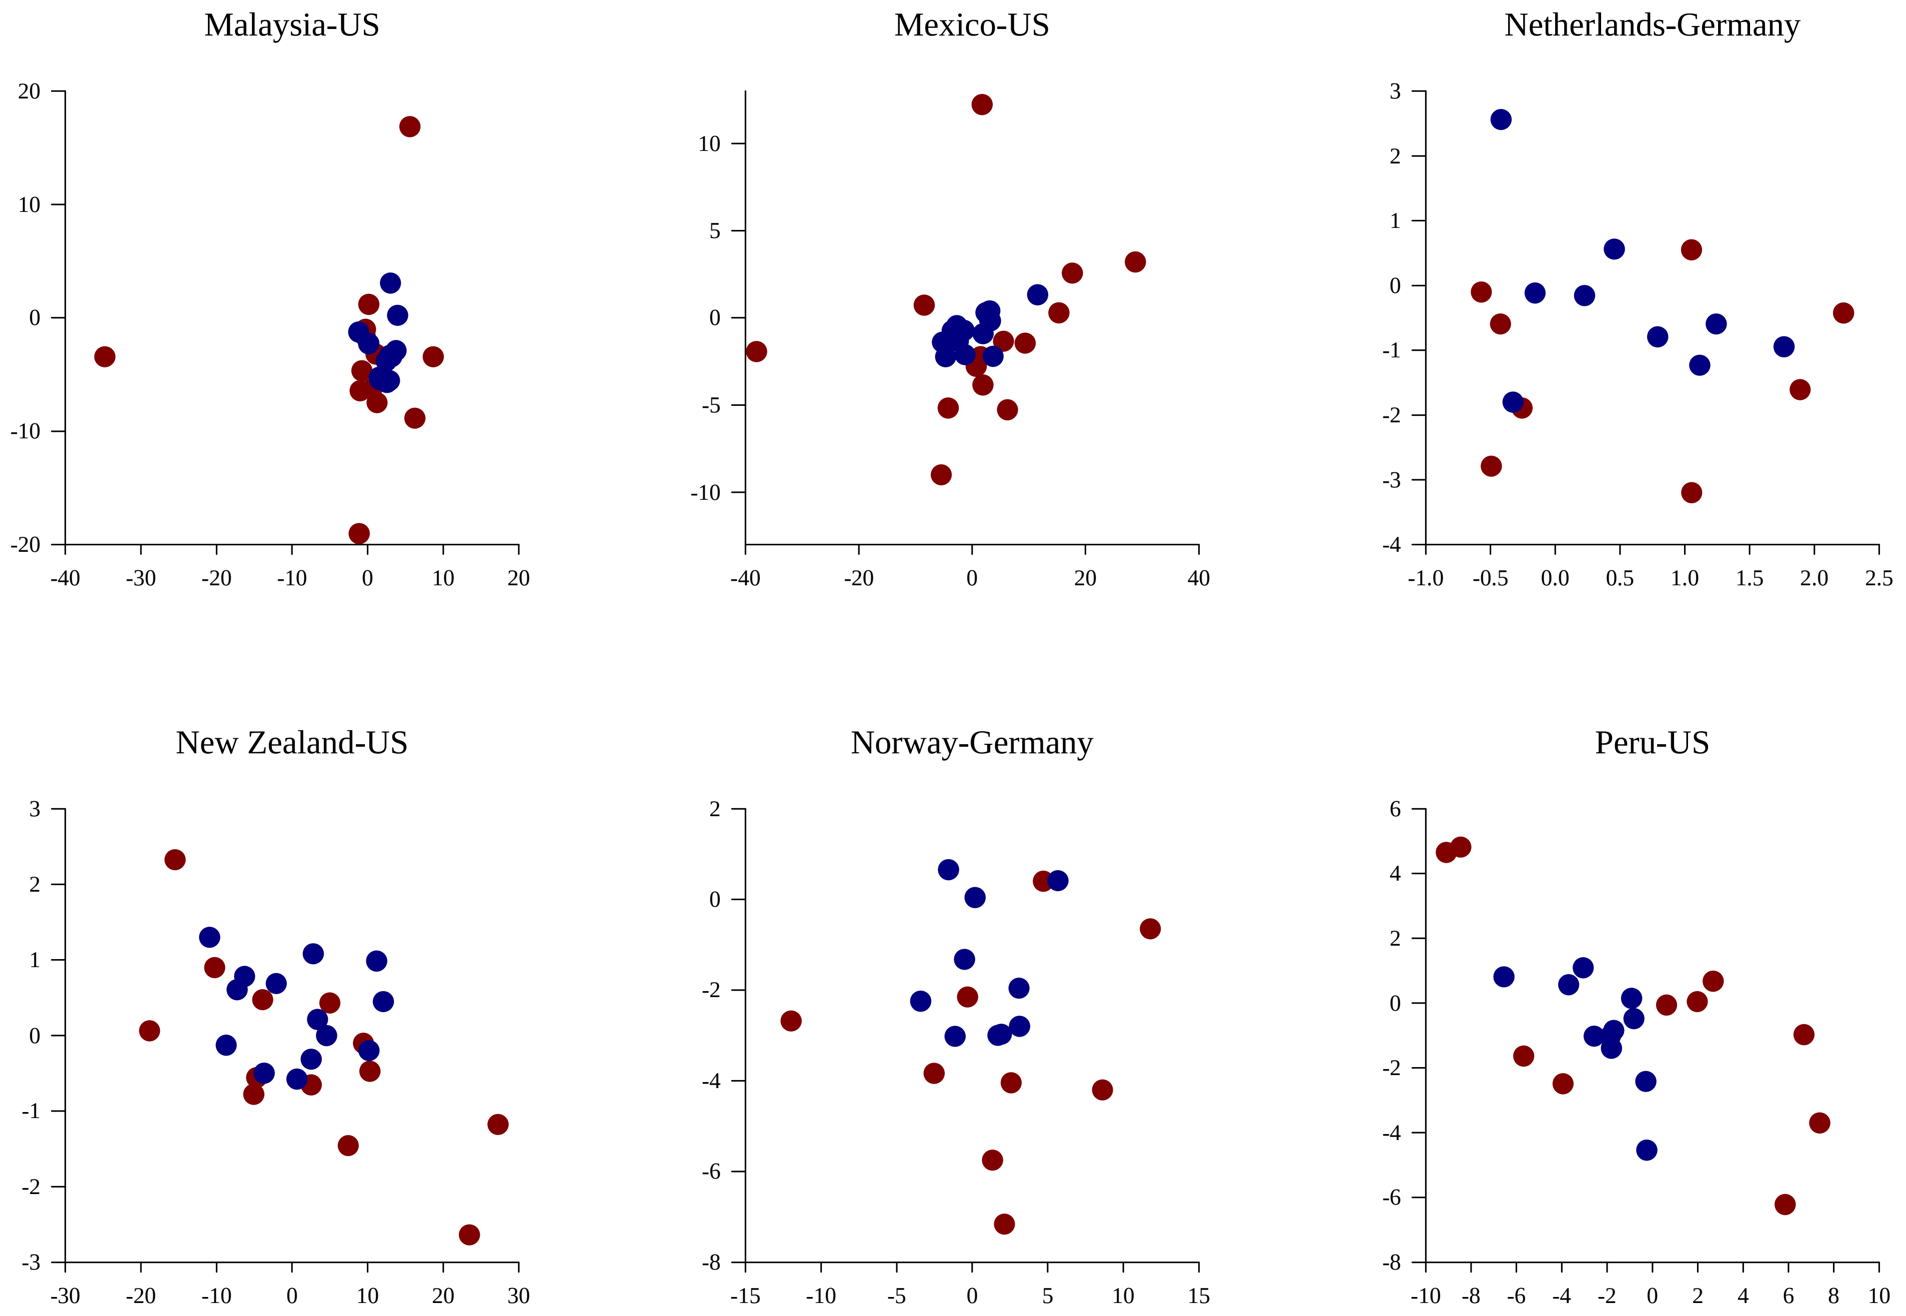
Malaysia-US
20 10 0 -10 -20
-40 -30 -20 -10 0 10 20
Mexico-US
10 5 0 -5 -10
-40 -20 0 20 40
Netherlands-Germany
3 2 1 0 -1 -2 -3 -4
-1.0 -0.5 0.0 0.5 1.0 1.5 2.0 2.5
New Zealand-US
3 2 1 0 -1 -2 -3
-30 -20 -10 0 10 20 30
Norway-Germany
2 0 -2 -4 -6 -8
-15 -10 -5 0 5 10 15
Peru-US
6 4 2 0 -2 -4 -6 -8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

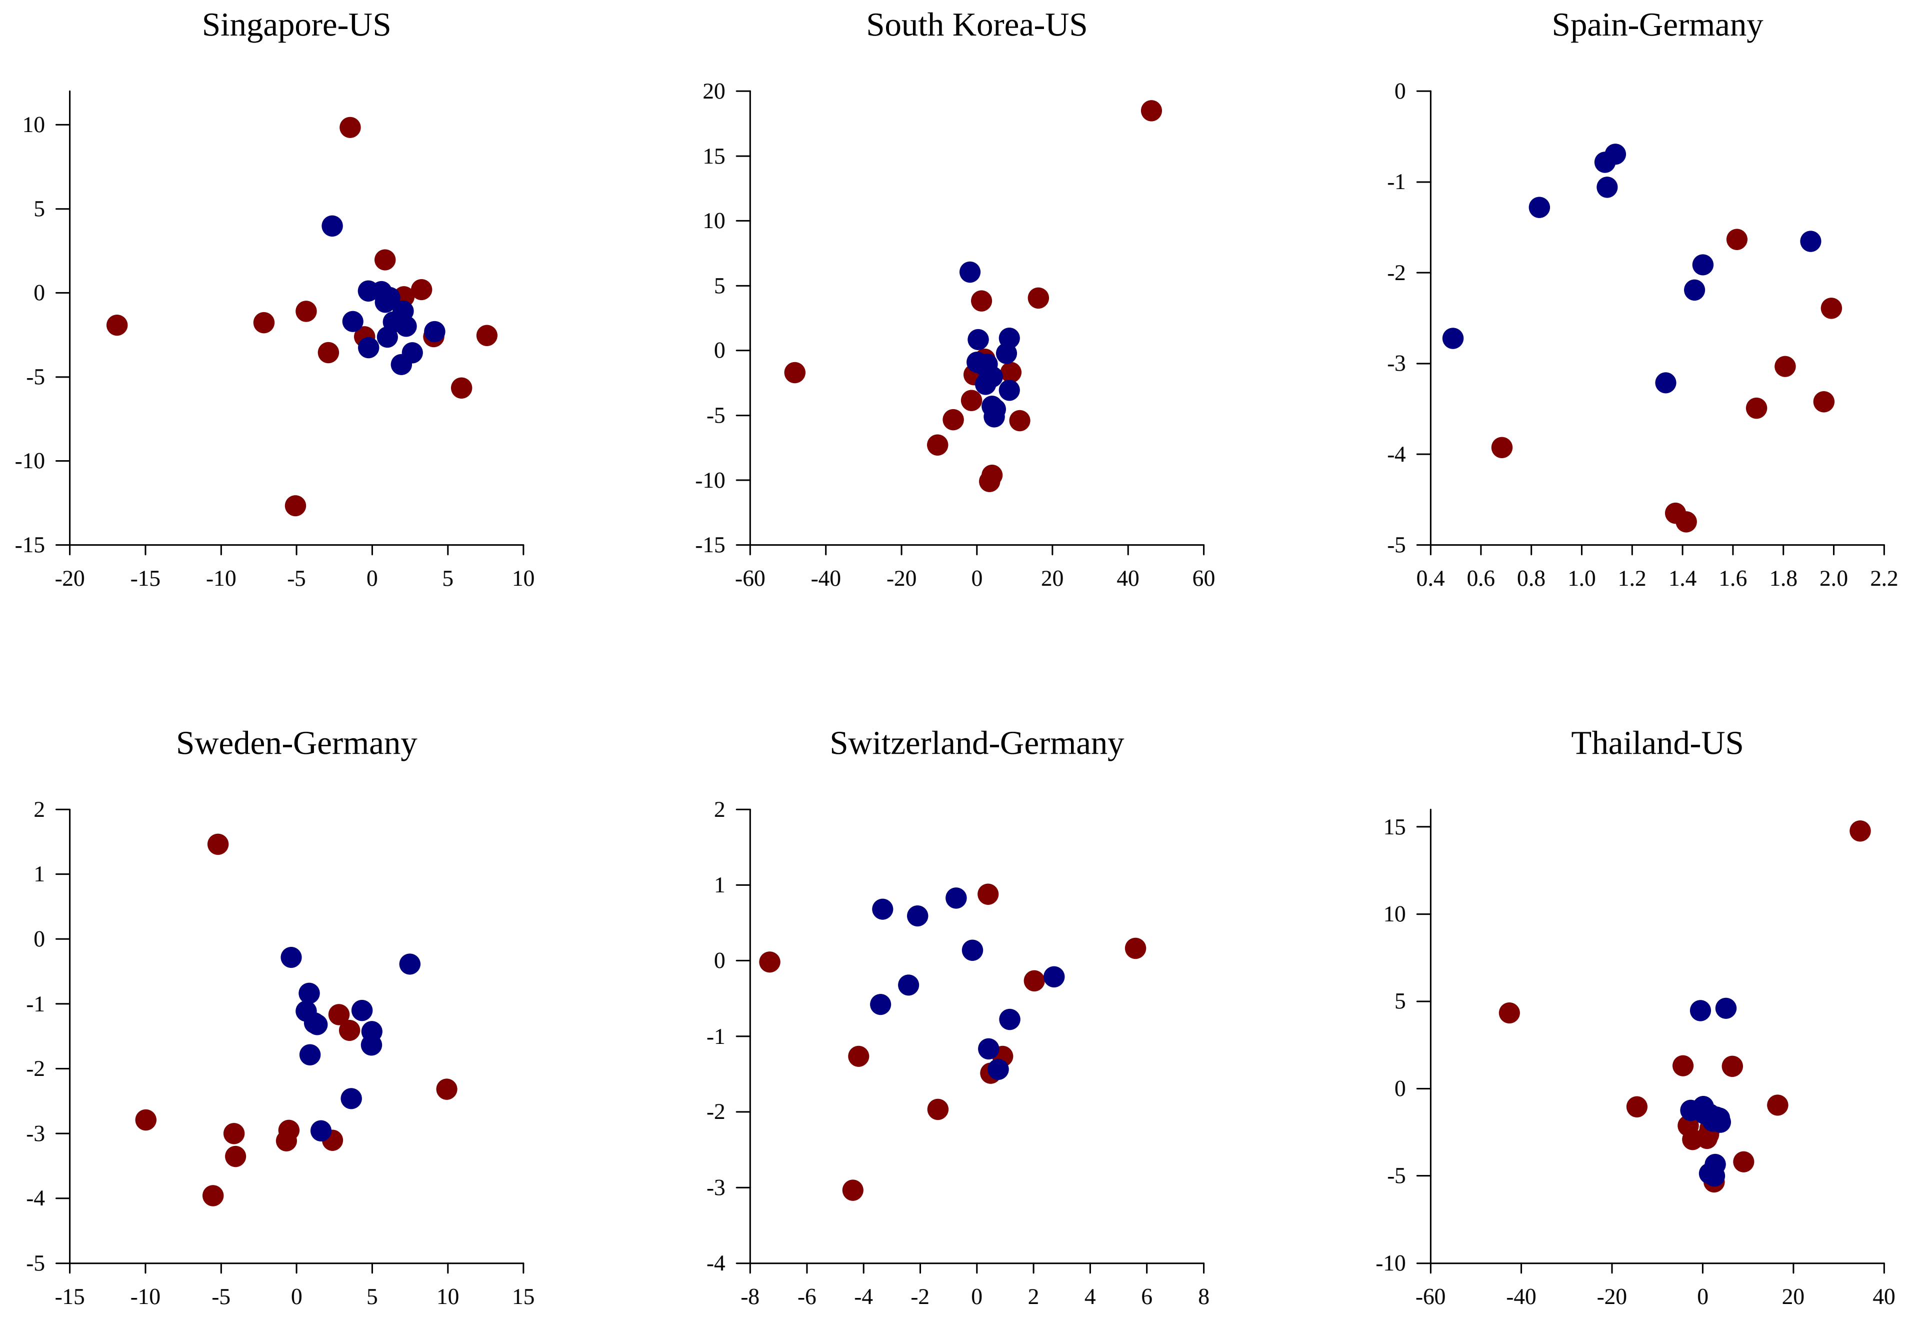
Singapore-US
10
5
0
-5
-10
-15
-20 -15 -10 -5 0 5 10
South Korea-US
20
15
10
5
0
-5
-10
-15
-60 -40 -20 0 20 40 60
Spain-Germany
0
-1
-2
-3
-4
-5
0.4 0.6 0.8 1.0 1.2 1.4 1.6 1.8 2.0 2.2
Sweden-Germany
2
1
0
-1
-2
-3
-4
-5
-15 -10 -5 0 5 10 15
Switzerland-Germany
2
1
0
-1
-2
-3
-4
-8 -6 -4 -2 0 2 4 6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

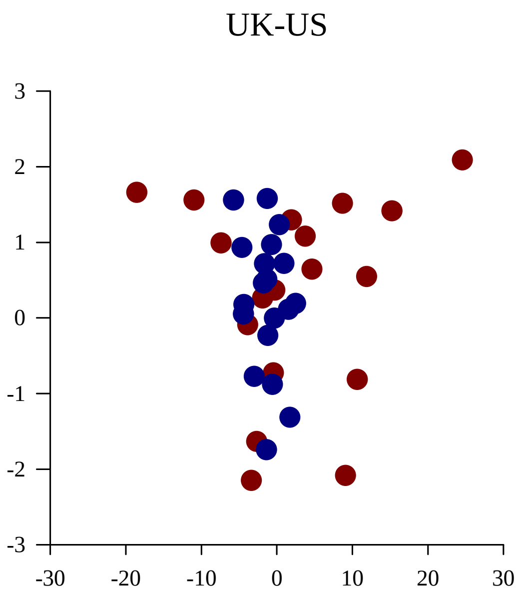
UK-US
3
2
1
0
-1
-2
-3
-30
-20
-10
0
10
20
30

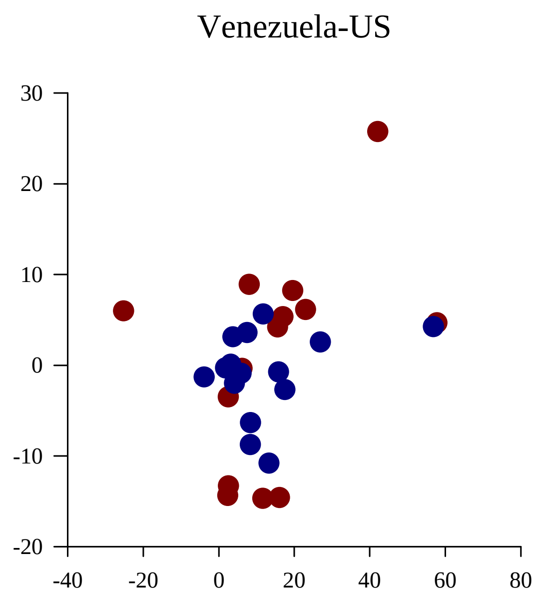
Venezuela-US
30
20
10
0
-10
-20
-40
-20
0
20
40
60
80

Figure 4: Currency-Union Economies

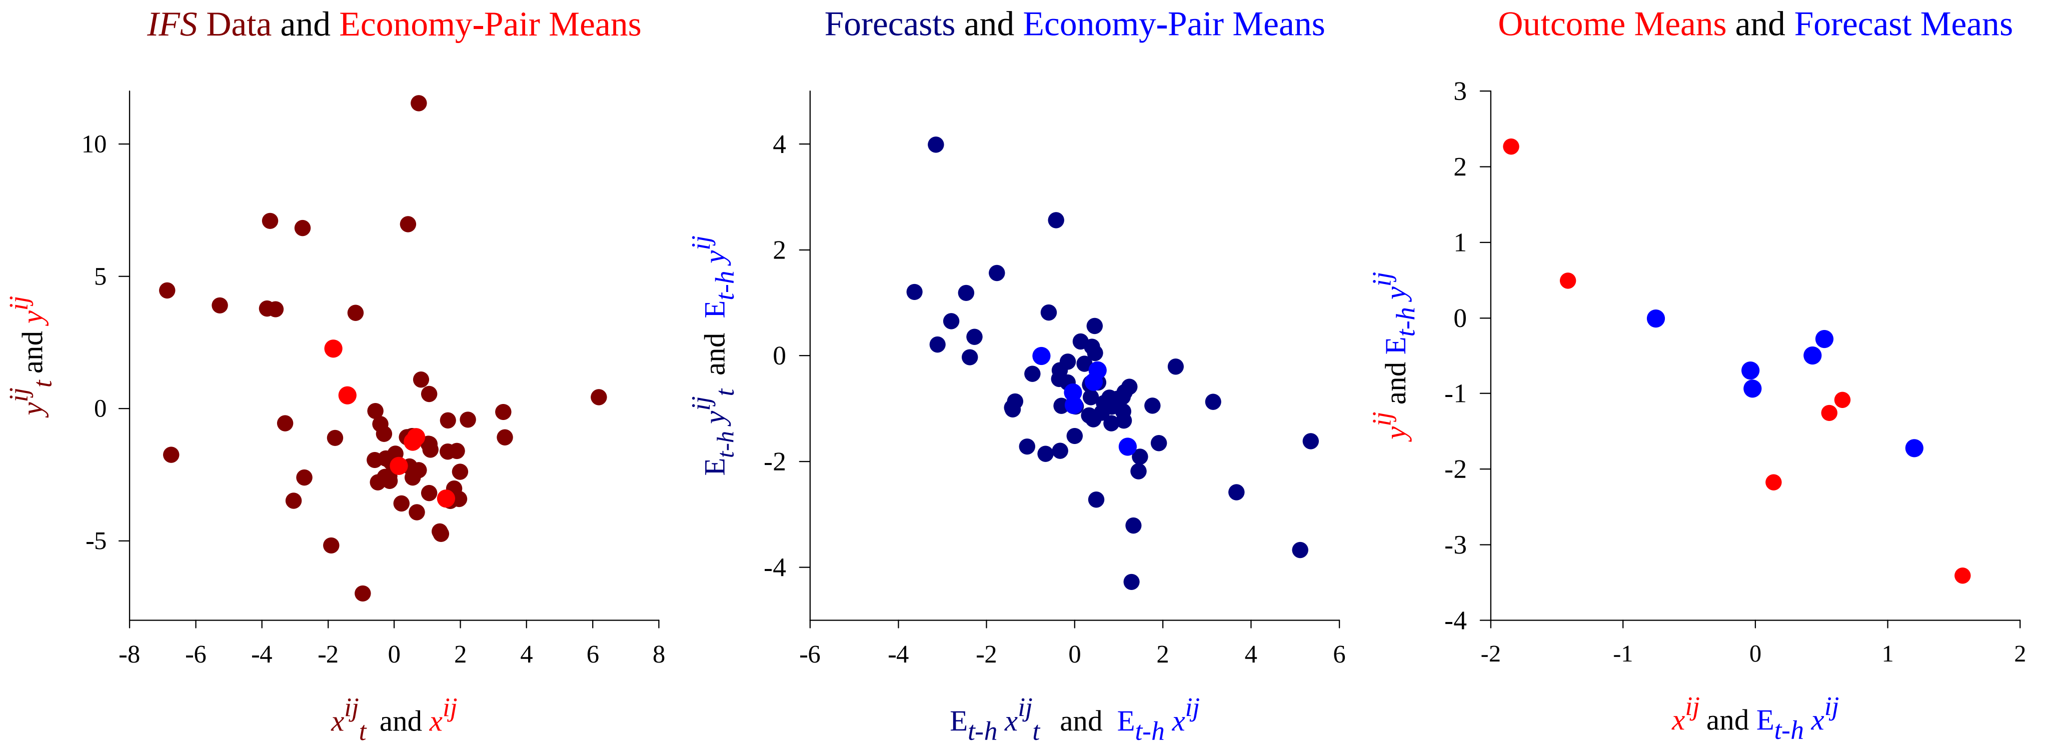

Figure 5: Euro Area Forecasts

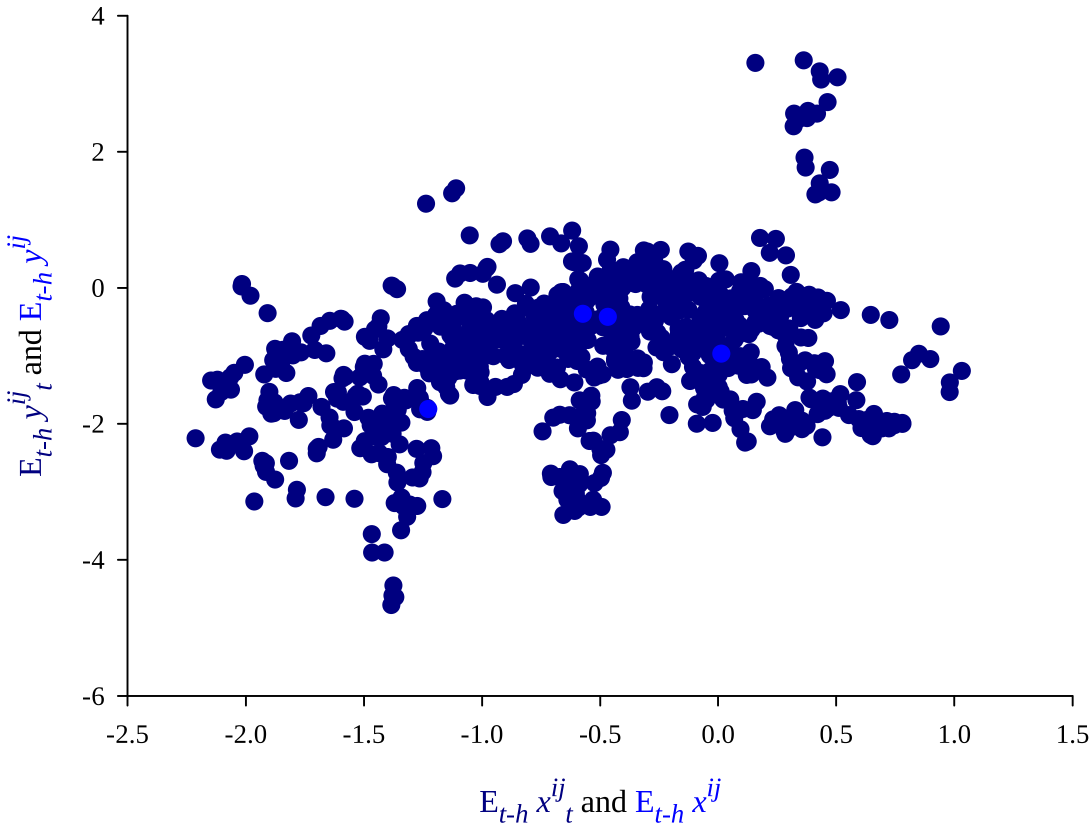

Figure 6: Euro Area Economy Pairs

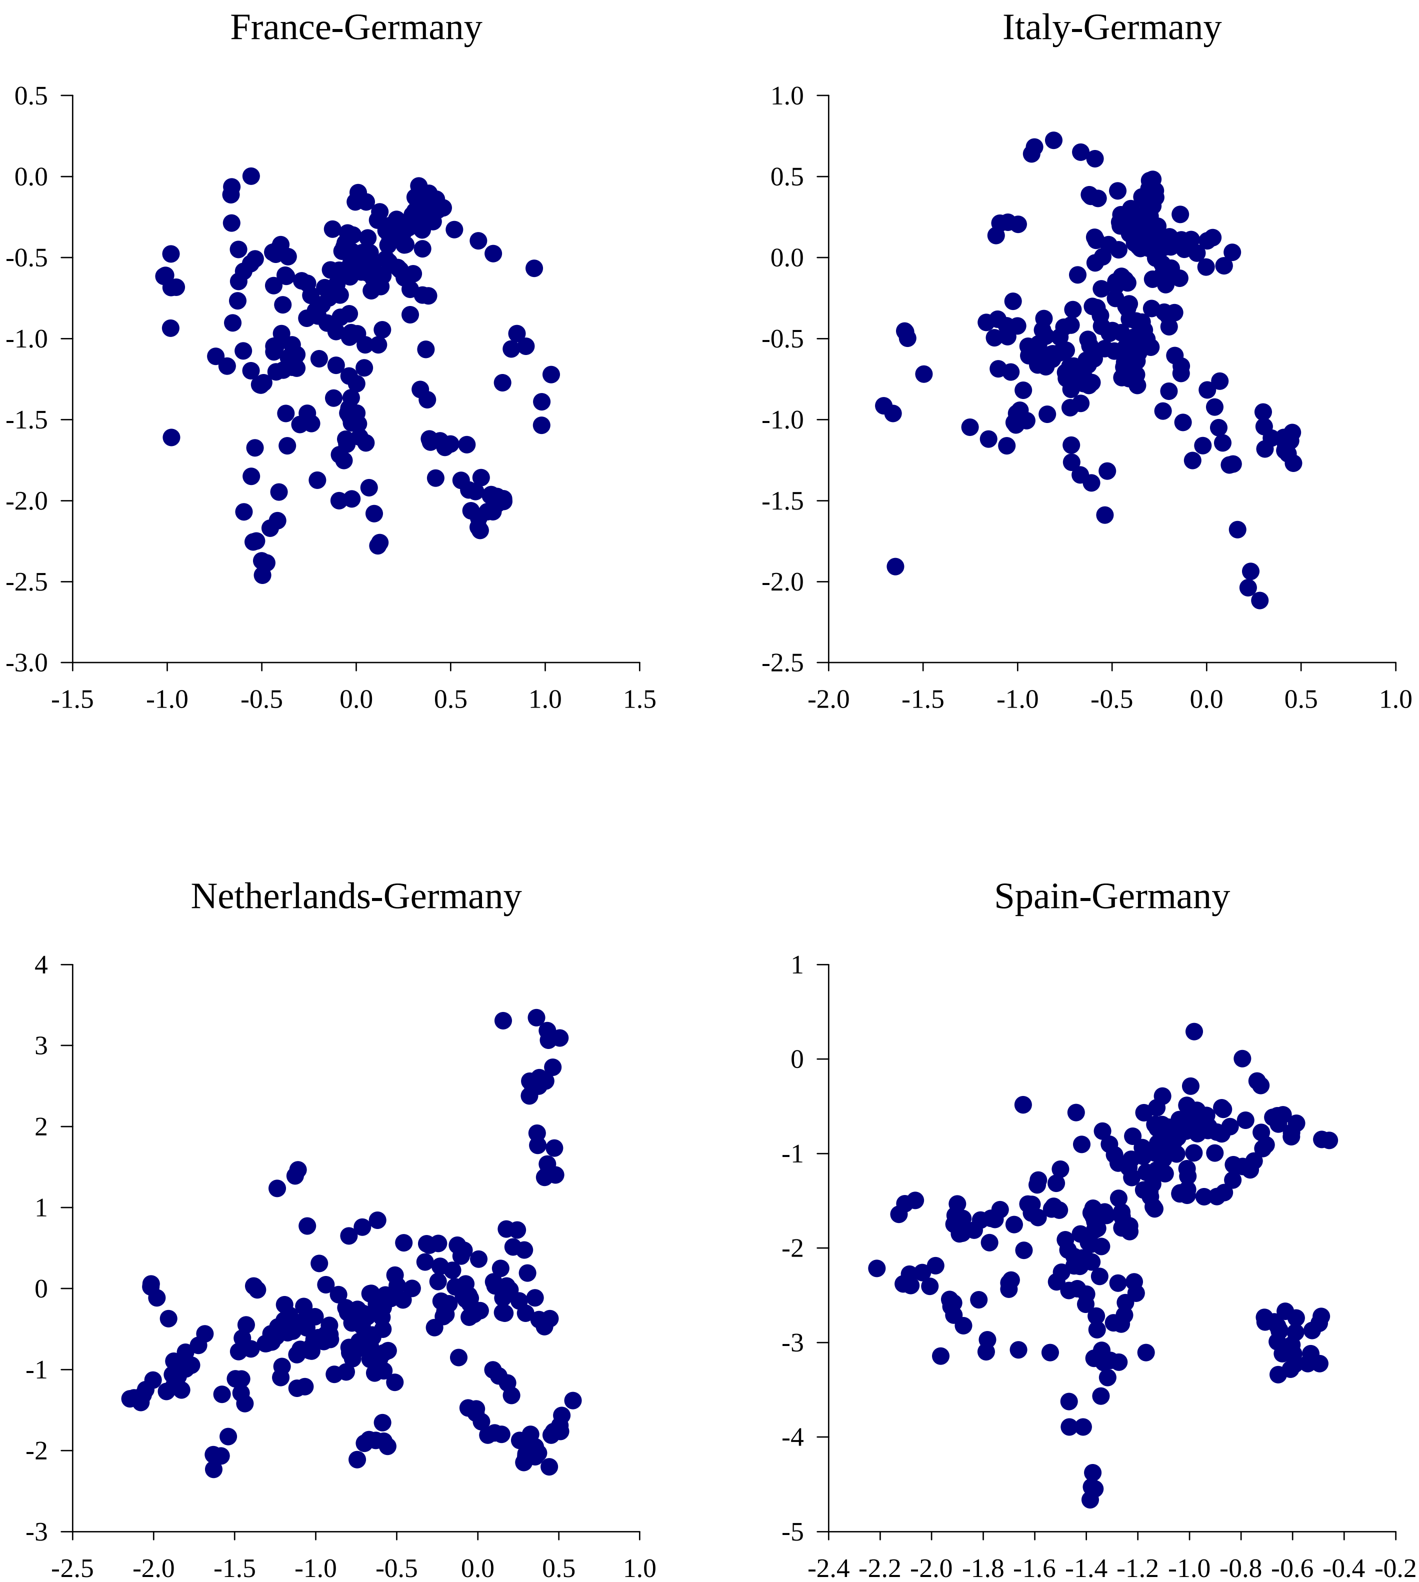

Figure 7: Euro Area Results by Year

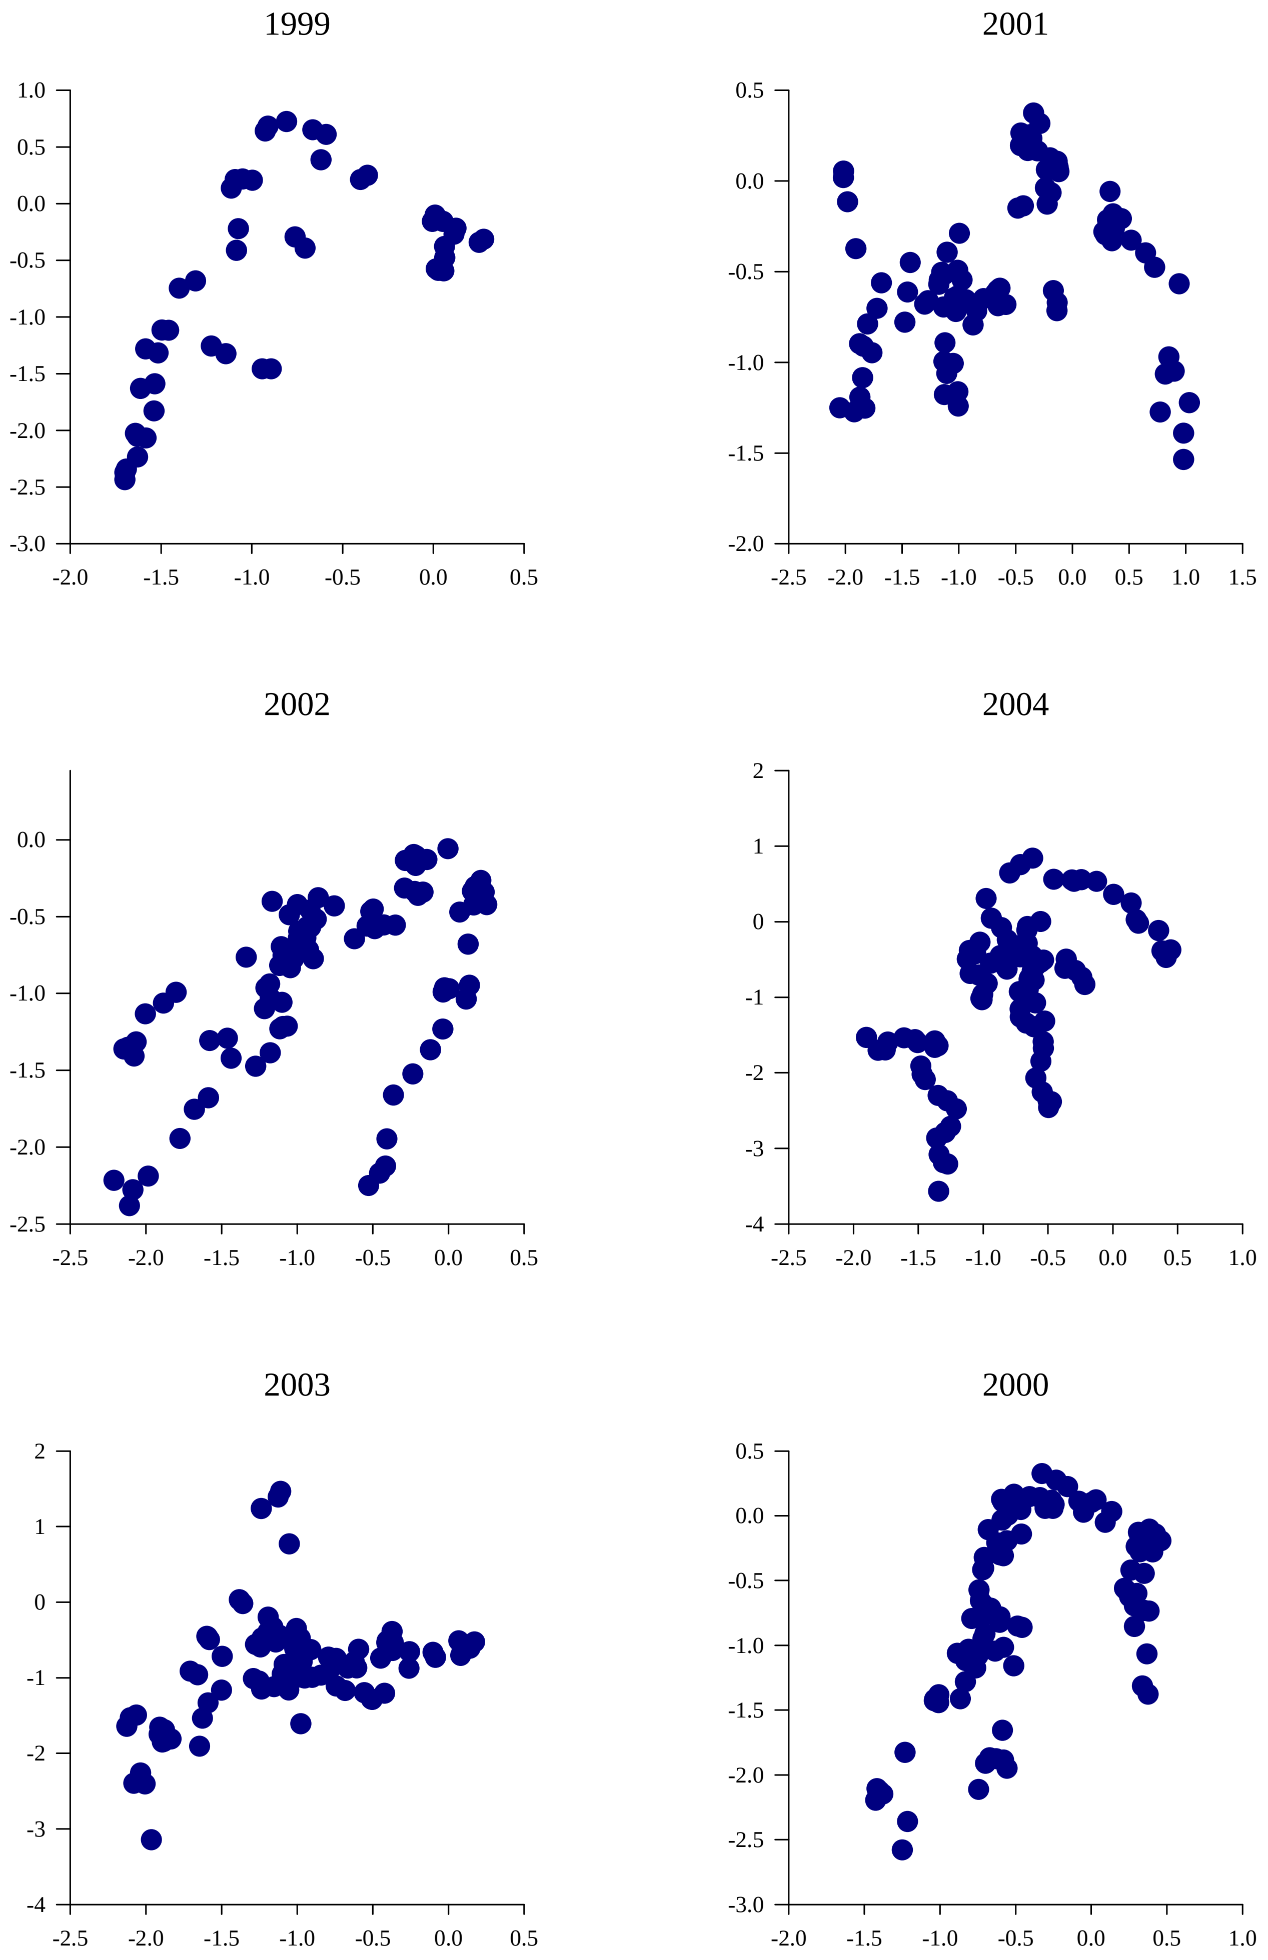

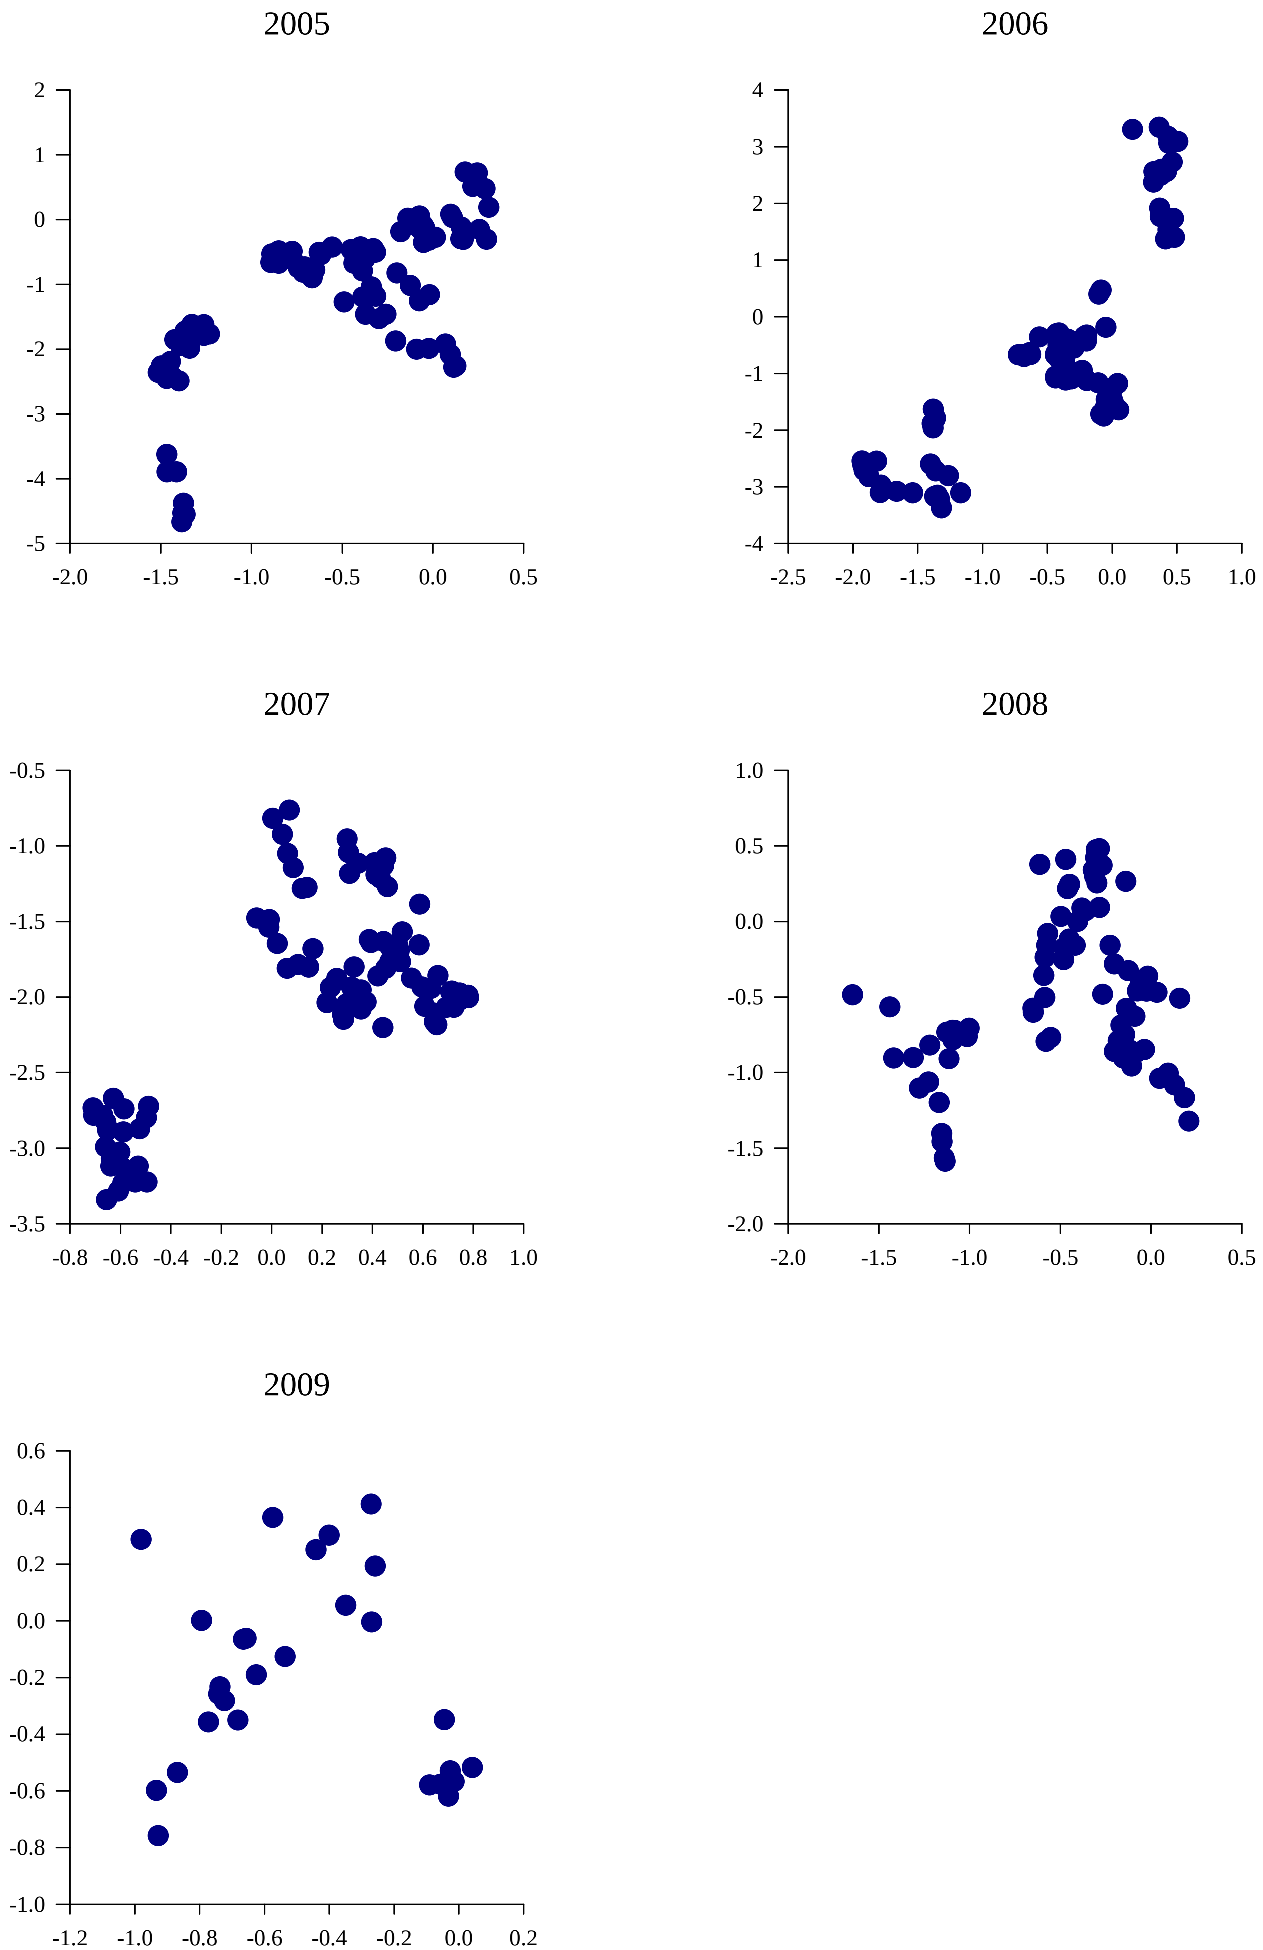
2005
2
1
0
-1
-2
-3
-4
-5
-2.0
-1.5
-1.0
-0.5
0.0
0.5
2006
4
3
2
1
0
-1
-2
-3
-4
-2.5
-2.0
-1.5
-1.0
-0.5
0.0
0.5
1.0
2007
-0.5
-1.0
-1.5
-2.0
-2.5
-3.0
-3.5
-0.8
-0.6
-0.4
-0.2
0.0
0.2
0.4
0.6
0.8
1.0
2008
1.0
0.5
0.0
-0.5
-1.0
-1.5
-2.0
-2.0
-1.5
-1.0
-0.5
0.0
0.5
2009
0.6
0.4
0.2
0.0
-0.2
-0.4
-0.6
-0.8
-1.0
-1.2
-1.0
-0.8
-0.6
-0.4
-0.2
0.0
0.2

Fig 8.1 Conditional Expectations,n=1

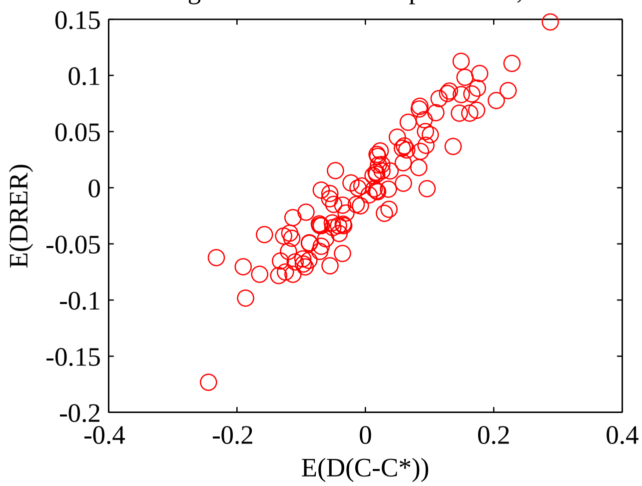

Fig 8.2 Growth Rates,n=1

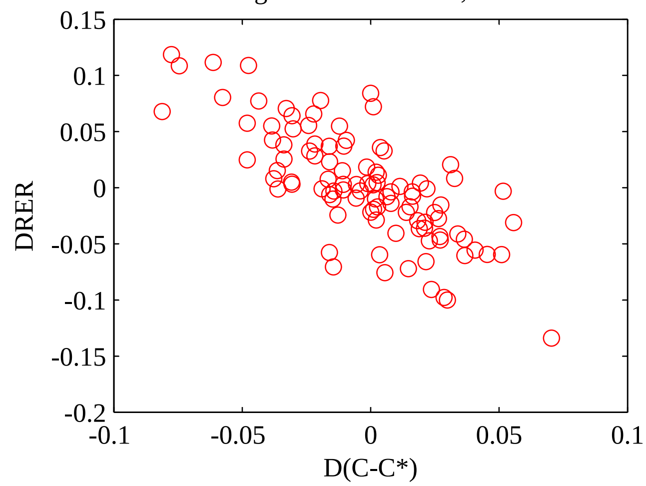

Fig 8.3 Conditional Expectations,n=0.5

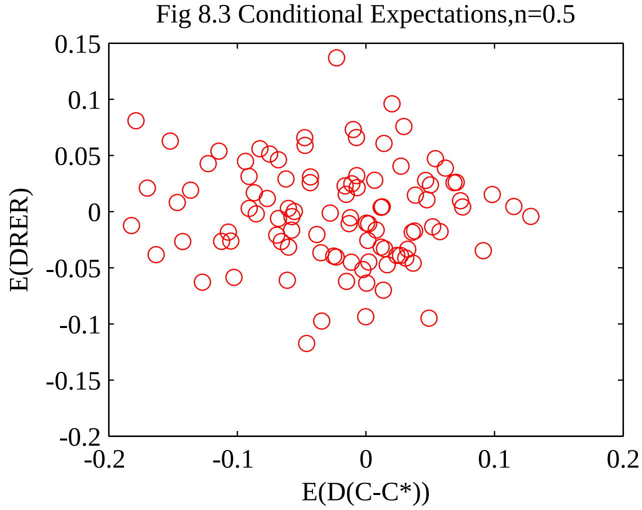

Fig 8.4 Growth Rates,n=0.5