

University of Wollongong

Research Online

National Institute for Applied Statistics
Research Australia Working Paper Series

Faculty of Engineering and Information
Sciences

2015

Investigating foreign portfolio investment holding: gravity model with social network analysis

Luke Muzur
University of Wollongong

Thomas Suesse
University of Wollongong

Pavel N. Krivitisky
University of Wollongong

Follow this and additional works at: <https://ro.uow.edu.au/niasrawp>

Recommended Citation

Muzur, Luke; Suesse, Thomas; and Krivitisky, Pavel N., Investigating foreign portfolio investment holding: gravity model with social network analysis, National Institute for Applied Statistics Research Australia, University of Wollongong, Working Paper 17-15, 2015, 24.
<https://ro.uow.edu.au/niasrawp/37>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Investigating foreign portfolio investment holding: gravity model with social network analysis

Abstract

Foreign Portfolio Investment (FPI) has been modelled with the aim of observing patterns, discovering potential inefficiencies and to provide evidence for theories. One particular point of interest is that over the course of many studies, bilateral FPI appears to increase with an increase in the correlation between the GDP growth rates of the sender and the receiver. The Capital Asset Pricing Model (CAPM) in financial theory predicts the opposite - and this is known as the correlation puzzle. We fit a number of gravity models for bilateral FPI holdings from the CEPII Coordinated Portfolio Investment Survey data using linear mixed models and latent space position models in order to test these theories. We use Maximum Likelihood Heckman Sample Selection estimators to account for potential bias in estimators that can be caused by frequent zeroes in our data. This results in two separate sets of response variables: the presence or absence of FPI between the two countries and the level of FPI between two countries conditional on its presence. Using seven cross-sections of data for the years 2000-2006, we estimate regression coefficients and apply model selection procedures to explain the correlation puzzle by accounting for a number of higher order dependencies such as sender and receiver random effects and then including transitivity, clustering and balance. Previous work had ignored these dependencies and had incorrectly assumed independence of residuals. We show that once these factors are captured the correlation fixed effect is occasionally negative, and often not significant, which matches more closely with the predictions of models used in finance, most notably the CAPM. It appears that if there is a presence of FPI between country s and country r , then as the correlation between the economic growth of these countries increases, the level of investment will decrease. However this correlation has no significant impact on the decision of country s to invest in country r .

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



National Institute for Applied Statistics Research Australia

The University of Wollongong

Working Paper

17-15

Investigating Foreign Portfolio Investment Holding: Gravity Model with Social Network Analysis

Luke Muzur, Thomas Suesse and Pavel N. Krivitsky

*Copyright © 2015 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong,
Wollongong NSW 2522. Phone +61 2 4221 5435, Fax +61 2 4221 4845.
Email: anica@uow.edu.au

Investigating Foreign Portfolio Investment Holdings: Gravity Model with Social Network Analysis

Luke Mazur¹, Thomas Suesse, Pavel N. Krivitsky

University of Wollongong

Abstract

Foreign Portfolio Investment (FPI) has been modelled with the aim of observing patterns, discovering potential inefficiencies and to provide evidence for theories. One particular point of interest is that over the course of many studies, bilateral FPI appears to increase with an increase in the correlation between the GDP growth rates of the sender and the receiver. The Capital Asset Pricing Model (CAPM) in financial theory predicts the opposite—and this is known as the correlation puzzle. We fit a number of gravity models for bilateral FPI holdings from the CEPII Coordinated Portfolio Investment Survey data using linear mixed models and latent space position models in order to test these theories. We use Maximum Likelihood Heckman Sample Selection estimators to account for potential bias in estimators that can be caused by frequent zeroes in our data. This results in two separate sets of response variables: the presence or absence of FPI between the two countries and the level of FPI between two countries conditional on its presence. Using seven cross-sections of data for the years 2000—2006, we estimate regression coefficients and apply model selection procedures to explain the correlation puzzle by accounting for a number of higher order dependencies such as sender and receiver random effects and then including transitivity, clustering and balance. Previous work had ignored these dependencies and had incorrectly assumed independence of residuals. We show that once these factors are captured the correlation fixed effect is occasionally negative, and often not significant, which matches more closely with the predictions of models used in finance, most notably the CAPM. It appears that if there is a presence of FPI between country s and country r , then as the correlation between the economic growth of these countries increases, the level of investment will decrease. However this correlation has no significant impact on the decision of country s to invest in country r .

Keywords: international trade, gravity models, latent space, correlation puzzle

Email address: lmazur@uow.edu.au (Luke Mazur)

¹Corresponding Author

1. Highlights

- Modelling of presence and level of bilateral FPI holdings to account for bias in estimators as a result of frequent zeroes in the data set
- Previous work had assumed the residuals were independent, which may not be the case. We consider the countries involved as a network and add random effects to capture residual dependencies between FPI holdings with the same sender and/or the same receiver. We also add latent space position random effects to capture higher order dependencies in this network of countries.
- These models feature estimates far closer to zero for the correlation fixed effect than models fitted via Ordinary Least Squares (OLS), and, at times, it was even reported to be negative—potentially one step closer to solving the correlation puzzle.
- Validation of the new models in the form of Out Of Sample Predictive Performance (OOSPP)

2. Introduction

Foreign Portfolio Investment (FPI) refers to purchasing a portion of a target, usually a company, in another country, with the primary purpose of obtaining monetary remuneration in the future. This can be contrasted with Foreign Direct Investment (FDI), which may or may not take place on a public market, and is primarily undertaken for control of companies. FPI is far more inclusive in who is able to participate and benefit from it, and as such has different patterns compared to the FDI.

FPI holdings are the amount that country s has invested into country r at a specific point in time (usually end of financial year), not for the purpose of controlling management. Flows, in contrast, refer to the movement of funds used to engage in this investment over a given time period.

A large amount of existing literature postulates that the amount of international bilateral goods trade between two nations is a product of their size (often captured using GDP or Market Capitalization) divided by the distance between them (Head et al., 2013). The origins of such a model, known as the *gravity model*, can be traced back to Savage and Deutsch (1960), who developed a theoretical multiplicative model of bilateral trade. Shortly afterwards Tinbergen (1962) analysed the development of third world countries and how interactions with first world investment partners, colonial or otherwise, affected their development. A small part of that paper featured an empirical study in modelling bilateral trade flows using a gravity model not dissimilar to the one currently in use. The purpose of this study was that in theory, if the gravity model employed was a good fit in the majority of cases then countries that had positive or negative deviations from this could be identified and studied further. This could lead to the uncovering of inefficiencies in global trade patterns that

could be hampering overall growth. Certain other covariates were added to the analysis but were deemed to have far weaker explanatory power than the three key covariates of importer size, exporter size, and geographic distance. The distance covariate was used to account for transaction costs, although there was speculation that it could be capturing informational asymmetry between sender and receiver as well. It was a static analysis and all time components were effectively ignored.

Despite these findings, the gravity model was deemed too closely related to physics, with limited economic theoretical support, and so was not well accepted amongst economists. Anderson (1979) derived the gravity model using properties of expenditure systems by starting from a framework in which every country is completely specialized in the production of a unique good, with no transport costs, tariffs, or other transaction costs. Disregarding error structure it can be shown that the flows from country s to country r are proportional to the product of their outputs. Though a number of strong assumptions were made, these were gradually removed to produce a solid economic foundation for the gravity model. This was considered too complex to be used in regular economic analysis, however.

It was not until the work of McCallum (1995) that the gravity model was considered as a legitimate tool for modelling trade flows. This paper featured a study of trade flows between provinces and states of USA and Canada. Despite the two countries being so closely integrated and similar in terms of culture and institutions, the gravity model (with a dummy variable for whether the states/provinces share a border) proved to have very strong explanatory power. The free trade agreement present at the time, as well as the integration suggested that the -1.5 coefficient of distance (far larger than predicted by economic theory), was not merely representing transaction costs.

Distance proxying for something other than transaction costs suggested that distance could have implications for equity flows as well as trade—in particular when considering the Capital Asset Pricing Model (CAPM). The CAPM assumed that investors will diversify as best they can in order to remove all firm-specific risk and only be vulnerable to market risk that affects all securities. Bilateral cross border equity flows should not be affected by such transaction costs, and, in fact, according to the CAPM, flows should increase with distance. This is because countries that are farther apart tend to have less correlation between their respective GDPs, and so investing in geographically distant countries would maximise the benefits of risk sharing. These predictions have been empirically disappointing: the proportion of domestic investment was far higher than would be expected under such a model. It seemed unlikely that in the early 2000s, significant barriers were preventing such flows given the amount of globalisation and integration between countries. This became known as the “correlation puzzle”.

Portes et al. (2001) and Portes and Rey (2005) hypothesised that distance was proxying for informational asymmetry, which leads to increased transaction costs and increased risk for less informed parties. They found that the gravity model was a good fit, and that controlling for a covariate that is highly

correlated with telephone communication between country s and country r lessened the effect of geographic distance. They went on to propose that since FPI flows and FPI holdings are highly correlated, the link between goods trade and FPI holdings should be studied further. Since both trade and equity holdings (and flows) are influenced by informational asymmetry, they should be strongly correlated, or even influence each other as an increase in one leads to more information that could, in turn, increase the other. Aviat and Coeurdacier (2007) showed that this was indeed the case, and that causality runs both ways. They argued that this is so strong that distance affects FPI holdings mainly through its effect on trade in goods. They found that when goods trade is included in the model, the impact of distance is vastly diminished and concluded that the two cannot be modelled separately. This showed that bilateral trade affected bilateral FPI holdings, but this was not well explained in terms of economic theory. Furthermore, this led to the question of whether countries that are more open to trade are more open to FPI. Peter (2012) developed a theoretical three-country model that explained the observed behaviour starting from basic economic principles. She then fit a model that included bilateral trade but also total trade of the investor and found that this was a significant factor adding credibility to her theory.

This led researchers to consider if there were higher order dependencies between the dyadic data, and how to model such links. Ward and Hoff (2007) applied a social network model to international bilateral goods trade. They found that latent covariates representing the inner products of k dimensional vectors had significant explanatory power. As far as the authors are aware, such models have not been applied to bilateral FPI holdings. We investigate them here, and go one step further by incorporating the random sender and receiver effects into the model along with the latent covariates to capture these dependencies more robustly using a method developed by Krivitsky et al. (2009).

A criticism of past combined theoretical and empirical studies was that many considered Ordinary Least Squares (OLS) estimates, but these have been shown to be inappropriate where the data sets contain many observations that are zero. This is naturally a problem in models using logarithmic transformations, and a number of methods can be used to address this issue. Martin and Pham (2009) empirically studied a number of methods using simulation and showed that Maximum Likelihood Heckman Sample Selection estimators performed best in the majority of situations. Furthermore, certain implementations lead to large amounts of bias in estimators, and great care was taken in resolving the frequent zeros issue in this paper.

The key purpose of this paper is to effectively model FPI holdings and move a step closer to solving the correlation puzzle, by applying models that more accurately capture the higher order dependencies between countries. The FPI can be represented in the form of a matrix showing the amount each country invested in another for a given year in millions of nominal US dollars. These network data intrinsically contain many potential dependencies between FPI values, which calls for models that can account for them. The covariates are obtained from a number of reliable sources providing information about the

countries in question and other measures of interaction among them.

The remainder of the paper is structured as follows: Section 3 discusses the data, then Section 4 describes the models fitted, the assumptions made, and the estimation methods. The results of these models are shown in Section 5. Finally, some concluding remarks as well as potential areas of improvement are presented in Section 6.

3. Data

Network data is often presented in the form of a *sociomatrix*, where the element in the s th (sender) row and r th (receiver) column is some measure of interaction between the s th node and the r th node. In the case of a directed network this is the amount that node s sends to node r and is represented as $Y_{s,r}$. If there are n nodes then theoretically there can be $n \times (n - 1)$ potential data values. In practice, a lot of the information for these nodes is missing, and the number of usable observations is significantly smaller—call it N . For convenience, the response network can be vectorised into an $N \times 1$ matrix \mathbf{Y} , and this notation is used throughout this paper. Here $\mathbf{x}_{s,r}$ is a row vector of covariates specific to the relationship between country s and country r , such as distance between s and r . These are stacked to form the design matrix $\mathbf{X} = (\mathbf{x}_{1,2}^\top, \mathbf{x}_{1,3}^\top, \dots, \mathbf{x}_{1,q}^\top, \dots, \mathbf{x}_{q-1,q}^\top)^\top$ of dimension $N \times p$ where each $\mathbf{x}_{s,r}$ is a vector with p elements. We consider the following covariates:

Gross Domestic Product (GDP) $_i$

GDP of country i in millions of nominal USD. This is used as a proxy for the “size” of a country, as it is assumed that “larger” countries are both more likely to invest more, in larger amounts and be more likely to be invested into and in larger amounts than “smaller” countries.

tradebyGDP $_{s,r}$

The trade between countries s and r divided by the product of their GDP, all in millions of nominal USD. It has been shown that trade and FPI holdings are closely related.

weighted distance (distw) $_{s,r}$

Distw between country s and country r . This is used as a proxy for information asymmetry between sender and receiver. For more information, see http://www.cepii.fr/PDF_PUB/wp/2011/wp2011-25.pdf. (Mayer and Zignago, 2011)

contiguity (contig) $_{s,r}$

Indicator of whether country s and country r share a border. It seems likely that this is a special case of distance and as such should be included separately.

time difference (timediff) $_{s,r}$

The difference in time zones in hours between country s and country r (between 0 and 12).

common official language (comlang_off)_{*s,r*}

Indicator of whether country *s* and country *r* have a common official language. It seems likely that sharing a common language would assist in mitigating information asymmetry as the communication barrier is lessened.

General Agreement on Tariffs and Trade (GATT)_{*i*}

Indicator of whether country *i* is part of the General Agreement on Tariffs and Trade (GATT). Countries that are part of the GATT are more likely to be open to trade, and potentially to investment as well.

correlation_{*s,r*}

The correlation of the economic growth of country *s* and country *r* over the last 10 years. Theoretically countries that have less correlation in their growth rates are more likely to invest in one another so as to maximise the benefits of diversification.

Regional Trading Agreement (RTA)_{*s,r*}

Indicator of whether country *s* and country *r* are part of the same regional trading agreement (RTA). Countries that are part of the same RTA are more likely to trade with each other which lessens the informational asymmetry between them leading to increased FPI between them.

common currency (comcur)_{*s,r*}

Indicator of whether country *s* and country *r* share a common currency. It is assumed that a common currency would make investment from sender to receiver less difficult.

common legal origin (comleg)_{*s,r*}

Indicator of whether country *s* and country *r* share a common legal origin. Countries that have the same legal origin are more likely to have similar laws, resulting in less information asymmetry regarding various legislation and its impacts on investment.

We used the natural logarithm transformation on tradebyGDP, GDP, and distw, with other covariates used unchanged.

The dependent variable was obtained from the CEPII Coordinated Portfolio Investment Survey in the form of a sociomatrix and was transformed into a vector. These data were taken for the years 2001–2007. There were 57–60 sender countries and 145–160 receiver countries, depending on the year. These were then merged with two other data sets: the CEPII data set “Network Trade” and the “Life During Growth” data set from the World Bank. The former contains almost all of the covariates, including the key covariates of GDP and distw. However, it did not extend sufficiently far into the past to calculate the correlation, so it was instead calculated from the “Life During Growth” data set. For more details, see Appendix B.

Due to the potential dual causation between trade and FPI, the natural logarithm of the FPI of the following year is used as the response. This means

that the FPI is caused to a degree by the covariates by virtue of them being from the previous year, but the covariates are not caused by the FPI.

Empirically, the temporal nature of the data is not very salient—the changes over time are minimal. For this reason, the paper will focus on a single year (covariates from 2001 and FPI from 2002, denoted as year 2001) and discuss any deviations for other years. It is also worth noting that financial offshore centers have been excluded from the analysis where they were easily identified, because these present unique investment patterns that are not of interest in this paper. We have chosen to err on the side of caution in labelling a country as a financial offshore center, so it is likely that some remain in the analysis, but as the sender and receiver random effects will absorb this, the fixed effects are unlikely to be affected to a large degree.

4. Methodology

4.1. Gravity Model

A gravity model postulates the strength of an interaction between two parties to be proportional to the size of each party divided by the distance between them. In this case, we conjecture that the amount of FPI holdings between country s and country r is proportional to their GDP (Gross Domestic Product) divided by the distance between them. It is worth noting that the response and the covariates are expressed in nominal USD:

$$\text{FPI}_{s,r} \propto \frac{\text{GDP}_s \times \text{GDP}_r}{\text{distw}_{s,r}}$$

$$\log(\text{FPI}_{s,r}) = \log(\text{GDP}_s) + \log(\text{GDP}_r) - \log(\text{distw}_{s,r}) + C$$

Models such as this are the benchmark and are fitted using Ordinary Least Squares (OLS), which assumes that the residuals are independent. This is unlikely to be the case, however: if country s is particularly keen on investing, then its sender-residuals are likely to be generally positive, and, similarly, if country r presents good investment opportunities, then its receiver-residuals are likely to be generally positive. It seems more likely that a large amount of residual dependence can be attributed to the specific sender and receiver in each observation and that these can be captured through the use of Linear Mixed Models.

4.2. Linear Mixed Model (LMM)

The LMM that incorporates sender and receiver random effects has the form

$$Y_{s,r} \equiv \log(\text{FPI}_{s,r}) = \mathbf{x}_{s,r}\beta + u_s + u_r + \epsilon_{s,r} \quad (1)$$

where $\mathbf{x}_{s,r}$ is a row of covariates and β is the vector of parameter estimates. Additionally, u_s and u_r are the sender and receiver random effects, both normally distributed with mean 0 and variances σ_s^2 and σ_r^2 respectively. $\epsilon_{s,r}$ is the residual error normally distributed with mean 0 and variance σ_ϵ^2 . `lmer` from the

R package `lme4` (Bates et al., 2014) was used to obtain the Restricted Maximum Likelihood Estimates for these parameters.

These models assume that two separate values of FPI that do not share the same sender or the same receiver country will be independent of each other, which is more reasonable than assuming all observations are independent, as OLS does. The potential for reciprocity is still ignored however. Given that our observations form a large network and each country can potentially affect every other country, it is worthwhile considering higher order dependencies.

4.3. Latent Space Position Model (LSPM)

One way to capture higher order dependencies is by assuming each vertex (or node) has a latent position in a k -dimensional social space. This position represents unobserved factors affecting its interactions, including but not limited to homophily (i.e., “birds of a feather flock together”) in ways not captured by the fixed effects. We assume that the edge values are stochastically independent given a function of their latent position vectors. Ward and Hoff (2007) used a bilinear coefficient representing the inner product of the latent positions of the sender and receiver to account for transitivity (i.e., “friend-of-a-friend is a friend”), clustering, and balance within the dyads. Following them, we will assume that these position vectors are independent k -dimensional random vectors, multivariate normal with a mean of zero and a covariance matrix $\sigma_l^2 I_n$ where I_n is the $n \times n$ identity matrix. In the following, these latent space position vectors are considered as row vectors.

The LSPM that captures these higher order dependencies has the form

$$Y_{s,r} \equiv \log(\text{FPI}_{s,r}) = \mathbf{x}_{s,r}\beta + u_s + u_r + \mathbf{z}_s \mathbf{z}_r^\top + \epsilon_{s,r}, \quad (2)$$

where $\mathbf{z}_s \mathbf{z}_r^\top$ are the random latent space position row vectors of the sender and the receiver respectively. All other variables are as in equation (1).

Maximum Likelihood estimation including REML estimation for these models is much more difficult because the bilinear form $\mathbf{z}_s \mathbf{z}_r^\top$ is not normally distributed, even though \mathbf{z}_s and \mathbf{z}_r are. As a result, `ergmm` from the R package `latentnet` (Krivitsky and Handcock, 2008) employs a Markov Chain Monte Carlo (MCMC) technique of Gibbs Sampling over the covariates to obtain parameter estimates.

4.4. Model Selection and Validation

The fixed effects in LMM and LSPM will be selected based on backward elimination via BIC (Schwarz et al., 1978). Coefficient estimates of each model were used to make predictions for all seven separate datasets used and the sum of squared errors were calculated. These were used to compare both types of models to each other, and to a baseline model fitted via OLS. See Appendix C for more details.

4.5. Generalized Models for Presence of FPI

Some pairs of countries have no measurable FPI. In our data in particular, any FPI value below 500K nominal USD was coded a 0i (no FPI), which requires special consideration. Here, we model the binary response of whether there was significant investment by country s into country r or not. Let

$$\tilde{Y}_{s,r} = \begin{cases} 1 & \text{if FPI}_{s,r} > 0 \\ 0 & \text{if FPI}_{s,r} = 0 \end{cases}.$$

We are interested in modelling the probability $\phi_{s,r}$ that $\tilde{Y}_{s,r}$ is 1. The covariates are continuous and, for the most part, have no upper bound, so we turn to models based on logistic regression. Let ϕ be the expectation of $\tilde{\mathbf{Y}}$ and g as the link function.

Apart from the change of \mathbf{Y} to $\tilde{\mathbf{Y}}$ as the response variable, all other notation remain the same. As before, independence of residuals can still not be assumed, and, for this reason, we turn to generalised linear mixed models and generalised latent space position models.

4.6. Generalised Linear Mixed Model (GLMM)

In GLMMs, sender and receiver random effects are added to account for the dependence between relationships that share the same sender or receiver respectively, in much the same way as the LMMs:

$$g(\phi_{s,r}) = \mathbf{x}_{s,r}\boldsymbol{\beta} + u_s + u_r + \epsilon_{s,r}. \quad (3)$$

Function `glmer` from the R package `lme4` (Bates et al., 2014) was used to calculate the Maximum Likelihood Estimates, while `blme` (Dorie, 2014) was used to verify the results that had difficulty converging. These GLMMs have the same drawbacks as the LMMs, in that they do not account for reciprocity, nor do they account for higher order dependencies such as transitivity, clustering or balance.

4.7. Generalized Latent Space Position Model (GLSPM)

GLSPMs capture these effects through the use of latent space position random vectors for each node. The inner products of these are incorporated into the model in the same way as the latent space position models (2):

$$g(\phi_{s,r}) = \mathbf{x}_{s,r}\boldsymbol{\beta} + u_s + u_r + \mathbf{z}_s\mathbf{z}_r^\top + \epsilon_{s,r}. \quad (4)$$

Similarly, we use `ergmm` from the R package `latentnet` (Krivitsky and Handcock, 2014) to estimate this model.

4.8. Model Validation for Generalised Models

Backward elimination via BIC is used to select the optimal GLMM and GLSPM as before. Receiver Operating Characteristic (ROC) curves were also examined and used to compare the two types of models to each other, as well as to a baseline model fit via OLS. See Appendix C for more details.

Table 1: GLMM and GLSPM Summaries

<i>Dependent variable: Presence of FPI</i>				
<i>Fixed Effects</i>	<i>Coefficient (Standard Error)</i>		<i>Coefficient (Standard Error)</i>	
Constant	-15.359***	(3.099)	-0.730	(0.455)
log(GDP _s)	0.725**	(0.223)	0.368***	(0.108)
log(GDP _r)	1.327***	(0.084)	0.235*	(0.115)
log(tradebyGDP)	0.317***	(0.072)	0.202*	(0.096)
log(distw)	-0.340*	(0.172)	-0.213	(0.204)
comlang_off	0.917***	(0.245)	1.066	(0.590)
GATT_r	1.470***	(0.411)	-	
correlation	-0.114	(0.207)	0.279	(0.393)
RTA	1.491***	(0.299)	0.791	(0.426)
<i>Random Effects</i>				
	<i>Coefficient</i>		<i>Coefficient</i>	
Sender Variance σ_s^2	4.818		5.257	
Receiver Variance σ_r^2	2.120		0.764	
Latent Space Position Variance σ_l^2	-		3.938	

Significance 0.05 > * > 0.01 > ** > 0.001 > ***

5. Results

In the following, the results for presence of FPI are given first, then the results for the level of log(FPI). Together, they can be used in forecasting because

$$E(\text{FPI}) = \phi \times \exp \left\{ E(\log(\text{FPI})|\text{FPI} > 0) + \frac{1}{2} V(\log(\text{FPI})|\text{FPI} > 0) \right\}.$$

5.1. Results for Presence of FPI

Next we report the models for whether FPI is present between country s and country r for a given year.

5.1.1. GLMM Results

Typically, the coefficients of log(GDP_s), log(GDP_r), log(tradebyGDP), and log(distw) are around 1, 1, 0.3, and -0.6, respectively, in gravity models fitted via OLS for the value of FPI. Because we are modelling only its presence probability on the logit scale, we would expect similar direction of the effect, but

not necessary magnitude and, indeed Table 1 (left) shows considerably different magnitudes, but similar coefficient signs. This was true for other years as well, although the magnitude of the estimate of $\log(\text{distw})$ was generally larger. The coefficient on correlation is negative, but it is not significant.

5.1.2. *GLSPM Results*

Table 1 (right) resembles the gravity models even less than Table 1 (left) did—the gravity model is almost indiscernible here. This is true for all other years although the model for 2001 resembles the gravity model the least. This particular model has the correlation coefficient positive, but it is not significantly different from 0.

In addition, we plotted the latent positions of each country in Appendix B.2. Certain nodes are isolated because they had no non-missing action as an investor or investee. The pattern of a large number of countries that act as hubs for investment and a smaller number of the developed countries that act as large investors and invest broadly is quite pronounced. There are several smaller groups of countries within the investees that act in similar ways and these do not seem to be as a result of being in the same region. The plots appear almost identical from year to year, even individual countries keep relatively the same location (remembering that these positions are invariant to rotation). The strong and persistent pattern of countries with large amounts of investment as opposed to countries that mainly acted as investees suggests that another model that considers latent space positions of each country in its capacity as an investor and an investee separately may be fruitful.

Table 2 shows posterior means for the results of the GLSPM while the plot of latent space positions deals with the latent space positions that satisfy Minimum Kullback-Leibler (MKL) divergence. This is because the posterior means demonstrate a better view of the entire posterior distribution of coefficients, while for the purposes of visualization the single “best” latent space positions is of most interest (Shortreed et al., 2006).

5.1.3. *Correlation Coefficient Comparison for Presence of FPI*

Figure 1 (top) shows a plot of the correlation coefficients for the GLMMs. Although there are two cases of negative correlation coefficients, neither of these are significantly different from zero at the 0.05 significance level, while the positive coefficients for the years 2003, 2004 and 2005 are significant. Looking at the wider confidence intervals, corresponding to the Bonferroni adjustment taking all years into account simultaneously (Dunn, 1961), the 2004 coefficient shows that overall the models suggest that the higher the correlation between the economic growth of the sender and receiver the more likely it is for FPI to be present between them. Figure 1 (bottom) contains the corresponding plot of the correlation coefficients for the GLSPMs; some estimates that are negative are lower than for the GLMMs while others that are positive have higher values. None of the coefficients were significant, and the simultaneous inference suggests there is no effect of correlation between the economic growth of the sender and the receiver on the presence of FPI between them.




Figure 1: Correlation Coefficients for GLMM (top) and GLSPM (bottom). The inside whiskers represent the 95% confidence intervals for the hypothesis test of $H_0 : \beta_{\text{correlation}} = 0$ vs. $H_a : \beta_{\text{correlation}} \neq 0$ while the outside whiskers represent the bonferroni correction confidence intervals for a simultaneous hypothesis test with the corresponding null and alternate hypotheses.

5.1.4. Results of Model Validation for Generalized Models

Both GLMMs and GLSPMs outperformed the baseline models, as their AUCs were never lower and significant in the vast majority of cases. GLSPMs only bettered the GLMMs in in-sample predictive performance, meaning that they have little use for prediction and it is likely they have been overfitted. This may be a consequence of binary data containing relatively little information compared to continuous measurements. As a result their correlation coefficient estimates are likely to be less reliable than those from the GLMMs.

5.2. Results for Level of $\log(\text{FPI})$

Next, we model $\log \text{FPI}$ for positive FPI from country s to country r in millions of nominal USD.

5.2.1. LMM Results

Table 2 (left) gives the LMM fit with sender and receiver random effects (Bates et al., 2014). As expected from the literature, the coefficients of $\log(\text{GDP}_s)$, $\log(\text{GDP}_r)$, $\log(\text{tradebyGDP})$, and $\log(\text{distw})$ are near 1, 1, 0.3, and -0.6, respectively, except that the coefficient $\log(\text{GDP}_s)$ is perhaps slightly higher. (This is not surprising, since the GLMM and GLSPM were fit to binary data.) The same holds true for other years. The most important observation, however, is that the estimate of correlation, while still positive, is much lower, yet removing it meant a drastic increase in BIC. This suggests that correlation may be itself correlated with certain higher order dependencies between the observations, and as we capture more of them, its estimate may decrease further. Our next model will capture dependencies of an even higher order.

Table 2: LMM and LSPM Summaries

<i>Dependent variable: Level of log(FPI)</i>				
<i>Fixed Effects</i>	<i>Coefficient (Standard Error)</i>		<i>Coefficient (Standard Error)</i>	
Constant	-13.357***	(2.019)	1.918***	(0.264)
log(GDP _s)	1.210***	(0.151)	0.283***	(0.052)
log(GDP _r)	1.029***	(0.069)	0.309***	(0.048)
log(tradebyGDP)	0.292***	(0.042)	0.139**	(0.051)
log(distw)	-0.618***	(0.073)	-0.468***	(0.083)
correlation	0.283*	(0.117)	-0.147	(0.159)
comcur	1.110***	(0.176)	1.092***	(0.194)
comleg	0.531***	(0.086)	0.326***	(0.099)
GATT _r	-		-0.696	(0.511)
nodematch(GATT)	-		0.528	(0.417)
<i>Random Effects</i>	<i>Coefficient</i>		<i>Coefficient</i>	
Sender Variance σ_s^2	2.350		1.788	
Receiver Variance σ_r^2	1.925		0.879	
Latent Space Position Variance σ_l^2	-		2.687	
Residual Variance σ_e^2	1.785		1.278	

Significance 0.05 > * > 0.01 > ** > 0.001 > ***

5.2.2. LSPM Results

The LSPM results are presented in Table 2 (right) (Butts et al., 2014; Krivitsky and Handcock, 2014): Firstly, the LSPMs still resemble the gravity models albeit with much lower estimates—not dissimilar to the LMM, and this was true for all LSPMs fit. The correlation is negative, however, which provides evidence for the CAPM. This is potentially the answer to the correlation puzzle—that the higher order dependencies were simply masking the true effect of the correlation.

We also plot the latent positions of each country in Appendix B.2. As before, MKL divergence latent space positions were used rather than posterior mean as the former provides the “best” configuration. Certain nodes are isolated because they had no non-missing data regarding their actions as an investor or an investee although there are far more isolated nodes than in the corresponding plot for the GLSPM. There appears to be a large number of countries that act as hubs for investment and a smaller number of the developed countries that act as large investors and invest globally on a large scale. There are several smaller groups of countries within the investees that act in similar ways, although these do not appear to be as a result of being in the same region. This effect is far less pronounced than in the corresponding plot for the GLSPMs however, which is not surprising given that any amount of investment is treated the same way in the GLSPM. The plots show a similar pattern from year to year (accounting for invariance to rotation) which shows the way the actors behave in the network does not drastically change from year to year.

5.2.3. Correlation Coefficient Comparison for Level of $\log(FPI)$

Figure 2 (top) shows the correlation coefficients for the LMMs. All the coefficients are positive but lower than is typical for OLS models. Furthermore only the correlation coefficient for the year 2006 is significant, but this is also sufficient to suggest that the level of $\log(FPI)$ from sender to receiver is higher with more correlation between their respective economic growths. Figure 2 (bottom) presents the correlation coefficients for the LSPMs; the majority are negative unlike the correlation coefficients for the LMMs. It is worth noting that only the coefficients for 2005 and 2006 are significant. This provides evidence for the idea that higher order dependencies between the nodes mask the true effect of correlation between the GDP growth rates on bilateral FPI. The correlation coefficient for the year 2005 also suggests that there is an overall effect of decreasing level of $\log(FPI)$ between the sender and the receiver as the correlation between their economic growths increases. The above plots indicate that as we capture more of these dependencies the correlation coefficient continues to decrease, and this is potentially the answer to the correlation puzzle.

5.2.4. Results of Model Validation

Both LMMs and LSPMs outperformed the baseline models, their sum of squared errors were almost never higher, and often below half of those calculated for the baseline models. The LSPMs had better in sample predictions, and better predictions one and two steps out than those of the LMMs.




Figure 2: Correlation Coefficients for LMM (top) and LSPM (bottom). The inside whiskers represent the 95% confidence intervals for the hypothesis test of $H_0 : \beta_{\text{correlation}} = 0$ vs. $H_a : \beta_{\text{correlation}} \neq 0$ for any given year while the outside whiskers represent the Bonferroni correction confidence intervals for a simultaneous hypothesis test with the corresponding null and alternate hypotheses.

5.3. Linear Trends in Coefficients

Although some trends were found, these were not consistent across years. Given that there were only 7 models fit for any type of model, these were deemed to be largely due to chance.

6. Concluding Remarks

The modelling of bilateral FPI holdings involved two steps, first the binary response of presence or absence of FPI from country s (sender) to country r (receiver), then the level of FPI from country s to country r conditional on the presence of FPI on the log scale. A key aim was to assess how the coefficient of the fixed effect of the correlation in GDP growth rates between country s and country r changed between models, particularly its sign.

6.1. Conclusions for presence of FPI

For the generalized models we applied a logit link to map the binary response of presence of FPI between country s and country r , to the continuous covariates. Adding sender and receiver effects led to correlation coefficient estimates close to zero, but there was sufficient evidence to suggest that higher correlation between the economic growths of countries s and r is positively associated with presence of FPI between the countries. When latent space position effects were added, we observed two instances of the correlation coefficient below zero, but none of them were significant. Comparing Out Of Sample Predictive Performance (OOSPP)

through the use of ROC curves of the two types of models fit showed that both outperformed the baseline models fit with GLM, while the Generalized Latent Space Position Models seldom bettered the Generalized Linear Mixed Models except in their in-sample predictive performance. This suggests over-fitting in the former; and based in the latter, there is evidence of a positive link between correlation of economic growths and likelihood of FPI from country s to country r .

6.2. Conclusions for level of $\log(FPI)$

By adding sender and receiver effects to the benchmark models fit via Ordinary Least Squares (OLS), we were able to partially replicate the correlation puzzle, albeit that the correlation coefficient was much lower than is typical for the benchmark models. There was sufficient evidence to suggest an overall effect of a positive link between higher correlation of economic growth between country s and country r and the level of $\log(FPI)$ from s to r . When latent space position effects were added, all of the correlation coefficients except one became negative. These models showed sufficient evidence to suggest the opposite effect, that lower correlation of the economic growths corresponded with higher levels of $\log(FPI)$. OOSPP was used to validate the models, and both types of models outperformed the OLS baseline models. The Latent Space Position Models proved superior for predicting one or two years ahead but were worse than the Linear Mixed Models otherwise. It appears that the former are a better fit, however parameter instability leads to their prediction quality deteriorating more rapidly than those of the latter—not necessarily evidence of a worse fit. As a result we tentatively reject the latter in favour of the former and conclude there is sufficient evidence to suggest that there is a negative link between correlation of the economic growth rates of country s and country r and the level of $\log(FPI)$ from s to r .

6.3. Recommendations

- Add separate latent space positions for each country’s capacity as a sender and as a receiver. This recommendation is based on certain unreciprocated ties present in the latent space position plots.
- Apply a stronger screening for financial offshore centers than was carried out in this paper, although the authors are doubtful it will make a significant difference to the results.

7. Acknowledgements

We would like to acknowledge the ITMS at the University of Wollongong for the use of their High Performance Cluster, without which it would be exceedingly difficult to fit the Latent Space Position Models and Generalized Latent Space Position Models. We would also like to acknowledge the support from the University of Wollongong and the National Institute for Applied Statistics Research Australia (NIASRA) in particular for their financial contribution during the submission process.

References

- Agresti, A., 2002. Categorical data analysis. A John Wiley and Sons, Inc. Publication, Hoboken, New Jersey, USA.
- Anderson, J. E., 1979. A theoretical foundation for the gravity equation. *The American Economic Review*, 106–116.
- Aviat, A., Coeurdacier, N., 2007. The geography of trade in goods and asset holdings. *Journal of International Economics* 71 (1), 22–51.
- Bates, D., Maechler, M., Bolker, B., Walker, S., 2014. lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-6. URL <http://CRAN.R-project.org/package=lme4>
- Butts, C. T., Handcock, M. S., Hunter, D. R., 2014. network: Classes for Relational Data. Irvine, CA, r package version 1.10.2. URL <http://statnet.org/>
- DeLong, E. R., DeLong, D. M., Clarke-Pearson, D. L., 1988. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, 837–845.
- Dorie, V., 2014. blme: Bayesian Linear Mixed-Effects Models. R package version 1.0-2. URL <http://CRAN.R-project.org/package=blme>
- Dunn, O. J., 1961. Multiple comparisons among means. *Journal of the American Statistical Association* 56 (293), 52–64.
- Head, K., Mayer, T., et al., 2013. Gravity equations: Workhorse, toolkit, and cookbook. *Handbook of International Economics* 4.
- Hlavac, M., 2014. stargazer: LaTeX/HTML code and ASCII text for well-formatted regression and summary statistics tables. Harvard University, Cambridge, USA, R package version 5.1. URL <http://CRAN.R-project.org/package=stargazer>
- Krivitsky, P. N., Handcock, M. S., 2008. Fitting position latent cluster models for social networks with latentnet. *Journal of Statistical Software* 24 (5).
- Krivitsky, P. N., Handcock, M. S., 2014. latentnet: Latent position and cluster models for statistical networks. The Statnet Project (<http://www.statnet.org>), R package version 2.6-2132.1-2014.05.25-03.32.19. URL CRAN.R-project.org/package=latentnet
- Krivitsky, P. N., Handcock, M. S., Raftery, A. E., Hoff, P. D., 2009. Representing degree distributions, clustering, and homophily in social networks with latent cluster random effects models. *Social Networks* 31 (3), 204–213.

- Martin, W., Pham, C., 2009. Estimating the gravity model when zero trade flows are frequent. World Bank manuscript.
URL https://www.deakin.edu.au/buslaw/aef/workingpapers/papers/2008_03eco.pdf
- Mayer, T., Zignago, S., 2011. Notes on CEPII's distances measures: The GeoDist database.
- McCallum, J., 1995. National borders matter: Canada–US regional trade patterns. *The American Economic Review*, 615–623.
- Peter, A., 2012. Bilateral trade, openness, and asset holdings. *Open Economies Review* 23 (4), 713–740.
- Portes, R., Rey, H., 2005. The determinants of cross-border equity flows. *Journal of International Economics* 65 (2), 269–296.
- Portes, R., Rey, H., Oh, Y., 2001. Information and capital flows: The determinants of transactions in financial assets. *European Economic Review* 45, 783–796.
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., Mller, M., 2011. proc: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics* 12, 77.
- Savage, I. R., Deutsch, K. W., 1960. A statistical model of the gross analysis of transaction flows. *Econometrica: Journal of the Econometric Society*, 551–572.
- Schwarz, G., et al., 1978. Estimating the dimension of a model. *The annals of statistics* 6 (2), 461–464.
- Shortreed, S., Handcock, M. S., Hoff, P., 2006. Positional estimation within a latent space model for networks. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences* 2 (1), 24.
- Tinbergen, J., 1962. An analysis of world trade flows. *Shaping the world economy*, 1–117.
- Ward, M., Hoff, P., 2007. Persistent patterns of international commerce. *Journal of Peace Research* 44 (2), 157–75.

Appendix A. Software

All analysis was performed with R, with special mention to the following:

- Mixed Modelling—lme4 (Bates et al., 2014)
- Setting Up Networks—network (Butts et al., 2014)

- Latent Space Position Models—latentnet (Krivitsky and Handcock, 2014, 2008)
- Bayesian Mixed Modelling (for checking lme4 output when convergence was in doubt)—blme (Dorie, 2014)
- ROC curves—pROC (Robin et al., 2011)
- Tables—stargazer (Hlavac, 2014)

R code can be obtained on request.

Appendix B. Data

Appendix B.1. Datasets

The FPI network data were obtained from the CEPII Coordinated Portfolio Investment Survey. It can be accessed from

<http://cpis.imf.org/>

date accessed: 24/03/2014.

The data on the majority of covariates were obtained from the CEPII data set “Network Trade”. This included the trade network data, GDP, distw, GATT, comleg, contig, comcur, RTA and comlang_off. It can be accessed from

http://www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=27

date accessed: 30/03/2014.

The data on correlation were obtained from the “Life During Growth” dataset from the World Bank. It can be accessed from <http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/0,,contentMDK:23407813~pagePK:64214825~piPK:64214943~theSitePK:469382,00.html> date accessed: 30/05/2014.

Data for timediff were collected from

<http://www.timeanddate.com/worldclock/custom.html?low=c&sort=1>

date accessed 30/04/2014.

Data for inflation were collected from

<http://www.multpl.com/inflation/table>

date accessed 20/09/2014.

Appendix B.2. Geographical Sample

Sender Countries: Aruba(ABW), Argentina(ARG), Australia(AUS), Austria(AUT), Belgium(BEL), Bulgaria(BGR), Bolivia(BOL), Brazil(BRA), Canada(CAN), Switzerland(CHE), Chile(CHL), Colombia(COL), Costa Rica(CRI), Czech Republic(CZE), Germany(DEU), Denmark(DNK), Egypt(EGY), Spain(ESP), Estonia(EST), Finland(FIN), France(FRA), United Kingdom(GBR), Greece(GRC), Hungary(HUN), Indonesia(IDN), India(IND), Ireland(IRE), Iceland(ISL), Israel(ISR), Italy(ITA), Japan(JPN), Kazakhstan(KAZ), Republic of Korea(KOR), Kuwait(KWT), Lithuania(LTU), Latvia(LVA), Mexico(MEX), Mongolia(MNG), Malaysia(MYS), Netherlands(NLD), Norway(NOR), New Zealand(NZL), Pakistan(PAK), Panama(PAN), Philippines(PHL), Poland(POL), Portugal(PRT), Romania(ROM), Russian Federation(RUS), Singapore(SGP), Slovakia(SVK), Slovenia(SVN), Sweden(SWE), Thailand(THA), Turkey(TUR), Ukraine(UKR), Uruguay(URY), United States of America(USA), Republica Bolivariana de Venezuela(VEN), South Africa(ZAF)


Receiver Countries: Aruba (ABW), Afghanistan(AFG), Angola(AGO), Albania(ALB), United Arab Emirates(ARE), Argentina(ARG), Armenia(ARM), Australia(AUS), Austria(AUT), Azerbaijan(AZE), Burundi(BDI), Belgium(BEL), Benin(BEN), Burkina Faso(BFA), Bangladesh(BGD), Bulgaria(BGR), Bosnia and Herzegovina(BIH), Belarus(BLR), Bolivia(BOL), Brazil(BRA), Brunei Darussalam(BRN), Bhutan(BTN), Botswana(BWA), Central African Republic(CAF), Canada(CAN), Switzerland(CHE), Chile(CHL), Cote d'Ivoire(CIV), Cameroon(CMR), Republic of Congo(COG), Colombia(COL), Comoros(COM), Cape Verde(CPV), Costa Rica(CRI), Czech Republic(CZE), Germany(DEU), Djibouti(DJI), Denmark(DNK), Dominican Republic(DOM), Algeria(DZA), Ecuador(ECU), Egypt(EGY), Eritrea(ERI), Spain(ESP), Estonia(EST), Ethiopia(ETH), Finland(FIN), Fiji(FJI), France(FRA), Gabon(GAB), United Kingdom(GBR), Georgia(GEO), Ghana(GHA), Guinea(GIN), The Gambia(GMB), Guinea-Bissau(GNB), Equatorial Guinea(GNQ), Greece(GRC), Guatemala(GTM), Guyana(GUY), Honduras(HND), Croatia(HRV), Haiti(HTI), Hungary(HUN), Indonesia(IDN), India(IND), Ireland(IRE), Islamic Republic of Iran(IRN), Iraq(IRQ), Iceland(ISL), Israel(ISR), Italy(ITA), Jamaica(JAM), Jordan(JOR), Japan(JPN), Kazakhstan(KAZ), Kenya(KEN), Kyrgyz Republic(KGZ), Cambodia(KHM), Kiribati(KIR), Republic of Korea(KOR), Kuwait(KWT), Lao, P.D.R.(LAO), Liberia(LBR), Libya(LBY), Sri Lanka(LKA), Lesotho(LSO), Lithuania(LTU), Latvia(LVA), Morocco(MAR), Moldova(MDA), Madagascar(MDG), Maldives(MDV), Mexico(MEX), Macedonia, FYR(MKD), Mali(MLI), Mongolia(MNG), Mozambique(MOZ), Mauritania(MRT), Malawi(MWI), Malaysia(MYS), Namibia(NAM), French Territories: New Caledonia(NCL), Niger(NER), Nigeria(NGA), Nicaragua(NIC), Netherlands(NLD), Norway(NOR), Nepal(NPL), New Zealand(NZL), Oman(OMN), Pakistan(PAK), Panama(PAN), Peru(PER), Philippines(PHL), Papua New Guinea(PNG), Poland(POL), Portugal(PRT), Paraguay(PRY), French Territories: French Polynesia(PYF), Qatar(QAT), Romania(ROM), Russian Federation(RUS), Rwanda(RWA), Saudi Arabia(SAU), Sudan(SDN), Senegal(SEN), Singapore(SGP), Solomon Islands(SLB), Sierra Leone(SLE), El Salvador(SLV), Sao Tome and Principe(STP), Suriname(SUR), Slovakia(SVK), Slovenia(SVN), Sweden(SWE), Swaziland(SWZ), Syrian Arab

Republic(SYR), Chad(TCD), Togo(TGO), Thailand(THA), Tajikistan(TJK), Turkmenistan(TKM), Tonga(TON), Trinidad and Tobago(TTO), Tunisia(TUN), Turkey(TUR), Tanzania(TZA) Uganda(UGA), Ukraine(UKR), Uruguay(URY), United States of America(USA), Uzbekistan(UZB), Republica Bolivariana de Venezuela(VEN), Vietnam(VNM), Republic of Yemen(YEM), South Africa(ZAF), Zambia(ZMB), Zimbabwe(ZWE)

sectionLatent Space Position Figures

Appendix B.3. Generalized Latent Space Position Models

Figure B.3: Latent Space Positions for Generalized Latent Space Position Model for 2001




Appendix B.4. Latent Space Position Models

Appendix C. Model Validation

Appendix C.1. Model Selection

Model selection for all models was performed by minimising the Bayesian Information Criterion (BIC). Correlation between GDP growths was kept in all of the models, as it was required for the research question. Due to the time taken to fit some of the models, a full model selection process could not be performed although several sets of covariates were selected based on covariates with strong explanatory power based on the BIC in simpler models for the data.

Figure B.4: Latent Space Positions for Latent Space Position Model for 2001



Appendix C.2. Out Of Sample Predictive Performance (OOSPP) for Level of $\log(\text{FPI})$ Models

The BIC is not a good indicator to use to compare models of different types, even once the optimal LMMs and LSPMs can be obtained (Schwarz et al., 1978). Furthermore, the goal of modelling and statistics in general is to make general observations and predictions from a sample. In this case, however, one cannot consider the data for 2001, for example, as a sample—it is in fact the census of FPI for the year 2001. Seeing as the effect of time is quite small, we can consider each year as a sample of a population of potential investment values and the relevant covariates. As a result, a measure of predictive accuracy can be obtained by using the coefficient estimates of each model fit for one year to predict $\log(\text{FPI})$ (given that some investment did occur) for other years; and then taking the squared differences between the observed values and the predicted, giving the sum of squared errors $\text{SSE} = \sum_{\text{all observed}} (\hat{Y}_{s,r} - Y_{s,r})^2$.

This SSE can be compared to that for a baseline linear model (fitted using the OLS method) with an additional parameter for the correlation between the GDP growth rates of country s and country r over the previous ten years, that is considered a good fit:

$$\log(Y_{s,r}) = \beta_1 \log(\text{GDP}_s) + \beta_2 \log(\text{GDP}_r) - \beta_3 \log(\text{distw}_{s,r}) + \beta_4 \log\left(\frac{\text{trade}_{s,r}}{\text{GDP}_s \times \text{GDP}_r}\right) + \beta_5 \text{correlation}_{s,r} + \epsilon_{s,r}$$

Then, a model's predictive power may be quantified with $\delta_{\text{model}} = \text{SSE}_{\text{model}} / \text{SSE}_{\text{baseline}}$:

$\delta_{\text{model}} > 1$ if the model in question has a larger sum of squared residuals from its predictions than the baseline model—a worse OOSPP.

A possible issue that arises is what to predict for a new country, which did not feature in the fitted model so and so does not have a sender or a receiver effect, nor a latent space position. The approach used here is to integrate (by simulation) over its random effects distribution (normal with a mean of zero and a variance estimated from the extant countries), and similarly for the latent space position.

As the response as well as several of the covariates are expressed in nominal USD, inflation may be a factor. This is further complicated by the fact that the response variable is one year ahead of the covariates, so each is inflated by a different amount. This was accounted for by adjusting the predictions for inflation relative to the year for which the model was fit. The actual year is largely irrelevant since we are observing the ratio of the quality of predictions of one model to another, a unitless measure.

Appendix C.3. OOSPP Results for level of log(FPI) Models

Table C.3 shows the results of the LMMs which are all significantly better than the baseline model fitted by OLS. Table C.3 shows that $\delta_{\text{LMM}} < 1$ particularly for the 1 and 2 step predictions. This means that the LMMs provide better predictive power than the benchmark. The predictions become worse as the time difference between the data the model was fitted on and the data upon which the model was used to make predictions increases, although in all cases these predictions were still better than the benchmark. The models also appear to be better at predicting into the future than estimating the past.

Table C.3: LMMs OOSPP

		Model Based on Year						
		2000	2001	2002	2003	2004	2005	2006
Predict for Year	2000	0.385	0.460	0.514	0.563	0.634	0.674	0.784
	2001	0.472	0.372	0.508	0.540	0.596	0.652	0.738
	2002	0.518	0.484	0.352	0.419	0.460	0.503	0.592
	2003	0.579	0.542	0.461	0.370	0.420	0.468	0.545
	2004	0.610	0.567	0.480	0.426	0.364	0.434	0.502
	2005	0.656	0.615	0.543	0.491	0.450	0.379	0.459
	2006	0.654	0.606	0.586	0.524	0.496	0.433	0.362

Finally, we consider the LSPMs. The elements of the main diagonal of Table C.4 all have $\delta_{\text{LSPM}} > \delta_{\text{LMM}}$, indicating that residuals are lower than for the previous type of models, and at one step out, they are slightly better, too. After this, they are comparable, although they seem to be much worse when estimating the past—worse even than the benchmark. This is still a useful finding because they are still the best model for short term predictions out of the two types fitted.

Table C.4: LSPMs OOSPP

		Model Based on Year						
		2000	2001	2002	2003	2004	2005	2006
Predict for Year	2000	0.316	0.335	0.466	0.622	0.788	0.878	1.082
	2001	0.384	0.255	0.431	0.547	0.705	0.788	0.994
	2002	0.397	0.311	0.264	0.402	0.526	0.589	0.796
	2003	0.480	0.388	0.363	0.297	0.417	0.502	0.689
	2004	0.544	0.464	0.430	0.375	0.306	0.402	0.581
	2005	0.614	0.547	0.541	0.472	0.409	0.295	0.484
	2006	0.607	0.578	0.612	0.534	0.469	0.366	0.374

Appendix C.4. OOSPP for Presence of FPI Models

The next logical step is to go through model selection processes for the generalized models. They first undergo similar backwards elimination as discussed in (Appendix C.1). Then, an OOSPP procedure is used to compare these models to each other and to investigate for how many years ahead do they outperform the predictions made by a baseline model. In this case, the baseline model has a logit link function mapping the binary response to the predictors (fitted by the Iteratively Reweighted Least Squares method):

$$g(\mu_{s,r}) = \beta_1 \log(\text{GDP}_s) + \beta_2 \log(\text{GDP}_r) - \beta_3 \log(\text{dist}_{w_{s,r}}) + \beta_4 \log\left(\frac{\text{trade}_{s,r}}{\text{GDP}_s \times \text{GDP}_r}\right) + \beta_5 \text{correlation}_{s,r} + \epsilon_{s,r}$$

This paper compares Area Under Curve (AUC) of the ROC curves created using the predictions from these generalized models (Agresti, 2002). Software for creating these curves and testing for significant differences between the predictors can be found in the R package `pROC` (Robin et al., 2011). Comparison of ROC curves is done using the DeLong method. (DeLong et al., 1988)

Appendix C.5. OOSPP Results for Presence of FPI

All the AUC values were higher for the GLMM ROC curves than those generated by the baseline models fitted using logistic regression. In every case the GLMMs outperformed the baseline models.

The GLSPM ROC curves had higher AUC than each corresponding generalized baseline ROC curve and in all but three cases the differences in AUC were significant at the 0.05 level.

Finally, comparing the ROC curves for the GLSPMs with the GLMMs in Table C.5, we discover that the former do not always outperform the latter.

The table demonstrates that the GLSPMs only provide better predictive performance for the data they were fitted on, and on occasion are superior for one step predictions. The vast majority of differences were significant at the 0.05 level with the key exception of the differences outside of the main diagonal

Table C.5: Positive Difference in AUC between Generalized Latent Space Position and Generalized Linear Mixed—denoted by +, − or 0 if difference not significant at 0.05 level

		Model Based on Year						
		2000	2001	2002	2003	2004	2005	2006
Predict for Year	2000	+	0	−	0	−	−	−
	2001	−	+	+	0	−	−	−
	2002	−	0	+	+	−	−	−
	2003	−	0	+	+	−	−	−
	2004	−	−	$\hat{+}$	−	−	−	−
	2005	−	−	−	−	−	$\hat{+}$	$\hat{+}$
	2006	−	−	0	−	−	−	$\hat{+}$

Characters with hats are significant at the 0.01 level.

Characters in bold are significant at the 0.001 level.

in favour of the GLSPM in all but two cases. The GLSPM for 2004 has no advantage over its GLMM counterpart—indeed this model had a particularly difficult time converging and this may be the reason for this behaviour. Table C.5 still shows that for the purposes of prediction one or two years into the future, the GLSPMs are occasionally superior to any other model fitted. Ultimately this illustrates that the GLSPMs by and large are not useful in terms of prediction, as they only have superior performance in prediction upon the data they were fitted to.

Appendix C.6. Reasoning Behind Separate Model Validation

Theoretically, the generalized linear model could have been used with its corresponding linear model and these could have been assessed together but if predictions are not good then this method does not enable one to determine which of the models is causing the problem, or if it is a mixture of the two. If the GLSPM was good for example, but the LSPM was deemed to be over-fitting and showed no improvement over the LMM for that year, then one could easily combine the two best models to create the best forecasts. As a result evaluating the two independently allows finding the ideal combination of the two.