

KIEL WORKING PAPER

Estimating the effects
of the “flight to
quality”, with an
application to German
bond yields and
interest payments



*Claire A. Boeing-Reicher and
Jens Boysen-Hogrefe*

ABSTRACT

ESTIMATING THE EFFECTS OF THE “FLIGHT TO QUALITY”, WITH AN APPLICATION TO GERMAN BOND YIELDS AND INTEREST PAYMENTS DURING THE EURO CRISIS

Claire A. Boeing-Reicher and Jens Boysen-Hogrefe

Recent calculations have suggested that the German federal government has saved roughly EUR 90-100 billion, cumulatively, due to low bond yields since the onset of the Euro crisis. In order to determine the contribution of the “flight to quality” to this sum, we define the flight to quality as a factor which has caused German bond yields and crisis country bond yields to move in opposite directions. Estimates show that only a small share is due to the flight to quality. Comparison with other approaches suggests that our factor approach is a promising way to look at the flight to quality.

Keywords: flight to quality, Euro crisis, bond yields, Germany, factor model
JEL classification: C32, C58, G12

Claire A. Boeing-Reicher

Amazon

E-mail: claire.anne.marie.reicher@gmail.com

Jens Boysen-Hogrefe

Kiel Institute for the World Economy

Kiellinie 66, 24105 Kiel, Germany

Phone: +49-431-8814-210

E-mail: jens.hogrefe@ifw-kiel.de

The responsibility for the contents of this publication rests with the author, not the Institute. Since working papers are of a preliminary nature, it may be useful to contact the author of a working paper about results or caveats before referring to, or quoting, a paper. Any comments should be sent directly to the author.

1 Introduction

In early 2015, the German federal ministry of finance, in conjunction with the German parliament (Deutscher Bundestag, 2015), stated that the German federal government had directly saved about EUR 94 billion due to persistently low bond yields since 2007. Other estimates, for instance, those of Boysen-Hogrefe and Filipczak (2015) and Dany et al. (2015), are of a similar order of magnitude. However, it has remained an open question how much of this fall in German bond yields has resulted from the “flight to quality” in response to the Euro crisis rather than resulting from other Euro area-wide factors, such as a fall in the natural real interest rate or changes in the ECB’s policy stance. In order to help disentangle these different phenomena, we formulate, estimate, and then compare three different models (including our own factor model) which capture the flight to quality in the Euro area. When we do this, we find that the flight to quality can account for a small portion of the fall in German bond yields since 2007 at short maturities, but a higher portion of that fall at longer maturities. Based on these estimates, we quantify the direct effects of the flight to quality on interest payments for the German federal government; we find that the flight to quality accounts for about one quarter to one third of the cumulative direct fall in interest payments, on the order of EUR 23 to 29 billion.¹

We estimate the effects of the flight to quality three different ways, in order to cope with model uncertainty. Our estimates are based on a definition of the flight to quality during the Euro crisis, whereby German bond yields and crisis-country bond yields move in opposite directions. To capture these movements, we first regress movements in German government bond yields on crisis country bond spreads. We do this because bond spreads offer a simple proxy for episodes in which crisis country bond yields and German bond yields move in opposite directions. Then, we set up our own more sophisticated factor model which defines the flight to quality as the effects of a factor which results in opposite movements in German and crisis-country government bond yields. We do this in order to explicitly extract movements in Euro area bond yields which are compatible with our definition of the flight to quality,

¹In what follows, the term “German government bonds” comprises all bills and bonds issued by the German federal government. Also, these figures do not represent total savings or costs to German taxpayers, since these calculations do not include effects of the crisis on German output or on implicit intra-European transfers, nor do they include indirect effects on future fiscal policy.

while recognizing that bond spreads are an imperfect proxy. Finally, we estimate an alternative factor model based on that of Diebold et al. (2006). That model links the yield curve to a set of macro factors (including an area-wide interest rate), while residual movements in German bond yields contain the effects of the flight to quality. We find that all three modeling strategies deliver similar results to each other, which suggests that they capture the same phenomenon. These estimates all indicate that the flight to quality can account for some of the fall in German bond yields, particularly at longer maturities, though the persistence of these effects is uncertain. When these effects are allowed to be persistent, then the flight to quality can account for an approximately EUR 23 to 29 billion cumulative direct reduction in interest payments for the German federal government since the onset of the crisis, on the order of about one percent of annual German GDP.

We base our definition of the flight to quality on definitions used in the recent literature.² One of the first explicit definitions of the flight to quality (or the “flight to safety”) is provided by Baele et al. (2014). Baele et al. straightforwardly define flights to safety as occurring on days when normalized returns to safe assets and the returns to risky assets both move in opposite directions, beyond a certain threshold. Based on this definition, they measure flight-to-safety episodes in stock and government bond returns over 23 countries, for a 32-year period ending in January 2012. They go on to discuss in detail how the returns of specific asset classes perform during these episodes, finding a particularly strong effect on the returns to riskier or more procyclical assets, and also on commodity prices. However, the application of their methodology to our dataset runs into some problems. The first problem is that when we apply their methodology, we find at most a small number of flight to quality episodes (4 at a one-year maturity), and the chronology of episodes is not consistent across maturities or the choice of a weekly vs. daily data frequency. Secondly, their definition selects for days on which German bond yields decrease strongly, which is not a problem when doing a descriptive analysis (as they do), but

²There are other possible definitions. For instance, Blatt et al. (2015) define the flight to quality as “the situation in which the level of interdependence decreases during a crisis, indicating a decoupling of safe haven countries that are believed to be unaffected by the crisis”, where a safe haven is defined as an asset that gains value during times of high uncertainty (Ranaldo and Söderlind (2010)) or that performs well under extreme negative market conditions (Baur and McDermott (2010)). Baele et al. (2014) give a good overview of these other definitions, and we discuss some of these phenomena in our results.

which is a problem when doing a causal analysis (as we do). In fact, based on our causal analysis, we find that flights to quality tend to show up as large movements in bond spreads or in crisis country bond yields but as small movements in German bond yields. Furthermore, our flight-to-quality factor is continuous-valued and symmetric, in that it captures episodes of normalization in which (all else equal) German bond yields rise and crisis-country bond yields fall.

In related work, Ehrmann and Fratzscher (2015) define the flight to quality during the Euro crisis as a situation in which yields move in different directions for German bonds and the bonds of the crisis countries. This definition allows for movements due to a preference for higher liquidity, a re-pricing of risk, or any other factors which would cause yields to move in different directions. Based on this definition, Ehrmann and Fratzscher estimate a vector autoregressive model using volatility breaks as an identification strategy, in order to decompose movements in Euro area bond yields by country-specific origins. They find an economically and statistically significant effect of shocks to crisis country bond yields on German yields during and, to some extent, after the crisis, but less so before the crisis. They also find a certain degree of “decoupling” of Italy and Spain from the rest of the Euro area during the crisis. While their identification strategy differs from ours, our identification strategy uncovers some of these same things, in addition to a possible more recent “recoupling” of Ireland with the rest of the non-crisis countries.

Moving beyond the flight to quality per se, there is additional evidence on the effects of liquidity and spillovers on German bond yields during the Euro crisis. For instance, De Santis (2014) finds a role for the “flight to liquidity” in pushing German bond yields downward during the early days of the crisis, based on a comparison between German central government bonds and less-liquid German agency bonds. DeSantis goes on to calculate spillover effects from the crisis to Euro area bond spreads, finding that news about the solvency of the crisis countries (particularly Greece) is related to large spillover effects for some countries, with smaller and less-precise effects for countries such as Austria, Finland, and the Netherlands. While De Santis does not concentrate on the effects of the crisis on German government bond yields per se, our results indicate that the flight to liquidity and spillovers are

empirically relevant phenomena. Furthermore, we obtain the same results for some of the other non-crisis countries. These results suggest that the flight to quality is a Euro area-wide phenomenon and not just a phenomenon that is limited to Germany. Furthermore, these results suggest that our factor approach can deliver a characterization of the “flight to quality” that is in line with other characterizations found throughout the literature.

2 Data and methodology

2.1 The path of German bond yields and yield spreads

We base our analysis on a dataset which contains weekly bond yields from Datasream, at maturities of 1, 2, 5, 7, and 10 years. We use weekly data in order to focus on a lower frequency than that implied by daily data; results derived using monthly data give similar results. The dataset covers Germany, the crisis countries (Portugal, Ireland, Italy, Greece, and Spain), and the other original Euro area countries (Austria, Finland, France, Belgium, and the Netherlands), from the beginning of January 2007 through the second week of October 2015. The dataset covers eleven countries altogether. In Germany and in the other non-crisis countries, these yields show a downward path over time, as shown for Germany in Figure 1. Our main question is to figure out how much of this downward path is driven by the sovereign debt crisis and the flight to quality, rather than by other determinants of bond yields.

As our main reduced-form crisis proxy, we use the average arithmetically-weighted bond spread between the crisis countries and Germany. These are shown in Figure 2. The main thing to note about this crisis proxy is that most of the crisis occurs in 2011 and 2012, with a slight relapse in 2015. Furthermore, there is a massive downward level shift in the yield spread between the weeks of March 5 and March 12, 2012. This level shift coincides with the restructuring of the Greek debt, and it represents a significant outlier in our dataset. Since this outlier is a special one-time event which accounts for the majority of variation in bond spreads, we control for this outlier in all that follows.

2.2 Calculating direct effects on fiscal balances

2.2.1 Basic methodology

To calculate the effects of the flight to quality induced by the crisis on German fiscal balances, it is necessary to estimate what German bond yields would have been in the absence of the crisis. For instance, Dany et al. (2015) assume as one of their baseline counterfactuals that the German yield curve would have remained at its 2000-2007 average from 2010 onward. Then based on the differences between observed bond yields and this counterfactual, they calculate the difference in interest payments on German government bonds issued between 2010 and 2015. They find a difference of about EUR 93 billion. Similar approaches by Boysen-Hogrefe and Filipczak (2015) or the numbers cited in Deutsche Bundestag (2015) come up with similar figures. Furthermore, Dany et al. (2015) set up counterfactuals in which an independent German monetary policy were to follow a Taylor rule, targeting German inflation and GDP. Based on these counterfactuals, they find a difference up to EUR 126 billion.

While we focus on a different set of counterfactuals in which the flight to quality (rather than the recession and crisis themselves) would have never happened, we start by adopting the same general approach, applied to the same data series for the same time periods (January 2010 through July 2015, which we later extend to January 2007 through October 2015). These data series are provided by the Federal Financing Agency *Deutsche Finanzagentur* (2015), in spreadsheet form. These spreadsheets contain information on the date of issue, date of maturity, volume, price, and yield of each bond issued by the German federal government during this period. As such, these spreadsheets do not contain data on bonds issued by state and local authorities, and so the results from this exercise are in some sense somewhat conservative. Still, we adopt this approach because it facilitates a comparison between our results and previous results.

In what follows, we derive the effects of the differential between actual bond yields and counterfactual bond yields in a similar manner. Namely, first, we multiply this differential by the value of a given bond (volume times price relative to par),

to derive the difference in interest payments between observed and counterfactual environments. We do this for every bond issue, for every year that the bond is outstanding.³ If the bond is outstanding for only part of a given year, we weight the interest payment by the number of days out of that year that the bond is outstanding. Then we sum up these differences in interest payments across issues and across years, to arrive at the effect of a given path of bond yields on the interest payments of the German federal government.

It is worth noting that we do not attempt to model the endogenous response of fiscal policy or overall indebtedness to this differential. While this would be a worthwhile thing to do, this would considerably complicate the current analysis, and this would make our results incomparable to other results. As a result, we consider this a useful avenue for future research.

2.2.2 Replicating previous estimates

In order to see how our methodology lines up with previous work, we check to see what would have happened if the change in bond yields the beginning of 2007 and the following years had not occurred. This serves as a check on our results. To do this, we assume that the interest rate differential for a given bond issue is given by the cumulative weekly changes in the yield curve since the beginning of 2007, based on data from Datastream for government bond yields of a maturity of 1, 2, 5, 7, or 10 years. For bond yields in between those maturities, we interpolate bond yields from the nearest maturities. For bond years greater than 10 years or below 1 year, we use those maturities, respectively.

When we do this for every bond issue for every year in the sample, we come up with a total cumulative effect on interest payments (from January 2010 to the end of 2015) of approximately EUR 88.0 billion. This also holds true when the dataset is extended from January 2007 to the first week of October 2015—the resulting figure is then EUR 91.0 billion. This is close to the EUR 93 billion figure originally reported by Dany et al. (2015) and the EUR 94 billion figure reported by the Bundestag

³We omit bonds denominated in U.S. dollars; these represent an insignificant share of bond issuances.

(2015).⁴ This implies that if the yield curve were to have remained at its 2007 levels, the German government would have faced pressure to run noticeably larger deficits over time, all else equal. That said, we reiterate that these are the direct benefits accruing to the German government and taxpayers from low bond yields, and not the benefits from the flight to quality per se. To estimate the direct effects of the flight to quality, we first have to estimate how large the flight to quality actually was, and then to use these estimates to derive a new set of counterfactuals.

3 Estimating counterfactual interest rate paths and their fiscal consequences

To disentangle the effects of the flight to quality, we offer three sets of counterfactuals that are compatible with our definition of the flight to quality: one counterfactual based on a regression approach, and two counterfactuals based on different factor models. All three of these approaches yield an effect of the flight to quality about one quarter to one third as large as the overall effects of low bond yields.

3.1 A counterfactual based on a regression approach

First, we set up a simple regression approach, which allows us to see how German bond yields are related to bond spreads. We start with a fairly general autoregressive model, which is equivalent to one line in a VAR:

$$r_t = k + \sum_{p=0}^P \alpha_p s_{t-p} + \sum_{p=1}^P \rho_p r_{t-p} + \sum_{p=0}^{P-1} \gamma_p d_{t-p} + \varepsilon_t. \quad (1)$$

This model relates a German yield r_t to contemporaneous and lagged average yield spreads across the crisis countries (Portugal, Ireland, Italy, Greece, and Spain) given by s_{t-p} , lagged German yields r_{t-p} , a dummy for the March 2012 debt restructuring given by d_{t-p} , and a white noise residual ε_t . This model structure implicitly orders the shock to yield spreads before the shock to German bond yields. The parameters governing the model are a constant k and regression coefficients α_p , ρ_p , and γ_p ; these

⁴When we experiment with different timing and pricing conventions, our results remain within this general range.

parameters need to be estimated.

In what follows, we start by fixing the lag length P at one week. This is based on the results from an OLS estimation of the model with a lag length P of two weeks (results available upon request). Based on such an estimation, we cannot reject the null hypothesis that the coefficients α_2 , ρ_2 , and γ_1 jointly equal zero, at separate bond maturities of 1, 2, and 10 years. Meanwhile, the p values for this null hypothesis at maturities of 5 and 7 years are 0.026 and 0.016. However, given the results for the other maturities and given concerns with multiple hypothesis testing (which tends to deflate observed p values relative to an ideal test), we consider it reasonable to adopt a lag length P of 1 week.

Next comes the issue of a possible unit root in German government bond yields. This is an issue because any estimation strategy that does not take unit roots into account would yield biased estimates, which in turn would bias our analysis of German interest payments. In fact, based on the visual evidence contained in Figure 1 and based on bootstrapped results from an ADF unit root test, we cannot reject the null hypothesis of a unit root in German government bond yields for any of the maturities. In fact, the smallest p value against the joint null hypothesis of $k = 0; \rho = 1$ is at a maturity of one year, with an estimated coefficient ρ of just above 0.995, an F statistic of 2.49, and two-sided p value of 0.288. Furthermore, the estimated value of ρ is larger than the expected estimate of ρ under the assumption of a unit root. Based on these findings, we argue that there is sufficient evidence to treat r_t as exhibiting a unit root with no drift. Furthermore, we find that an estimated unit root model with one lag of Δr_t and two lags of s_t yields a coefficient on Δr_{t-1} that is economically and statistically indistinguishable from zero but is incompatible with a unit root in Δr_t . This implies that a higher order of integration is highly unlikely.

Since the evidence is in line with a unit root, we re-cast equation (1) in terms of changes in r_t , such that:

$$\Delta r_t = \sum_{p=0}^P \alpha_p s_{t-p} + \sum_{p=0}^{P-1} \gamma_p d_{t-p} + \varepsilon_t. \quad (2)$$

Results from this estimation are included in Table 1. First of all, the estimated coefficients α_p indicate that an increase in the yield spread has had a negative contemporaneous effect on German yields, although the yield spread only accounts for a small share of changes in German yields. Based on these estimated effects, a counterfactual scenario that holds yield spreads at their levels as of the beginning of 2007 (also implying no rescue actions) would yield a change in German interest payments of about EUR 23.4 billion between 2007 and 2015, with almost all of the effect coming after 2010. (These counterfactual yields are shown in Figure 3 for one and ten year horizons.) This number is equal to about one quarter of the original figure of EUR 91.0 billion. Furthermore, across all maturity bands, an F test against the null hypothesis that the coefficients α_p jointly sum to zero gives a minimum two-sided p value of about 0.135. This implies that it is not possible to statistically distinguish the long-run effect of the crisis on bond yields from zero, though the current point estimates represent the best possible estimate of this effect.

Motivated by this, we then set up a version of equation (2) which imposes no long-run effect. This serves as a robustness check to our baseline results. This setup takes the form:

$$\Delta r_t = \sum_{p=0}^{P-1} \delta_p \Delta s_{t-p} + \sum_{p=1}^{P-1} \gamma_p d_{t-p} + \varepsilon_t. \quad (3)$$

Estimates of this equation can be found in Table 2. Crucially, these estimates are similar in magnitude to the estimates in Table 1 with respect to the effects of bond spreads on German government bond yields. Furthermore, a counterfactual scenario based on these estimates since the beginning of 2007 would yield a change in German interest payments of about EUR 15.0 billion between 2007 and 2015, with almost all of the effect coming after 2010. These counterfactual yields are shown in Figure 4 for the one and ten year horizons—at these horizons, the estimated effects of the crisis do not change substantially, while at some horizons (particularly the 5-year horizon, not shown), the estimated effects fall somewhat in absolute value. In turn, this implies that the estimated effects of the crisis are somewhat (but not highly) sensitive to whether or not that crisis is allowed to have lasting effects on German government bond yields.⁵

⁵Appendix A presents estimates for a dataset that excludes Greece. The estimates from such a dataset are broadly similar.

3.2 A factor model approach

3.2.1 The factor model

Next, we set up a factor model which aims to extend the reduced-form regression approach in two ways: by allowing government bond returns to co-move in a time-varying manner with time-varying idiosyncratic volatility, and by more explicitly separating common co-movements in Euro area bond yields (which are not influenced by the flight to quality) from common movements in opposite directions, which capture the flight to quality. The structure of this factor model is explicitly motivated by our definition of the flight to quality, and this model can also disentangle a crisis factor from a common Euro-area-wide factor. To implement these features, we set up a dynamic factor model with time-varying factor loadings and time-varying idiosyncratic variances. We then estimate this model using Bayesian techniques. We find that this model confirms the basic picture painted by the regression estimates, while providing more detail about other phenomena such as decoupling, recoupling, and country-specific factors.

The model uses the following notation. First, $y_{i,t}$ denotes the change in returns $\Delta r_{i,t}$, where $i = 1, \dots, n$ denotes the country, and $t = 1, \dots, T$ is a time index. Secondly, the dynamic factor model features K factors f_t with factor loadings $b_{i,t}$ and c and idiosyncratic components $e_{i,t}$, with the observation equation:

$$y_{i,t} = b_{i,t}f_t + cd_t + e_{i,t}. \quad (4)$$

The dummy d_t indicates the discontinuity associated with the rescue package, which is treated as exogenous. The factors f_t follow a VAR law of motion, with one lag and no constant, such that:

$$f_t = \Phi f_{t-1} + \epsilon_t, \quad (5)$$

where ϵ_t is a standard multivariate Gaussian white noise process and Φ is a $K \times K$ matrix. The idiosyncratic components $e_{i,t}$ follow a first-order autoregressive process with stochastic volatility for the innovations, such that:

$$e_{i,t} = \rho_i e_{i,t-1} + u_{i,t} \sigma_{i,t}, \quad (6)$$

where $u_{i,t}$ is white noise with a standard Gaussian distribution and $\sigma_{i,t}$ is a set of time-varying idiosyncratic volatilities. The factor loadings all follow independent random walks, such that:

$$b_{i,t} = b_{i,t-1} + \Sigma_i^{1/2} w_{i,t}, \quad (7)$$

where Σ is a diagonal $K \times K$ matrix with diagonal elements $\varsigma_{i,k}^2$, and $w_{i,t}$ is white noise with a standard Gaussian distribution. Finally, the idiosyncratic volatilities $\sigma_{i,t}$ follow the law of motion:

$$\log \sigma_{i,t} = \alpha_{i,0} + \alpha_{i,1} \log \sigma_{i,t-1} + \nu_{i,t}, \quad (8)$$

where $\nu_{i,t}$ follows a standard Gaussian white noise process with a mean of zero and a standard deviation of $\sigma_{i,\nu}$.

We estimate the model using a Gibbs sampler, with an augmented Metropolis-Hastings step to estimate the time varying volatilities. Details on the augmented Gibbs sampler, along with the weakly informative prior distributions, are provided in Appendix B. For inference we take 15,000 draws after an initial burn-in period of 5,000 draws. We conduct the estimation separately for each maturity (1, 2, 5, 7, and 10 years).

Our estimated model has three factors, to which we apply the following identifying assumptions. The first factor exhibits positive loadings for all countries at all times. This factor captures common co-movements in the Euro area yield curve. The second factor loads positively on Germany and negatively on each of the crisis countries (Greece, Italy, Spain, Portugal, and Ireland). This factor captures the flight to (or from) quality during and after the crisis. The third factor has a zero loading for Germany and negatively for all of the crisis countries. This factor captures co-movements in the Euro area yield curve that are orthogonal to movements in the German yield curve. Altogether, this factor model provides more structure than a simple regression or VAR model would provide, and in particular, this factor model distinguishes between Euro-area-wide movements in the yield curve and movements in the yield curve that are related to the crisis.

3.2.2 Model estimates

The three estimated factors are shown in Figure 5, for the model at a five-year horizon. (Other horizons are available from the authors upon request.) Factor loadings are shown in Figure 6 (for the non-crisis countries) and Figure 7 (for the crisis countries). The distributions of coefficient estimates at different horizons are available upon request; these distributions look broadly similar to those at a 5-year horizon.

The first factor (top panel) drives common movements in European bond yields. This factor shows an especially large downward movement during the peak years of the crisis in 2011 and 2012. What this implies is that some portion of the fall in German government bond yields during that time is due to common factors within Europe (such as the ECB’s monetary policy and Euro-area-wide financial market conditions), and not just due to the flight to quality. In addition, the influence of this factor has fallen more slowly for the crisis countries (except Ireland) than for some of the other countries, while it has fallen more quickly for the non-crisis countries. This indicates that the effects of monetary policy and/or Euro-area-wide financial market conditions have remained somewhat more important for the crisis countries than for the non-crisis countries. It is also worth noting that this divergence captures the same type of “decoupling” found by Ehrmann and Fratzscher (2015) over this time period.

The second factor (middle panel) depicts the crisis factor, which again showed heightened volatility between 2009 and the first half of 2013. All of the non-crisis countries except for Belgium load neutrally or positively onto this factor, particularly during the height of the crisis in 2011-12. Meanwhile, all of the crisis countries load neutrally or negatively by construction—Greece, Spain, Italy, and Portugal more so, Ireland less so. In fact, the crisis appears to have more or less stopped affecting Irish bond yields sometime by the beginning of 2012, when Irish bond yields began to “recouple” with the rest of Europe. Again, this captures some of the “decoupling” of Italy and Spain from European capital markets, but this also indicates that the benefits from the flight to quality have not only accrued to Germany, but to the non-crisis countries more generally. This is a phenomenon originally described by De Santis (2014), and this implies that a naive comparison between German and

other non-crisis country bond yields would yield a downward-biased estimate of the absolute effects of the flight to quality on German government bond yields. Furthermore, this means that the flight to quality is not just a flight to liquidity, since this affects small countries.

The third factor (bottom panel) drives common movements in European bond yields outside of Germany that are not correlated with German bond yields. For the non-crisis countries outside of Germany, this non-German factor is not in general an economically or statistically significant factor behind movements in bond yields, apart from an episode in Belgium during the crisis. This implies that this factor is in fact not related to the flight to quality, in line with the identifying assumptions of the factor model. However, for some of the crisis countries (mainly Italy, Portugal, Spain, Greece, and temporarily, Ireland), this is an economically significant factor. Again, this factor exhibits significant “recoupling” for Ireland—by assumption it is not allowed to go to zero, but in comparison with the other crisis countries, this factor falls in importance over time.

3.2.3 A counterfactual path for German government bond yields

Given these estimates, we go on to estimate a counterfactual path for German government bond yields by taking the posterior mean effect of the second factor on German government bond yields, cumulatively since 2007, across all maturities. These paths are shown in Figure 8, for maturities of 1 and 10 years. These estimates track the main regression-based estimates that can be found in Figure 3, with a slightly smaller effect at the shorter end of the yield curve. This implies that the simpler regression-based estimates pick up the main effects of the flight to quality. Furthermore, the counterfactual paths for bond yields implied by the factor model would indicate that the German taxpayers have saved approximately EUR 26.4 billion, cumulatively, from 2007 to 2015. These estimates imply that the flight to quality accounts for a bit more than one quarter (EUR 26.4 billion / EUR 91.0 billion) of the total savings to German taxpayers as a result of the current low interest rate environment. Furthermore, the flight to quality accounts more than all of the savings that are identified using the regression-based approach (EUR 23.4 billion).

3.3 Results from an alternative yield curve model

Finally, we estimate an alternative yield curve model for German government bonds, as a robustness check. This model is based on a model given by Diebold et al. (2006), based in turn on a model by Nelson and Siegel (1987). This model links the yield curve to other macro-financial variables. Given this setup, our approach follows Boysen-Hogrefe (2012), who estimates the effects of the flight to quality based the assumption that residual movements in German government bond yields represent the flight to quality, given realizations of macro-financial variables.

Following Nelson and Siegel (1987) the model for the yield curve takes the following form, whereby yields r_t^τ with maturity τ depend on the three factors given by L_t (level), S_t (slope), and C_t (curvature) such that:

$$r_t^\tau = L_t + S_t \left(\frac{1 - \exp\{-\lambda\tau\}}{\lambda\tau} \right) + C_t \left(\frac{1 - \exp\{-\lambda\tau\}}{\lambda\tau} - \exp\{-\lambda\tau\} \right) + u_t^\tau. \quad (9)$$

As in Diebold et al. (2006) the three estimated yield curve factors are augmented by three macroeconomic factors. These factors are the Euro area money market interest rate (3 month Euribor), the GDP deflator, and the output gap (both quarterly, output gap provided by Oxford Economics, interpolated to a monthly frequency). These six factors are assumed to follow a VAR(1).

We assume λ as in Diebold et al. (2006) to be 0.49. Given λ , we extract the factors L_t , S_t , and C_t via a linear transformation of the yields. We then estimate the resulting VAR using OLS. Then we put the model in state-space form, and conditional on the macroeconomic variables, we form projections for L_t , S_t , and C_t . The residuals between projected yields at maturity τ and realized yields represent the estimated effects of the flight to quality.

These counterfactual yields for maturity $\tau = 1$ and $\tau = 10$ are plotted in Figure 9. These counterfactual yields track the counterfactual yields from the baseline regression model and from the other factor model relatively closely. In fact, the estimated savings to German taxpayers based on this counterfactual yield path was EUR 29.0 billion from 2007 to 2015, which is also near the EUR 23.5 billion figure

derived from the baseline factor model as well as the EUR 23.4 billion from the regression model. This suggests that all three models capture similar phenomena.

4 Conclusion

When we estimate the effects of the flight to quality on German government bond yields using a regression model and two different factor models based on a common definition of the flight to quality, we consistently find that the flight to quality accounts for a significant share of the fall in German government bond yields since 2007. These effects are larger at longer maturities (5-10 years), where counterfactual bond yields are relatively close to their 2007 levels. However, at shorter maturities (1-2 years), we find counterfactual government bond yields closer to their current levels. Taken together, these counterfactual paths imply that German taxpayers have cumulatively saved approximately EUR 23 to 29 billion from the direct effects of flight to quality on interest payments since 2007, with our preferred estimates coming in at EUR 26.4 billion. This figure is of marginal economic significance (well below one percent of annual German GDP) and it is significantly smaller than the overall effects of the low interest rate environment (approximately EUR 91 billion).

These results exhibit a considerable degree of consistency across all three models, which indicates that these models all capture a similar phenomenon. Since Euro area bond spreads and our estimated second factor capture the same episodes, this means that Euro area bond spreads function relatively well as a reduced-form crisis indicator. On the other hand, the factor model provides more information than what a simple regression model would provide—namely, the “decoupling” of Italy and Spain discussed by Ehrmann and Fratzscher (2015), the “recoupling” of Ireland, the spillovers into other non-crisis countries described by De Santis (2014), and the ways in which Belgium does not completely resemble the other non-crisis countries. Also, the model indicates that the flight to quality is likely not just a flight to liquidity, since the smaller, less liquid non-crisis countries (except for Belgium) also participated in the flight to quality.

Given the success of the factor model in matching these results, this also suggests

directions for future work. For instance, future work can help to reconcile our factor approach with the threshold approach of Baele et al. (2014) over their dataset. While both approaches are designed to measure the same thing, these approaches seem to differ from each other in practice, and future work can help to show what drives these differences. Based on our results, some of these differences are likely due to the small effects of the flight to quality that we find, in practice, on German bond yields. Additionally, future work could look at a broader set of observables, in order to see how other asset prices have behaved over the course of the Euro crisis.

We would like to thank our colleagues for their helpful comments. We would also like to thank Geert Bekaert, Clemens Kool, Bert Smid, and the participants in the 2016 EUROFRAME Conference. This work has been financed in part by Thyssen Foundation grant Az.10.15.1.015IB. All opinions are ours and not those of our respective institutions.

A Appendix: Regression results excluding Greece

A.1 The results

Table A1 contains regression results for the baseline regression model using crisis country data excluding Greece. This implementation of the model does not contain a dummy variable for the Greek rescue package in March 2012. These regression results substantially mirror results from a dataset that includes Greece—an increase in yield spreads coincides with a fall in German bond yields, while the evidence on a permanent effect of bond spreads is ambiguous. These results indicate that the main results in the paper are not sensitive to the inclusion or exclusion of Greece.

This insensitivity also carries over to the calculated direct fiscal savings to the German government, based on the counterfactual paths of bond yields shown in Figure A1. These bond yields behave substantially like those from the other models, and if bond yields had followed their counterfactual paths, they would have resulted in a direct saving to German taxpayers of about EUR 34.5 billion. This figure is broadly in line with the estimates from the other models.

B Appendix: Bayesian estimation of the factor model

B.1 Prior distributions

The model contains the parameters: $\{\alpha_{0,i}\}_{i=1}^{11}$, $\{\alpha_{1,i}\}_{i=1}^{11}$, $\{\sigma_{i,\nu}^2\}_{i=1}^{11}$, $\{\rho_i\}_{i=1}^{11}$, Φ , $\{\{\varsigma_{i,k}^2\}_{k=1}^K\}_{i=1}^{11}$ and the latent variables $\{f_t\}_{t=1}^T$, $\{\{b_{i,t-1}\}_{t=1}^T\}_{i=1}^{11}$, $\{\{\sigma_{i,t}\}_{t=1}^T\}_{i=1}^{11}$. For these objects, we assume the following prior distributions:

- $(\alpha_{0,i}, \alpha_{1,i})' \sim N(\mu_\alpha, \Sigma_\alpha)$.
- $\sigma_{i,\nu}^2 \sim IGam(\beta_1^\nu, \beta_2^\nu)$, parameterized as the number of observations and sum of squares.
- $\text{vec}(\Phi) \sim N(0, \Sigma_\Phi)$.
- $\rho_i \sim N(0, \sigma_\rho)$.

- $\varsigma_{i,k}^2 \sim IGam(\beta_1^b, \beta_2^b)$, parameterized as the number of observations and sum of squares.

with the following prior moments:

- $\mu_\alpha = (0.8, 0.7)'$ and $\Sigma_\alpha = I$, following Tsay (2005), with some allowance for a looser prior.
- $\beta_1^\nu = 0.1$ and $\beta_2^\nu = 0.1$.
- $\Sigma_\Phi = 10 \times I$.
- $\sigma_\rho = 1$.
- $\beta_1^b = 0.1$ and $\beta_2^b = 0.1$.

We select these moments because they are weakly informative, and so they can “regularize” our posterior estimates without overwhelming them.

B.2 The MCMC algorithm

To estimate the factor model, we implement an MCMC algorithm. This algorithm cycles through the full conditional distributions of parameters and latent variables. In particular, we cycle through the following conditional distributions for the model parameters:

- $p(\rho_i | \{y_{i,t}\}_{t=1}^T, \{b_{i,t-1}\}_{t=1}^T, \{f_t\}_{t=1}^T, \{\sigma_{i,t}\}_{t=1}^T) \sim N((\mathbf{e}^*{}'\mathbf{e}^* + \sigma_\rho^{-1})^{-1}\mathbf{e}^*{}'\mathbf{e}^*, (\mathbf{e}^*{}'\mathbf{e}^* + \sigma_\rho^{-1})^{-1})$, where \mathbf{e}^* denotes the vector of residuals from $y_{i,t}$ and \mathbf{e}^* denotes the first lag. Note: both vectors are elementwise multiplied by the relevant inverse standard deviations $\sigma_{i,t}^{-1}$.
- $p(\Phi | \{f_t\}_{t=1}^T) \sim N((\mathbf{F}'\mathbf{F} + \Sigma_\Phi^{-1})^{-1}\mathbf{F}'\mathbf{f}, (\mathbf{F}'\mathbf{F} + \Sigma_\Phi^{-1})^{-1})$, where \mathbf{f} denotes the vectorized factor matrix and \mathbf{F} denotes a block diagonal matrix with K blocks containing the lagged factor matrices on the diagonal blocks, and zeros elsewhere.
- $p(\varsigma_{i,k}^2 | \{b_{i,k,t-1}\}_{t=1}^T) \sim IGam(T + \beta_1^b, \sum_{t=1}^T (\Delta b_{i,k,t})^2 + \beta_2^b)$.
- $p(\sigma_{i,\nu}^2 | \{\sigma_{i,t}\}_{t=1}^T, \alpha_{0,i}, \alpha_{1,i}) \sim IGam(T + \beta_1^\nu, \sum_{t=1}^T \epsilon_{t,i}^2 + \beta_2^\nu)$.

- $p((\alpha_{0,i}, \alpha_{1,i})' | \sigma_{i,\nu}^2, \{\sigma_{i,t}\}_{t=1}^T) \sim N((\mathbf{Z}'\mathbf{Z}/\sigma_{i,\nu}^2 + \Sigma_\alpha^{-1})^{-1}(\mathbf{Z}'\mathbf{z}/\sigma_{i,\nu}^2 + \Sigma_\alpha^{-1}\mu_\alpha), ((\mathbf{Z}'\mathbf{Z}/\sigma_{i,\nu}^2 + \Sigma_\alpha^{-1})^{-1}), \text{ where } Z \text{ is a } T-1 \times 2 \text{ matrix consisting of a constant and } \{\log \sigma_{i,t-1}\}_{t=2}^T.$
 z denotes the vector $\{\log \sigma_{i,t}\}_{t=2}^T$.

Then we cycle through the following conditional distributions for the latent variables:

- First, given all data and other parameters, the factors $\{f_t\}_{t=1}^T$, and the variances $\{\{\sigma_{i,t}\}_{t=1}^T\}_{i=1}^{11}$, the distribution of the time-varying loadings $\{\{b_{i,t}\}_{t=1}^T\}_{i=1}^{11}$ is given by a multi-move MCMC algorithm, separately for each country i . This multi-move algorithm draws, for each t separately in increasing order from 1 through T , $\{\{b_{i,t}\}_{i=1}^{11}\}$ given $\{\{b_{i,\tau} : \tau \neq t\}_{t=1}^T\}_{i=1}^{11}$, the other parameters of the model, and our observed data. If our proposed MCMC draw violates the identifying restrictions, we reject and re-draw. If we reject more than a set number of times (one thousand in our application), we impose our identifying restrictions exactly, and then we re-draw the less problematic parameters. In our calculations, this happens extremely rarely.
- Then, given all data and other parameters, the time-varying loadings $\{\{b_{i,t-1}\}_{t=1}^T\}_{i=1}^{11}$, and the variances $\{\{\sigma_{i,t}\}_{t=1}^T\}_{i=1}^{11}$, the distribution of the factors $\{f_t\}_{t=1}^T$ is given by the Kalman Filter. We draw from this distribution according to the forward-backward procedure discussed by Kim and Nelson (1999), without any modifications.
- Finally, given all data and other parameters, the factors $\{f_t\}_{t=1}^T$, and the time-varying loadings $\{\{b_{i,t-1}\}_{t=1}^T\}_{i=1}^{11}$, we draw the log time-varying variances $\{\{\sigma_{i,t}\}_{t=1}^T\}_{i=1}^{11}$ from their conditional distributions, according to the Metropolis-Hastings Algorithm as discussed by Tsay (2005).

Resulting Gibbs sweeps at a five-year horizon are shown for the non-time-varying parameters in Figures B1 through B9. In general, the model exhibits reasonably good mixing and is computationally stable. The one issue is that when the priors for the time-varying variance processes become too loose, these variances behave in unusual ways for those countries that do not exhibit much time variation in idiosyncratic variance or much idiosyncratic variance to begin with. This also makes it difficult to identify the idiosyncratic autocorrelations ρ_i . Apart from that, our estimation results are not sensitive to our priors.

Gibbs sweeps for the other parameters (like for some of the time-varying parameters) and Gibbs Sweeps at other horizons are available from the authors upon request.

References

- [1] Baele, Lieven, G. Bekaert, and K. Inghelbrecht, 2014. "Flights to Safety." Working paper, Finance and Economics Discussion Series 2014-26, Federal Reserve Board.
- [2] Baur, D.G., and T.K. McDermott, 2010. "Is gold a safe haven? International evidence." *Journal of Banking and Finance* 34, pages 1886-1898.
- [3] Blatt, D., B. Candelon, and H. Manner, 2015. "Detecting contagion in a multivariate time series system: An application to sovereign bond markets in Europe." *Journal of Banking and Finance* 59, pages 1-13.
- [4] Boysen-Hogrefe, J., 2012. "Die Zinslast des Bundes in der Schuldenkrise: Wie lukrativ ist der sichere Hafen?" *Perspektiven der Wirtschaftspolitik* 13 (Special Issue), pages 81-91.
- [5] Boysen-Hogrefe, J., and J. E. Filipczak, 2015. "Die Zinslast des Bundes in der Finanz- und Schuldenkrise: Chance und Risiko." *Kiel Fokus* 173.
- [6] Dany, G., R. Gropp, and G. von Schweinitz, 2015. "Germany's Benefit from the Greek Crisis." *IWH Online* 7/2015.
- [7] De Santis, R.A., 2014. "The euro area sovereign debt crisis: Identifying flight-to-liquidity and the spillover mechanisms." *Journal of Empirical Finance* 26, pages 150-170.
- [8] Deutscher Bundestag, 2015. *Stenografischer Bericht, 100. Sitzung*. Plenarprotokoll 18/100.
- [9] Deutsche Finanzagentur, 2015. *Auction Results*, data for years 2005-2015. Retrieved August 12, 2015 and November 13, 2015.
- [10] Diebold, F.X., G.D. Rudebusch, and B. Arouba, 2006. "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach." *Journal of Econometrics* 131, pages 309-338.
- [11] Ehrmann, M., and M. Fratzscher, 2015. "Euro Area Government Bonds—Integration and Fragmentation During the Sovereign Debt Crisis." Bank of Canada Working Paper 2015-13.

- [12] Kim, C.J., and C.R. Nelson, 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, first edition. MIT Press Books, The MIT Press, Number 0262112388.
- [13] Nelson, C.R., and A.F. Siegel, 1987. “Parsimonious modeling of yield curves.” *Journal of Business* 60, pages 473-489.
- [14] Rinaldo, A., and P. Söderlind, 2010. “Safe haven currencies.” *Review of Finance* 14, 385-407.
- [15] Tsay, R.S., 2005, *Analysis of Financial Time Series*, Second Edition. Wiley-Interscience.

Table 1: Estimates of the unit root regression model (2)

Coefficient	Description	1 yr	2 yr	5 yr	7 yr	10 yr
α_0	Yield Spread s_t	-0.0015	-0.0253	-0.0783	-0.1294	-0.1531
	(Std. err.)	0.0059	0.0076	0.0123	0.0151	0.0163
α_1	Yield Spread s_{t-1}	0.0015	0.0251	0.0768	0.1279	0.1531
	(Std. err.)	0.0061	0.0079	0.0124	0.0152	0.0165
γ_0	Rescue dummy	-0.8620	-0.9216	-0.3833	-0.3052	-1.3209
	(Std. err.)	0.3640	0.2981	0.1131	0.1057	0.1733
$H_0 : \alpha_0 + \alpha_1 = 0$	F stat.	0.00	0.03	2.24	2.08	0.00
	p value	0.958	0.855	0.135	0.150	0.990
R^2		0.017	0.028	0.087	0.143	0.165

This table presents estimates of the regression model (2), which relates changes to German yields to current and past crisis country bond spreads (in levels), and a dummy for the March 2012 outlier. The data series used is a set of weekly bond yield series for a chosen set of maturities, from January 2007 to October 2015. Source: Datastream and authors' calculations. Coefficient estimates are followed by standard errors.

Table 2: Estimates of the unit root regression model (3)

Coefficient	Description	1 yr	2 yr	5 yr	7 yr	10 yr
α_0	Δ Yield Spread Δs_t	-0.0147	-0.0257	-0.0786	-0.1294	-0.1531
	(Std. err.)	0.0052	0.0071	0.0124	0.0151	0.0162
γ_0	Rescue dummy	-0.8524	-0.9441	-0.4018	-0.3201	-1.3213
	(Std. err.)	0.3145	0.2713	0.1126	0.1054	0.1696
R^2		0.017	0.028	0.082	0.139	0.165

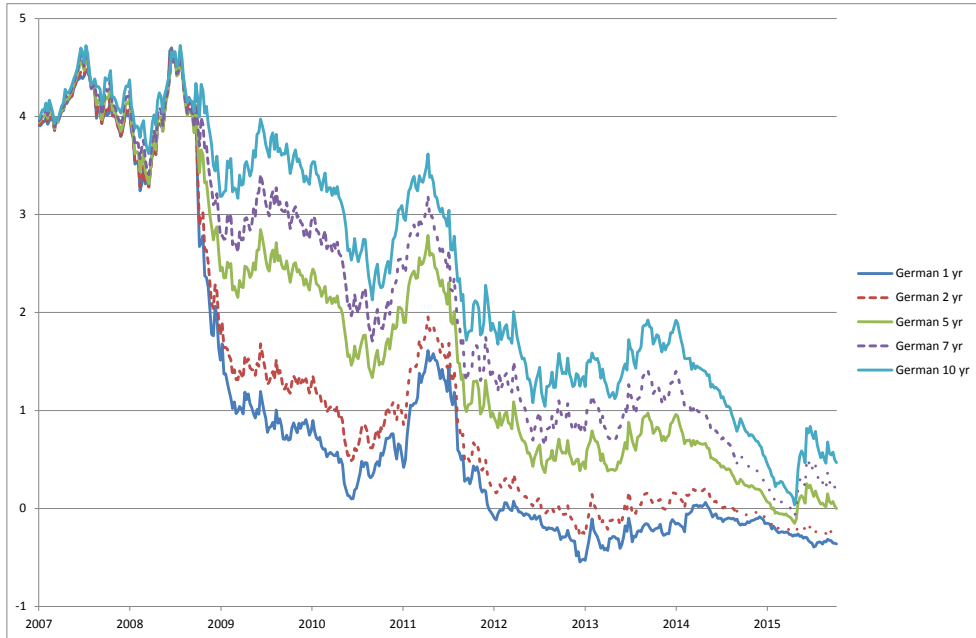
This table presents estimates of the regression model (3), which relates changes to German yields to current and past crisis country bond spreads (in first differences) and past changes to German yields, as well as a dummy for the March 2012 outlier. The data series used is a set of weekly bond yield series for a chosen set of maturities, from January 2007 to October 2015. Source: Datastream and authors' calculations. Coefficient estimates are followed by standard errors.

Table A1: Estimates of the unit root regression model, excluding Greece

Coefficient	Description	1 yr	2 yr	5 yr	7 yr	10 yr
α_0	Yield Spread s_t	-0.0651	-0.0755	-0.1164	-0.1733	-0.2506
	(Std. err.)	0.0140	0.0140	0.0159	0.0192	0.0223
α_1	Yield Spread s_{t-1}	0.0619	0.0724	0.1136	0.1706	0.2480
	(Std. err.)	0.0140	0.0140	0.0159	0.0192	0.0223
$H_0 : \alpha_0 + \alpha_1 = 0$	F stat.	3.60	4.16	3.93	3.38	2.78
	p value	0.058	0.042	0.048	0.067	0.096
R^2		0.051	0.067	0.110	0.156	0.220

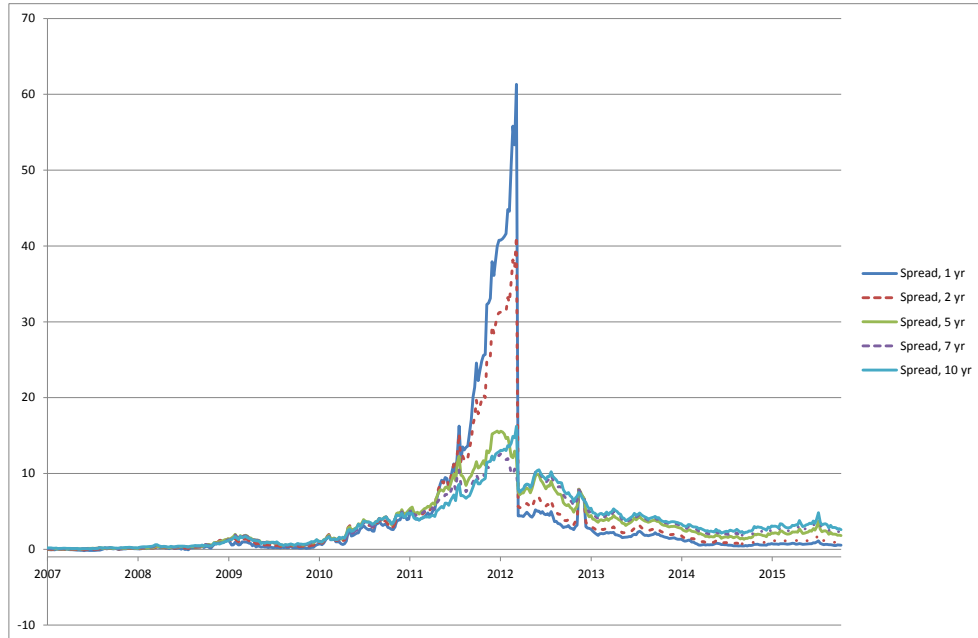
This table presents estimates of the regression model (2), which relates changes to German yields to current and past crisis country bond spreads (in levels), excluding Greece. The data series used is a set of weekly bond yield series for a chosen set of maturities, from January 2007 to October 2015. Source: Datastream and authors' calculations. Coefficient estimates are followed by standard errors.

Figure 1: German government bond yields



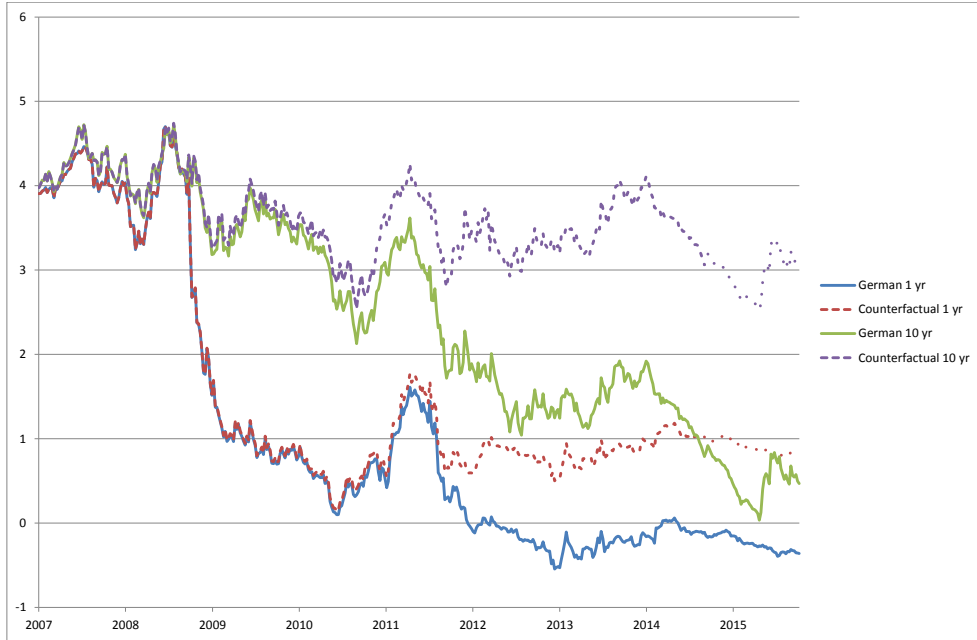
This figure shows German government bond yields at maturities of 1, 2, 5, 7, and 10 years, in percent, from 2007 to 2015. Source: Datastream.

Figure 2: Crisis country bond spreads



This figure shows crisis country bond spreads at maturities of 1, 2, 5, 7, and 10 years, in percent, from 2007 to 2015. These are given as the difference between the average bond yields of Portugal, Ireland, Italy, Greece, and Spain, and those of Germany. Source: Datastream and authors' calculations.

Figure 3: Baseline German counterfactual bond yields based on the regression equation (2)



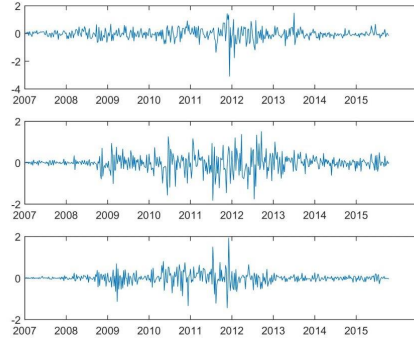
This figure shows German and counterfactual German 1-year and 10-year government bond yields, in percent, from 2007 to 2015. These are based on the effects of changes to the bond spread on German government bond yields implied by equation (2), under the counterfactual of no change since the beginning of 2007. The March 2012 rescue is also treated as part of the intervention.

Figure 4: Robustness check: German counterfactual bond yields based on the regression equation (3)



This figure shows German and counterfactual German 1-year and 10-year government bond yields, in percent, from 2007 to 2015. These are based on the effects of changes to the bond spread on German government bond yields implied by equation (3), under the counterfactual of no change since the beginning of 2007. The March 2012 rescue is also treated as part of the intervention.

Figure 5: Estimated factors (5 year maturity)



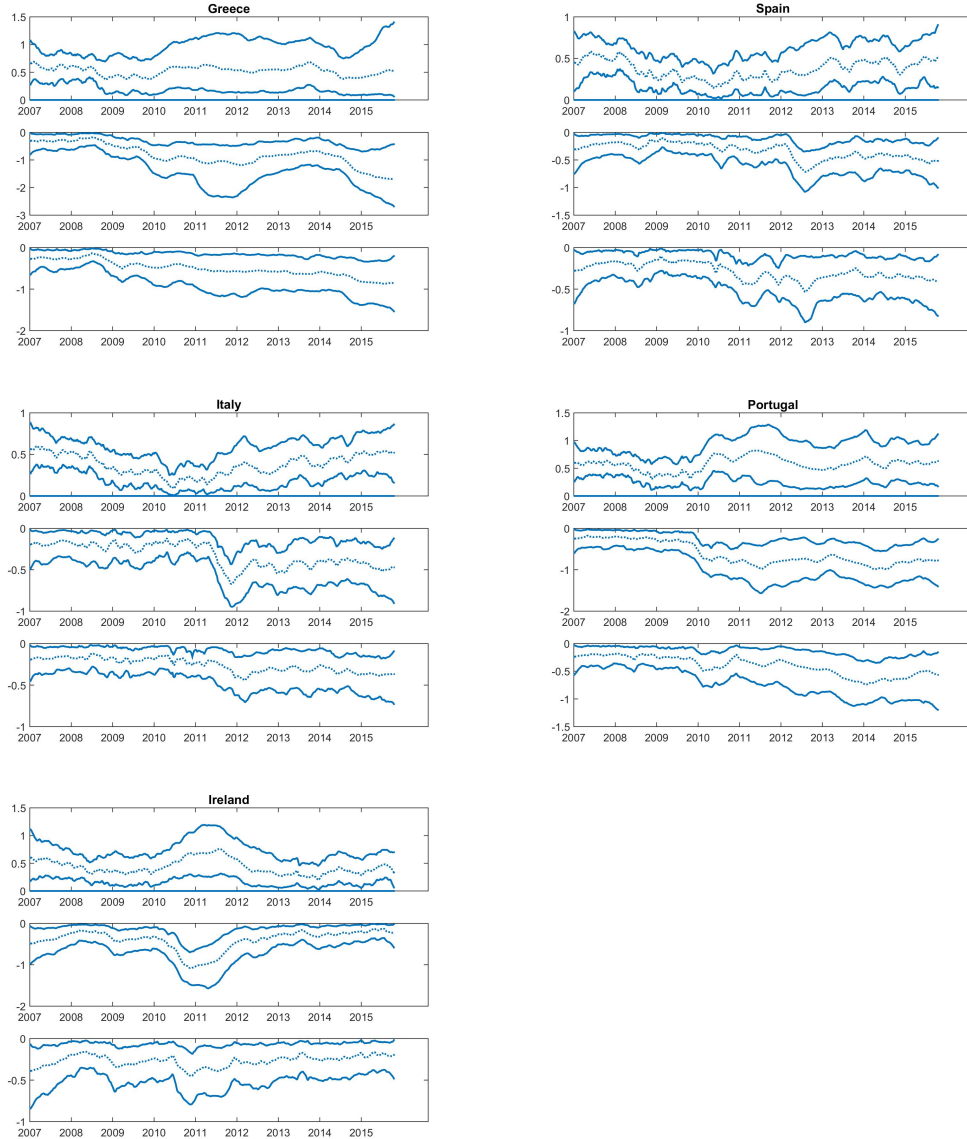
This figure shows the estimated factors. The first factor captures common co-movement in bond yields, which loads positively for every country. The second factor captures the flight to quality, which loads positively for Germany and negatively for the crisis countries. The third factor captures changes in bond yields that are orthogonal to changes in German yields, which load negatively for the crisis countries. The figure depicts results for yields with 5 years maturity only; results for other maturities are available upon request.

Figure 6: Estimated factor loadings (5 year maturity, non-crisis countries)



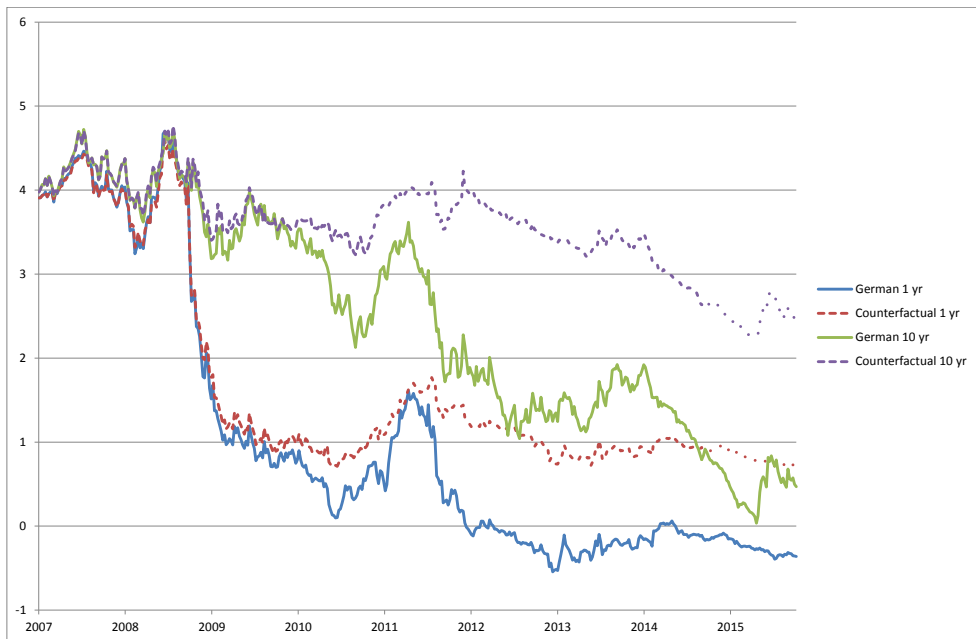
This figure shows the estimated factor loadings for a selection of countries, namely Austria, Belgium, Finland, France, Germany, and the Netherlands. For each country, the figure depicts the loadings for all three factors. The first factor captures common co-movement in bond yields, which loads positively for every country. The second factor captures the flight to quality, which loads positively for Germany and negatively for the crisis countries. The third factor captures changes in bond yields that are orthogonal to changes in German yields, which load negatively for the crisis countries. The figure depicts results for yields with 5 years maturity only; results for other countries and other maturities are available upon request.

Figure 7: Estimated factor loadings (5 year maturity, crisis countries)



This figure shows the estimated factor loadings for a selection of countries, namely Greece, Italy, Ireland, Spain, and Portugal. For each country, the figure depicts the loadings for all three factors. The first factor captures common co-movement in bond yields, which loads positively for every country. The second factor captures the flight to quality, which loads positively for Germany and negatively for the crisis countries. The third factor captures changes in bond yields that are orthogonal to changes in German yields, which load negatively for the crisis countries. The figure depicts results for yields with 5 years maturity only; results for other countries and other maturities are available upon request.

Figure 8: German counterfactual bond yields based on the baseline factor model



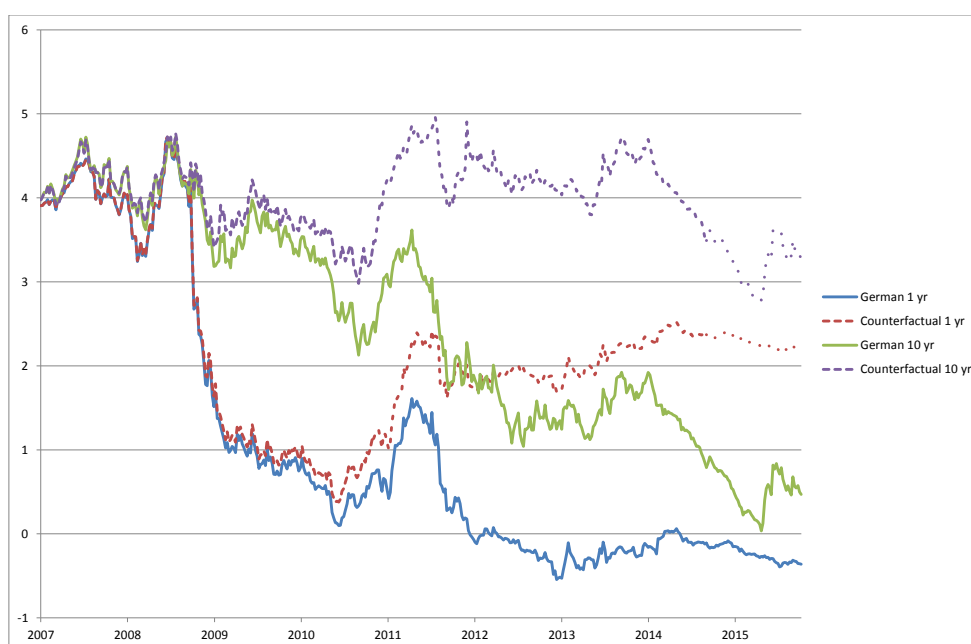
This figure shows German and counterfactual German 1-year and 10-year government bond yields, in percent, from 2007 to 2015. These are based on the effects of the crisis factor (second factor) on German government bond yields implied by the factor model, under the counterfactual of no shocks since the beginning of 2007. The March 2012 rescue is also treated as part of the intervention.

Figure 9: German counterfactual bond yields based on the alternative yield curve model



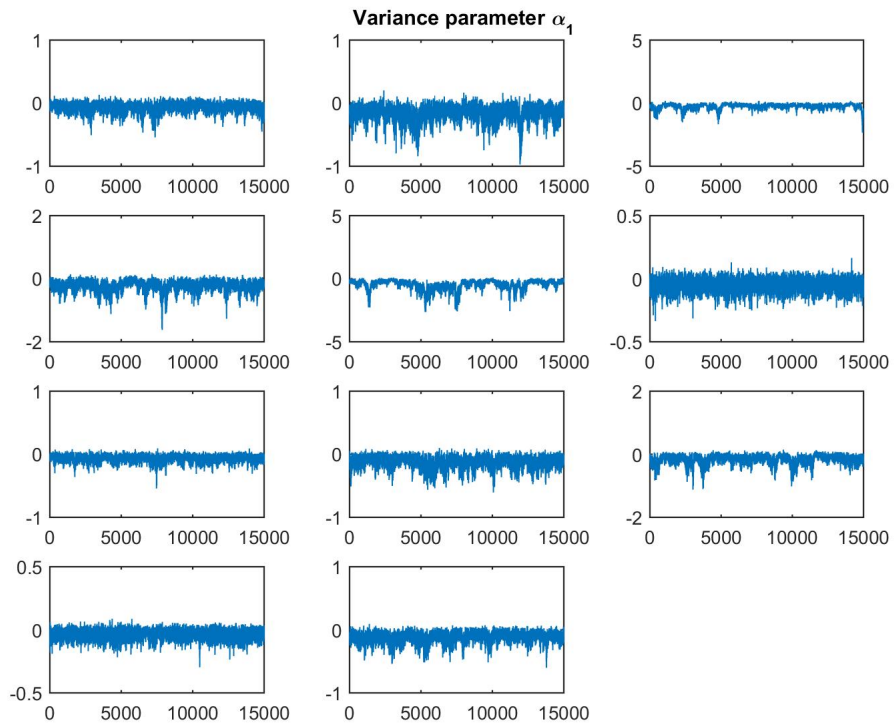
This figure shows German and counterfactual German 1-year and 10-year government bond yields, in percent, from 2007 to 2015. These are based on residual movements in German government bond yields implied by the model, under the counterfactual of no residual movements since the beginning of 2007.

Figure A1: German counterfactual bond yields, regression model excluding Greece



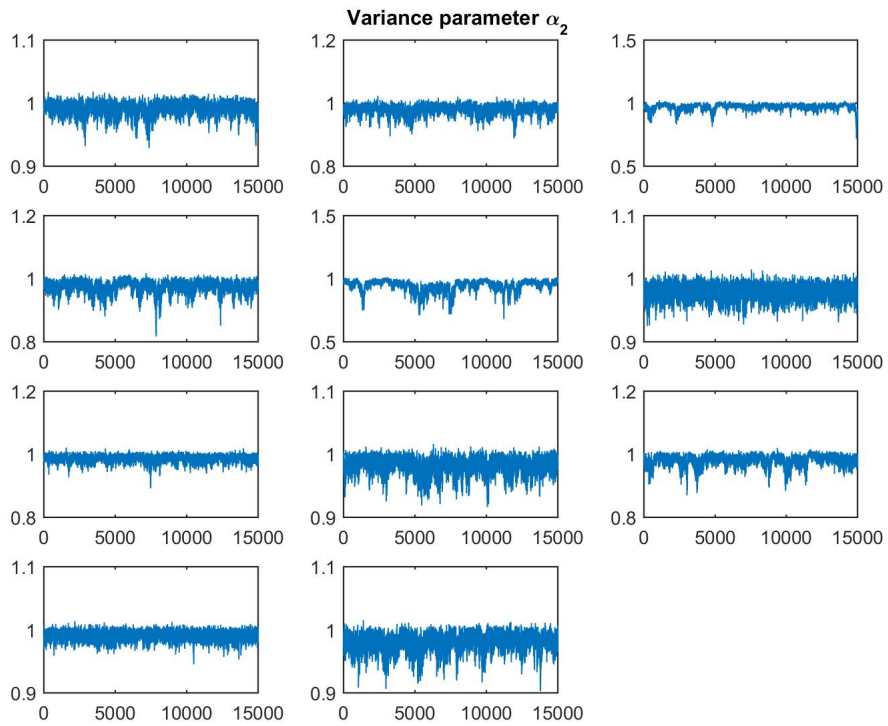
This figure shows German and counterfactual German 1-year and 10-year government bond yields, in percent, from 2007 to 2015. These are based on the effects of changes to the bond spread on German government bond yields implied by equation (2), not including a rescue dummy for March 2012, under the counterfactual of no change since the beginning of 2007. These calculations use data for the crisis countries excluding Greece.

Figure B1: Gibbs sweeps, factor model, five-year horizon



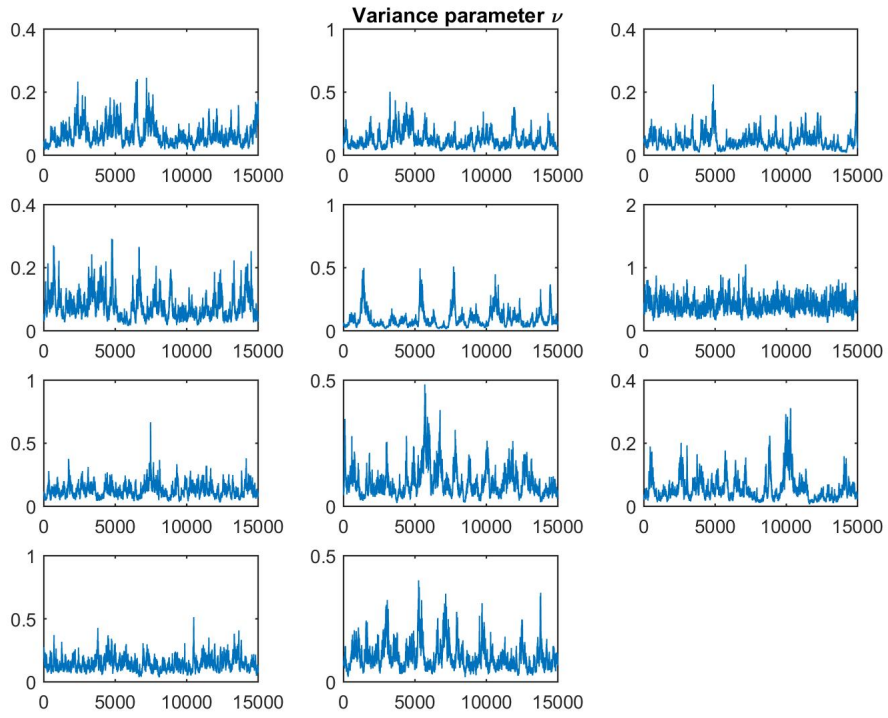
This figure shows posterior Gibbs sweeps for the parameter $\alpha_{1,i}$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B2: Gibbs sweeps, factor model, five-year horizon



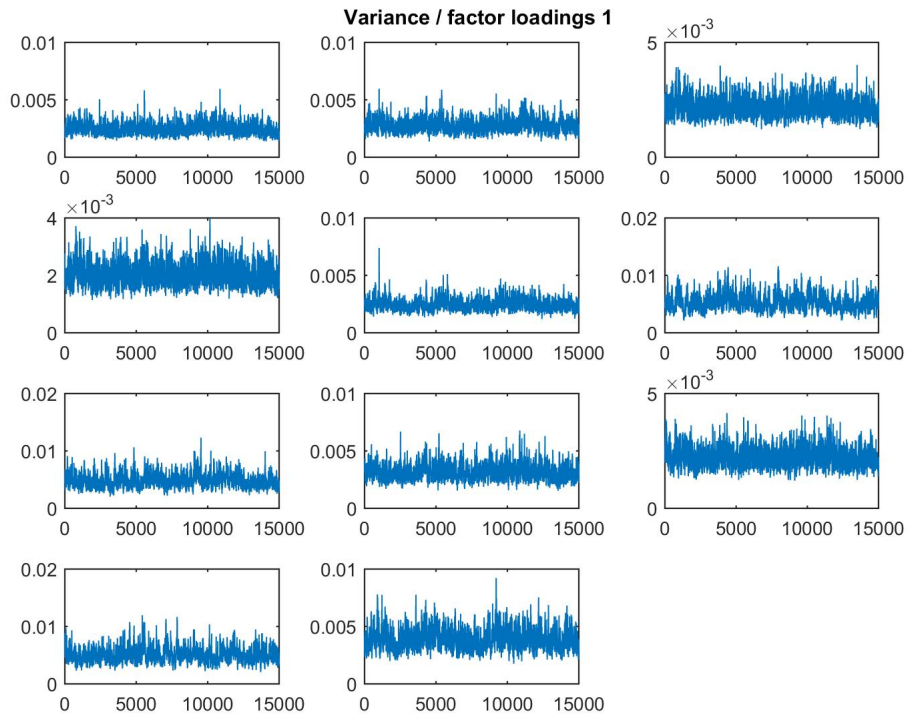
This figure shows posterior Gibbs sweeps for the parameter $\alpha_{2,i}$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B3: Gibbs sweeps, factor model, five-year horizon



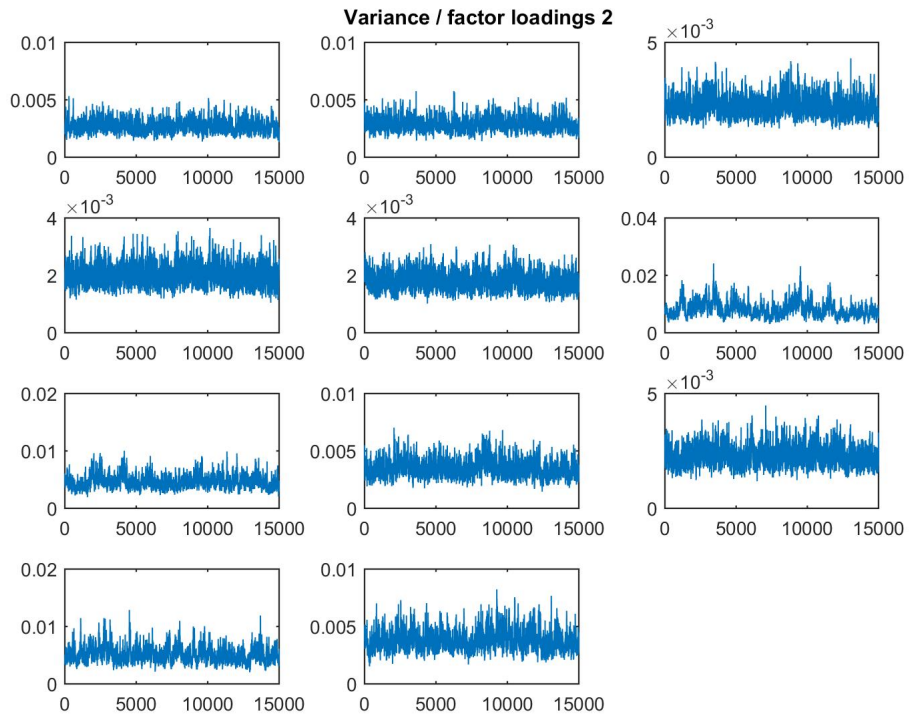
This figure shows posterior Gibbs sweeps for the parameter $\sigma_{i,\nu}^2$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B4: Gibbs sweeps, factor model, five-year horizon



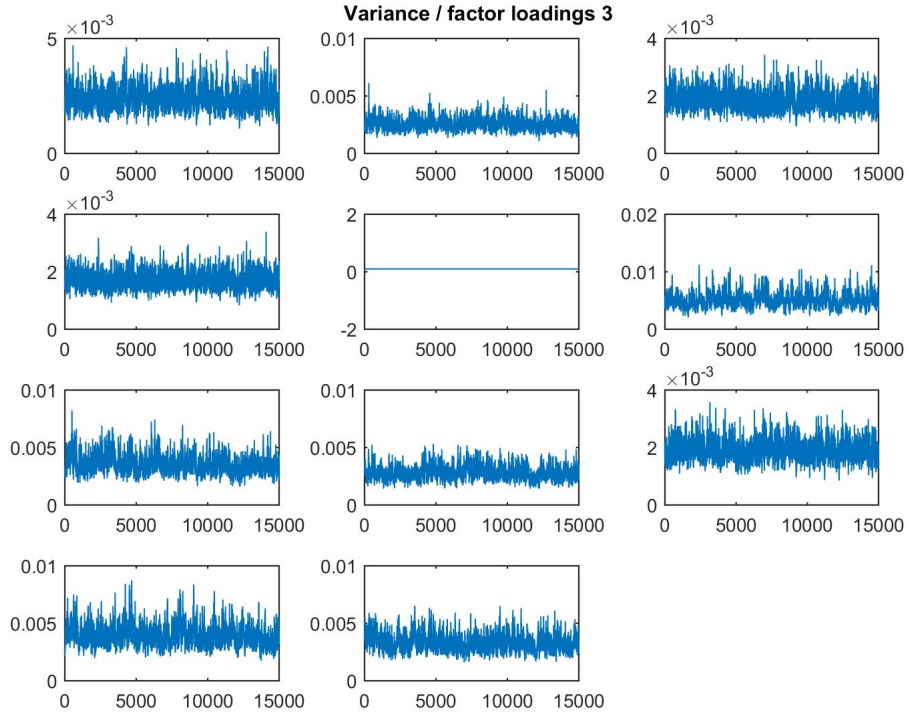
This figure shows posterior Gibbs sweeps for the parameter $\varsigma_{i,1}^2$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B5: Gibbs sweeps, factor model, five-year horizon



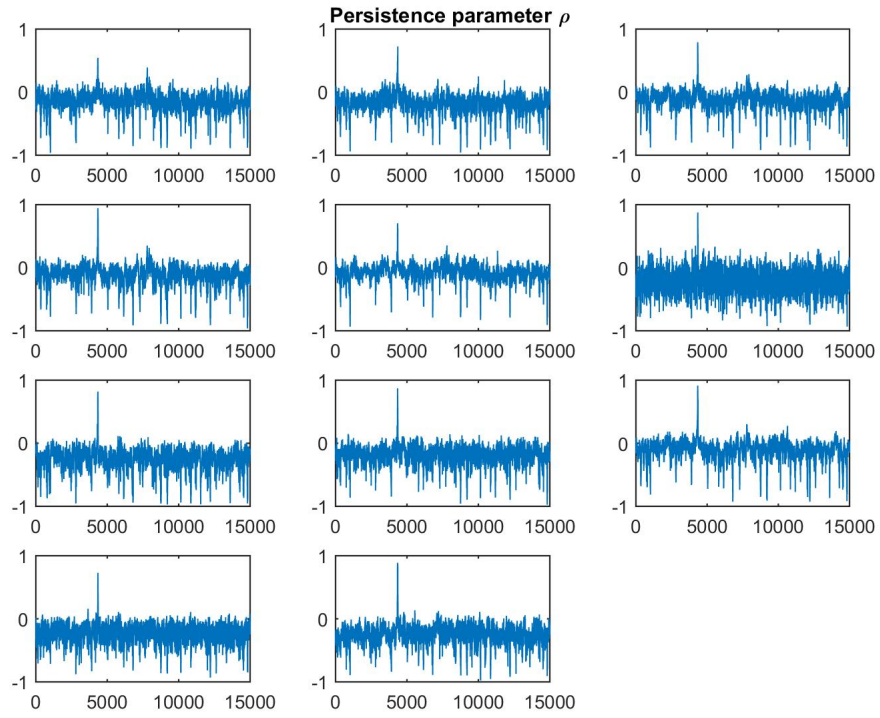
This figure shows posterior Gibbs sweeps for the parameter $\zeta_{i,2}^2$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B6: Gibbs sweeps, factor model, five-year horizon



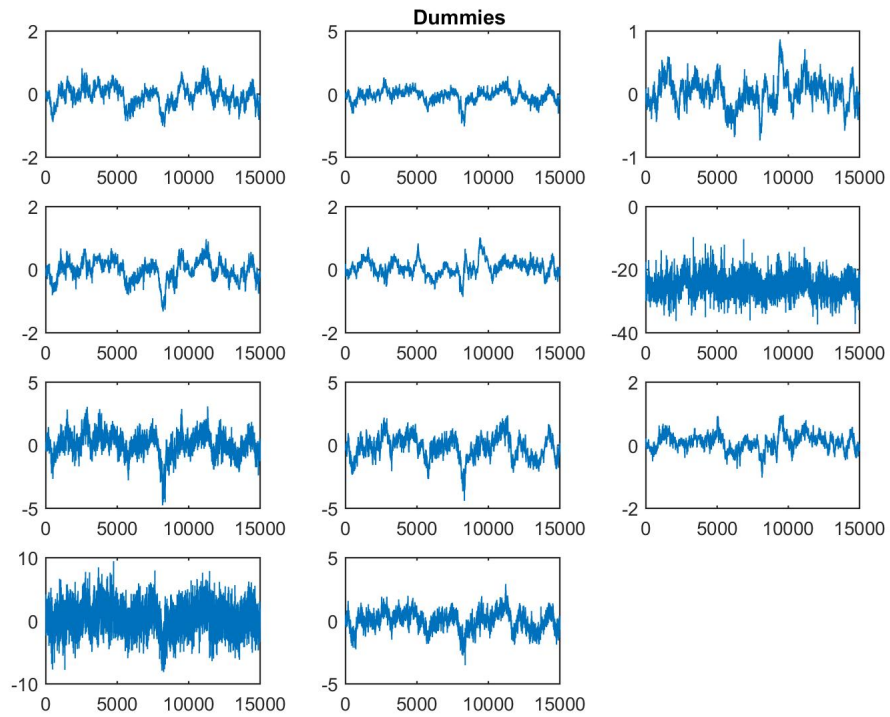
This figure shows posterior Gibbs sweeps for the latent variable $\varsigma_{i,3}^2$ for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain. Note that this parameter for Germany is undefined.

Figure B7: Gibbs sweeps, factor model, five-year horizon



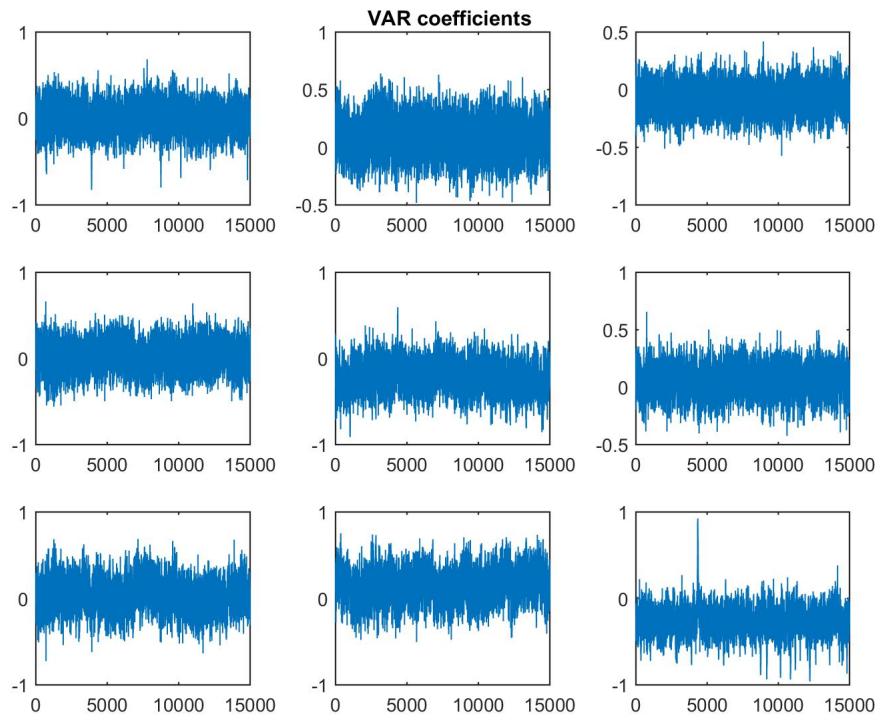
This figure shows posterior Gibbs sweeps for the parameter ρ_i for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B8: Gibbs sweeps, factor model, five-year horizon



This figure shows posterior Gibbs sweeps for the parameter c_i for the eleven countries, for 15,000 draws, at a five year horizon. The countries are read from left to right, then top to bottom. The countries are, in order: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain.

Figure B9: Gibbs sweeps, factor model, five-year horizon



This figure shows posterior Gibbs sweeps for the parameters contained in the 3 by 3 matrix Φ , for 15,000 draws, at a five year horizon. The elements in the first row correspond with the first row of Φ , the second row with the second row, etc.