Systemic risk and the Euro crisis: a narrative approach

(Preliminary and Incomplete)

Saleem A. Bahaj*

August 15, 2013

PhD Candidate, Faculty of Economics, University of Cambridge, Cambridge, United Kingdom.

Abstract

This paper uses an innovative identification approach to investigate the economic ramifications of changes government borrowing costs during the Euro-crisis. I construct a time series of critical news events in crisis hit countries and measure the impact of events on non-domestic yields at an intra-day frequency. To distinguish exogenous movements in yields from changes in macroeconomic conditions, it is postulated that domestic economic shocks have no casual effect on the high frequency bond market reaction to events abroad. The effect of external news on domestic yields can be viewed reflecting a systemic component to government borrowing costs, driven by cross-country interlinkages and the nature of the crisis. An aggregation of foreign events serves as a proxy variable for innovations in this systemic component which is used to identify a structural shock in a proxy SVAR. A counterfactual analysis is used to remove this component from bond yields of crisis hit countries: this provides evidence that yields had diverged from levels justified by local macroeconomic shocks during 2011 and early 2012 . Impulse response analysis confirm that systemic shocks were an important driver of asset prices, but the impact is relatively short-lived. The shocks were also an important driver of unemployment during the crisis.

Key words: High frequency identification, Narrative identification, Proxy SVARs,Panel VARs, Contagion, Sovereign Risk

Author's E-mail: sab202@cam.ac.uk.

JEL Classification Numbers: E44, E65, F42

*I would like to thank James Cloyne, Giancarlo Corsetti, Alex Kohlhas, Marek Jarocinski, Michele Piffer, Morten Ravn, Riccardo Trezzi, Stephen Wright for helpful comments and advice. I also thank seminar participants at the Cambridge macroeconomic workshop, the 7th Rimini Bayesian Econometrics Workshop and Goldman Sachs for their attention and feedback. Further gratitude, as always, to Donald Robertson for his continued guidance. I would acknowledge the Cambridge Endowment for Research in Finance for providing the funds to purchase the intraday data as well as Giancarlo Corsetti for his support in the funding application. The brokerage firm ICAP is attributed with providing said data. I am grateful to the Economic and Social Research Council and the Royal Economic Society for personal financial support. All remaining errors are my own.

1 Introduction

A hallmark and a gauge of intensity of the European sovereign debt crisis that started in the Autumn of 2009 has been the sharp and highly correlated movements in sovereign bond yields throughout the periphery of the monetary union. At least superficially, the gyrations in bond markets seem to have little direct correspondence to changes in macroeconomic or fiscal conditions in peripheral economies. And this observation has led to claim that yields reflect confidence factor and the crisis has been driven by self-validating expectations on behalf of market. A popular argument as put forward by De Grauwe (2011), amongst others, is that countries in a currency union are vulnerable to belief-driven crises as the sacrifice of monetary sovereignty also entails loosing access to a local lender of last resort willing to avert self-fulfilling runs in the sovereign bond market.[1] Beliefs also play role when the sustainability of the currency union as a whole comes under question; alterations in market confidence in the political commitment to the union may result in a convertibility premium during times of stress. While the political dynamics of the crisis have been a story about the design of ad-hoc insurance mechanisms between countries in the union accompanying the institutions to prevent moral hazard; in theory such insurance can be optimal, see Tirole (2013), but uncertainty over the terms under which it is offered and the eventual revenue base supporting the different varieties debt must also be a factor driving yields.

These channels are not mutually exclusive nor do I wish to pretend the above is exhaustive. But they are illustrative of the mechanisms by which Euro-zone sovereign bond yields can move separately from innovations to the local economy. One can think of the bond yield as having a systemic component, reflecting the union's fragilities, separate from the local economic conditions.

Accepting this, this paper is concerned with two empirical questions: first, how large was the systemic component in yields during crisis; second, what was the implication of changes in yields driven by this component on the local macroeconomies of crisis hit countries.

These two questions cannot be approached separately. Crises impact macroeconomic conditions; indeed, internalising this feedback is often important to generating self-fulfilling crisis in a theoretical setting (see Cohen and Portes (2004), Cohen and Villemont (2012)). To divorce these questions results in ignoring this endogeniety problem. To assess what proportion of a change in sovereign yields that can be attributed to changes in local macroeconomic conditions one needs to isolate changes in those conditions that are independent of the crisis.[2]

Dealing with this issue represents the primary contribution of this paper. There is no palatable identification strategy relying solely macroeconomic time series alone. One struggles to think of short-run, long-run or sign restrictions that are justifiable. Instead the approach taken here is to obtain identification via external information

[1]Corsetti and Dedola (2013) add nuance to this argument by pointing out that there is always a trade-off between inflation and default which may make monetary interventions ineffective. Furthermore, for monetary interventions to be credible, the fiscal authorities must be willing to make good losses by the monetary authority.

[2]The existing empirical literature on determinants of borrowing costs of the crisis (see, for example, Aizemann et al (2013), De Grauwe and Ji (2013), Percoli et al (2013), Manasse and Zevaloni (2013)) treat local economic conditions as exogenous. While Neri and Ropele (2013) consider the macroeconomic implications of movement in yields but treat the sovereign risk premia contemporaneously independent of local fundamentals

by constructing a narrative.

Narrative methods have been relied upon for empirical studies of the Euro crisis elsewhere in the literature (see Beetsma et al (2013) and Brutti and Saure (2013)). These studies have focused on the transmission of news between countries/markets on a daily basis and show that during the Euro-crisis financial markets reacted strongly to events and these reactions are transmitted across countries. I build upon this research by taking their results as given and using them to justify and motivate a narrative identification strategy for use in a structural VAR.

The transmission of foreign events serves as a source of exogenous variation. One could argue, perhaps naively, that from the point of view of a single euro-zone country, foreign policymakers are less likely to internalise local contemporaneous shocks when making their decisions and therefore foreign events could be thought of as exogenous. However, given the nature of the crisis this is an overly strong assumption. The identification strategy used here is strengthened by its reliance upon high frequency financial market reactions to foreign events to construct the narrative series. Under these circumstances, the exogeneity assumption can be justified because even if foreign agents are reacting to local economic innovations, assuming rational expectations one would expect market participants to be able to anticipate this reaction and therefore the the market move can be thought of as the "surprise" component of an announcement as is standard with a high frequency identification strategy (see Gurkaynak and Wright (2013)).

It is important to emphasise that this paper does not attempt to single out sources of the transmission between countries; some examples of other channels considered in the literature range from sunspots to changes in trade linkages, common financial holdings and changes in the beliefs of investors about policymakers decision rules (known as wake-up calls). Quantifying their relative importance is beyond the scope of the work. Instead, the issue is approached from a macroeconomic perspective. Given a set of identified shocks to the yield, it is possible to trace through how these shocks have fed into the private cost of finance and macroeconomic aggregates and construct an estimate of the overall systemic premium via a counterfactual analysis.

The empirical methodology consists of two main components, the construction of a narrative series to isolate exogenous movements in yields due to foreign crisis events and a reduced form econometric model, based around a panel vector-autoregression (VAR), to assess the impact of the systemic shocks on the macroeconomy.

Regarding the narrative series, the first step is to identify a time series of key events in each country affected by the Euro-crisis. I do this by considering the relevant news stories that make daily news summaries of European financial media outlets and can be considered an event (e.g. a policy announcement, a vote, an election result, speech etc.). The time that an event occurred is isolated and the impact on other Euro-area countries sovereign yields is calculated by looking at the response of the relevant sovereign bond yield in the immediate vicinity of the announcement. The narrative proxy variable for any given Euro-zone country is then taken to a monthly frequency by summing all the movements in yields from non-domestic events that occur over the month.

Regarding the reduced form model, the crisis provides only a short time series of observations for the purposes of estimation. In order to improve the precision of the estimates I use the information in the cross-section. However, the crisis has struck countries in a different ways with varying intensity - a homogenous parameter setup is an overly strong assumption. This problem is dealt with by using a partially pooled panel VAR model estimated using Bayesian methods (see Canova and Ciccarelli (2009, 2013) and, in particular, Jarocinski (2010)). This methodology allows for an estimate heterogeneous country specific models that make use of the information in the cross-section, as well as an average pooled model which the heterogeneous parameters are centered around. How close the parameters are to this cross-country average, i.e. the degree of pooling, is allowed to be data dependent.

Rather than including the narrative series directly in the VAR (as in for example Romer and Romer (1989, 2010)), the identification is carried out by using the narrative series as a proxy variable for the structural shock of interest in a similar vein to Mertens and Ravn (2013a,b). This approach accounts for the fact that the proxy is imperfectlt measured and suffers from scaling and censoring effects which are potentially problematic in the context of the narrative series constructed here.

The extension of the proxy-variable-based narrative identification regime into Bayesian model with heterogeneous parameters is a methodological contribution of this paper. In the frequentist setup (e.g. Stock and Watson (2012)) the proxy essentially serves as an instrument for the true structural shock. Here the proxy is treated and used differently; the distribution of the proxy is explicitly modeled and its conditional density in relation to the reduced form model is derived. I show under certain conditions the proxy has a linear relationship with reduced form residuals coupled with heavy-tailed errors. The relative size of the coefficients on the reduced form residuals is sufficient for identification up to a scaling assumption.

The main findings are the following. A counterfactual analysis on the sovereign bond yield reveals large systemic premia at certain points in the sample throughout the crisis hit countries. Premia peaked in mid 2011, the model estimates that the premium on 10 year Italian bonds hit 127bp with an equivalent figure of 381bp for Portugal, and spiked again in May 2012 around the Greek election. This implies that at the worst point of the crisis the Italian government was paying 1.3% in additional interest to borrow for 10 years compounded as a result factirs unrelated to local economic conditions.

However, from a policy perspective, the ECB interventions in the Summer/Autumn of 2012 appear to have been effective: by the end of the year, the systemic component was gone and yields appear to be at a neutral setting in line with local macroeconomic conditions.

In terms of the dynamic impact on the economy; interestingly, the changes in the systemic component tend to have a short-lived impact on the bond yields themselves: reflective of patterns of intensification followed by periods of relative calm that have marked the Euro crisis. Nonetheless, there are lingering macroeconomic implications not least in terms of unemployment which remains elevated a year after the shock. Qualitatively the responses conform to conventional economic wisdom regarding interest rate shocks: increases in yields reduce economic growth, increase unemployment and increase in private borrowing costs. Quantitatively, on average a 100bps increase in the sovereign yield corresponds to a 2ppt reduction industrial production growth - roughly speaking that is 0.6ppt off GDP -

and adds 0.9ppt to the unemployment rate (both are peak responses). From the point of view of a variance decomposition, this corresponds to around 40% of the forecast error in unemployment a year ahead. Suggesting these shocks were also an important driver of economic conditions in the sampled countries.

Curiously, there is no direct evidence that the increase in yields provokes governments to increase their primary fiscal balance and reduce their rate of rate of borrowing. Indeed, in some of the countries, the fiscal balance deteriorates on impact in response to a systemic shock. This may reflect weakening economic conditions and it is worth noting that the fiscal balance is at zero at the same point that the unemployment rate peaks. Therefore, if one defines austerity as a change in the cyclically adjusted balance then the relative co-movement of the fiscal balance and the unemployment rate is evidence of an austerity policy.

The rest of the paper is organised as follows. Following a brief discussion of the related literature. Section 2 provides an example of a narrative event in order to give context to the identification strategy. Section 3 lays out the empirical methodology including the partially pooled panel VAR and the narrative identification strategy. Section 4 discusses the construction of the narrative dataset, the assumptions justifying the identification strategy and the presents the graphs of the narrative series and how they compare with actual events. Section 5 presents the results from the VAR, including impulse responses, variance decompositions and counterfactual analysis. Section 6 presents a sensitivity analysis and section 7 concludes.

1.1 Related Literature

Notable examples of using narrative methods to identify macroeconomic shocks include Romer and Romer (1989) for monetary policy shocks, Ramey and Shapiro (1998) and Ramey (2011) for government spending shocks and Burnside, Eichenbaum and Fisher (2004), Romer and Romer (2010), Cloyne (2011) and Favero and Giavazzi (2012) for tax shocks. This paper is distinct from this literature in two ways. First, in that it focuses on changes in asset prices around market relevant events; Brutti and Saure (2012) take a similar approach by using the basic Romer and Romer (1989) specification with dummy variables on daily data. Second is that the narrative series is used as a proxy variable for the true shock rather than assuming it is observed directly, this is a similar identification setup as in Mertens and Ravn (2013) and Stock and Watson (2012). Discussion of how the reduced form panel model fits into the wider class of panel VARs can be found in Canova and Ciccarelli (2013) which offers a review of the relevant literature.

There is a rich literature of looking at high frequency market reactions to evaluate the news content in events. Fleming and Remolona (1999) and Andersen, Bollerslev, Diebold and Vega (2003, 2007) looked at the reaction to various market instruments around US macroeconomic data releases. Kaminsky and Schmukler (1999) carried out a study of how news events (both foreign and domestic) moved equity prices during the Asian crisis. Andersson (2007) studies how monetary policy meetings influence intra-day market volatility. To my knowledge, this is the first paper to look at intra-day market reactions with regards to the Euro-crisis and to use those reactions in the narrative identification framework described. However, there is a related literature assessing the determinants and transmission of risk during the crisis. Alongside the narrative studies already referenced, other contributions include Ang and Longstaff (2012), Constancio (2012), De Sanctis (2012) and Kallestrup, Lando and Murgoci (2012).

A final relevant branch is the theoretical literature related to transmission of risk in financial crisis. Earlier research related to the ERM crisis emphasised the importance various channels: from real trade based inter-linkages (Gerlach and Smets (1995)), the importance of multilateral cooperation (Buiter et al (1998) and concerns about multiplicity of equilibrium driven by changes in investors beliefs (Eichengreen and Wyplosz (1993), Jeanne and Mason (2000)). The idea of information content of foreign action changing investors beliefs is related to the idea of wake up of calls: Goldstein (1998) coined this term to capture the sudden awareness of risks in Asian financial systems during the 1997-98 crisis. This represents a major source of interdependence between Euro-area countries as well, where markets update their priors about the sustainability of the single currency in all countries based upon outcomes in one. Another important channel that is not dealt with in this paper is the interconnectedness of the banking system which has been discussed by, amongst others, Bolton and Jeanne (2011).

2 Anecdotal Evidence

This section provides an example event to illustrate the idea behind the identification strategy.

2.1 Catalonia Requests a Bailout, 28th of August 2012

At 13:01 on the 28th of August 2012 the *Reuters* news agency reported that the Spanish region of Catalonia, the wealthiest in the country with an economy the same size as Portugal, would request 5 billion euros of aid from the Spain's Regional Liquidity Fund. The report was confirmed officially by a Catalan government spokesman 5 minutes later, a long with the warning that the region would "*not accept political conditions for the aid*". The announcement sent jitters through financial markets. At 13:00 Spanish ten year bonds yielded 6.36% by 14:00 the yield was 6.43%, a 3 standard deviation move at an hourly frequency over the crisis period.

Spanish regions had been hit hard by the crisis. They relied upon the frothy construction and property sectors for their tax revenues and the bursting of Spain's housing bubble left them with large deficits. Approximately a third of Spain's overall deficit in 2011 was down to the regional governments. Spain estabilished it's Regional Liquidity Mechanism the previous month with 18 billion Euros of capital to support regions facing borrowing difficulties. It had already been tapped by Valencia and Murcia. Catalonia's decision would leave half the fund committed with other regions still in trouble. And the need of the country's richest region (and the most indebted) to request support for the central government was unlikely to be considered a positive signal that regional governments had their fiscal affairs in order. The decision also ignited regional tensions within Spain. Catalan elections in November led to a seperatist majority in the regional parliament. The Catalan Premier stated that the bailout money was

Figure 1: Catalonia Requests a Bailout, 28th August 2012: High Frequency Market Reaction

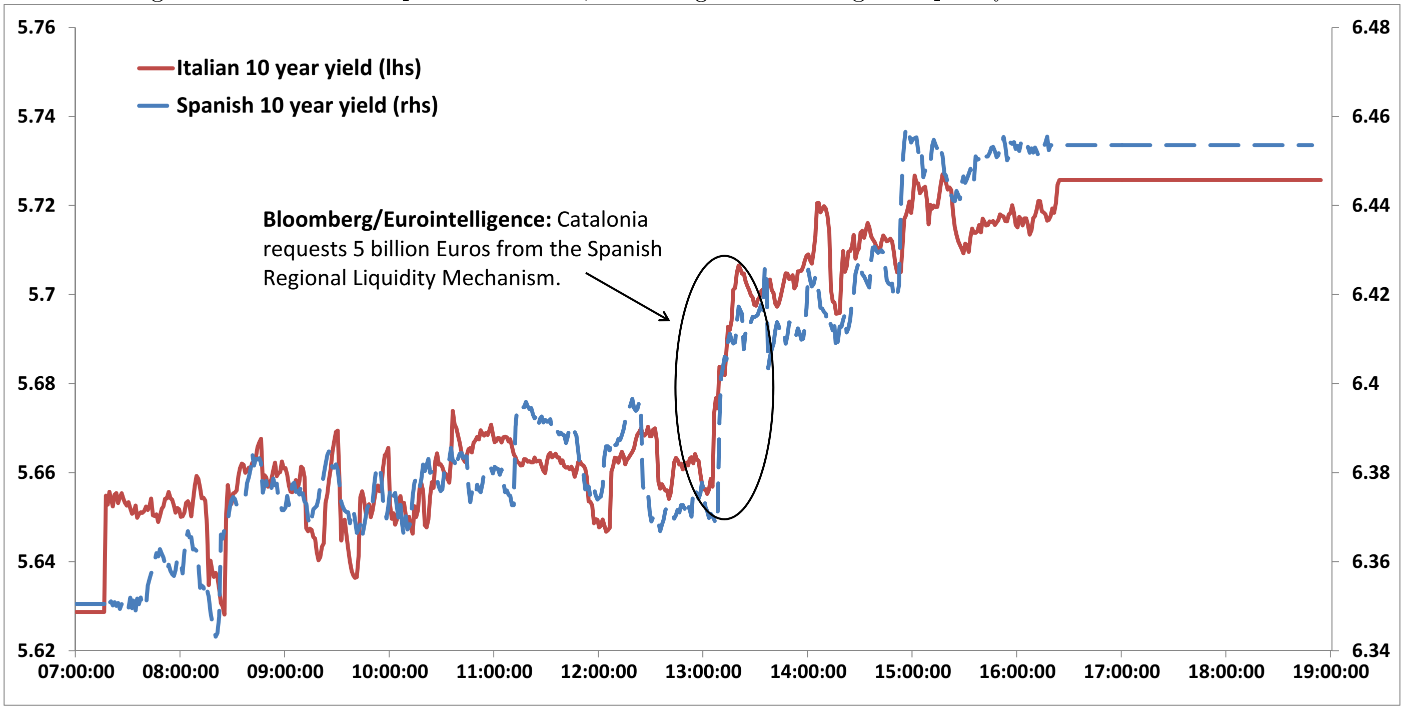

not a transfer from Madrid but simply a return of all the tax money that had flowed out of the region in previous years.

Still, the bailout decision was largely domestic policy matter. In effect, it represented a transfer of liabilities from the regions to the central government in Spain and a mutualisation of the Spanish public balance sheet. It had no direct international aspect. Nonetheless, the move in yields was not isolated to Spain. Italian 10 year bond yields increased by 5 basis points in the immediate vicinity of the announcement and yields in "core" countries declined. For example, German yields fell by 1.5bp in the few minutes following the announcement. Curiously, there was little response to the announcement in Irish and Portuguese bond markets. But by this stage in the crisis these two countries had been bailed out by the Troika and market attention was elsewhere. To give a sense of the market move around the announcement figure 1 presents the intraday bond yields in Spain and Italy on the 28th of August 2012.

2.2 Discussion

The point that this example is designed to illustrate is that there little is chance that this move in the Italian yield can be thought of as being a function of a change in macroeconomic conditions in Italy in August 2012. It was not a reaction to a weakening in Italian total factor productivity, for example. Indeed, it is unlikely that the Catalans were thinking too much about Italy when they made their announcement. But even if there was endogenous feedback between policy choices in one country from macroeconomic conditions in other, and such reactions may be reasonable in light of of how the crisis developed, any systematic reaction should be anticipated by market participants and therefore should be priced. Thus we can think of the move in the Italian yield observed above as not being in some way caused by new information about Italian economic conditions; it may cause changes in those conditions but was not caused by them.

One point to note is that this identification strategy breaks down if foreign agents are reacting to changes in local macroeconomic conditions that the market is not yet aware of, meaning that that the foreign event is informative about local economic shocks. Therefore, the critical identifying assumption is not that foreign events are not a reaction to local macroeconomic shocks but that markets are rational and there is no information asymmetry between market participants and foreign agents regarding these shocks at the point of an event. The information asymmetry problem is why a focus on foreign events is necessary; for example, the announcement of a new austerity package may lower the local bond yield but that move is correlated with the fiscal shock that hits the country at the same time.

Narrative strategy used in this paper attempts to isolate events like the example above, look at the reaction of foreign bond markets and use an aggregation of those reactions as a proxy for a shock to the yield. I refer to this as a systemic shock as it captures the fact that the bond yield is being driven by external events that transmit due to the systemic nature of the crisis. This is reflects that fact that this strategy cannot be thought of capturing a channel that is purely related to local sovereign risk; although that is perhaps the primary feedback mechansim to real economy. There may be other channels by which are captured by the proxy and feed through to the economy such as external demand or increased uncertainty. This is important to bear in mind when interpreting the results.

3 Econometric Methodology

This section lays out the econometric methodology. The reduced form model follows a Bayesian panel SVAR with cross-sectional heterogeneity in slope and covariance matrices estimated with partial pooling along the lines of Jarocinski (2010). To identify systemic shocks, the strategy relies upon a proxy SVAR approach of Mertens and Ravn (2013a,b) and Stock and Watson (2012). These contributions rely on a frequentist setup. A pure Bayesian methodology adds a layer of complexity in the sense that one cannot directly apply the two-stage procedure used in the frequentist approaches. However, it is advantageous as the information from the proxy is included in the VAR coefficients and confidence intervals can readily be constructed that include uncertainty from both the reduced form VAR estimation and the identification procedure.

As an additional methodological challenge, the VAR frequency is monthly yet, in the spirit of Mertens and Ravn, the proxy variable is best viewed as an aggregation of censored events measured with error. This is dealt with by using a leptokurtic distribution to capture the effect of the censoring process when estimating the relationship between the narrative instrument and the structural systemic shock. The remainder of this section is divided as follows: the first describes the panel VAR model used; specifically the prior structure used in the hierarchical model. The second details the narrative identification strategy and the econometric issues surrounding the use of an aggregated censored instrument. The third discusses the effect of this prior structure on the interpretation of the results and the fourth deals with the algorithm used in estimation.

3.1 The reduced form VAR

The primary feature of the panel VAR model used here is to allow for heterogeneity in the slope and covariance matrices of the country specific models. This is done by setting up the country specific parameters in the shape of a hierarchy with exchangeable priors, the country specific models are then estimated using partial pooling. The effect of the partially pooled estimator is to help improve the country specific model by including information from the cross section as well as allowing for an estimate for an average model. The exchangeable prior results in shrinkage of the country specific coefficients towards this common average with the degree of shrinkage is allowed to be data dependent. At the first level of the hierarchy, a prior is formed over the distribution of parameters of the individual country models. At the second level, hyperpriors are formed over the hyperparameters for the common components in the distribution of country-specific parameters. This section sketches the model structure and offers a brief justification for the selected priors; Jarocinski (2010) offers a fuller discussion in this regard.

In order to describe the VAR structure formally the following notation is adhered to: vectors are lower case symbols, matrices are uppercase symbols, the indices $c = 1, ..., C$, $l = 1, ..., L$ and $t = 1, ..., T$ denote countries, VAR lags and time periods (months, specifically) respectively. The dimension of the VAR is denoted N. For each country the reduced form VAR is of the form:

$$y_{c,t} = \sum_{l=1}^{L} B'_{c,l} y_{c,t-l} + \Gamma'_c z_t + u_{c,t} \tag{1}$$

Where y_{ct} is a $N \times 1$ vector of endogenous country variables, B'_{cl} is the matrix of country specific coefficients on lag l of the endogenous variables, z_t are deterministic variables with corresponding coefficient Γ_c and $u_{c,t}$ is the vector of VAR innovations at time t. These innovations are assumed to be i.i.d. and to have a prior distribution $u_{c,t} \sim N(0, \Sigma_{c,u})$, where $\Sigma_{c,u}$ is a covariance matrix to be estimated. As is standard in the Bayesian VAR literature, equation 1 can be rewritten in its SURE representation. Let $x_{c,t} = [y'_{c,t-1}, ..., y'_{c,t-L}]'$, stacking the t observations on $y_{c,t}$, $x_{c,t}$ and z_t vertically to create data matrices allows the model to be expressed as:

$$Y_c = X_c B_c + Z_c \Gamma_c + U_c$$

Where $B_c = [B'_{c,1}, ..., B'_{c,l}]$. Last, define the vectorised data and parameter terms as: $y_c = vec(Y_c)$, $\beta_c = vec(B_c)$ and $\gamma_c = vec(\Gamma_c)$.

3.1.1 The first level of the hierarchy

The first level of the hierarchy governs the statistical form of the individual country models. Given the prior over the VAR innovations, the likelihood for the model corresponding to country c is given by:

$$p(y_c|\beta_c, \gamma_c, \Sigma_c) = N((I_N \otimes X_c)\beta_c + (I_N \otimes Z_c)\gamma_c, (\Sigma_c \otimes I_{T_c})) \tag{2}$$

The country slope coefficients β_c are assumed to have prior normal distribution with common mean $\bar{\beta}$ and variance Λ_c which is country specific:

$$p(\beta_c|\bar{\beta}, \Lambda_{1c}) = N(\bar{\beta}, \Lambda_{1c}) \tag{3}$$

The parameter vector $\bar{\beta}$ serves as the cross-country average slope coefficients, as with a standard panel model. A non-informative prior is assumed for γ_c in each country:

$$p(\gamma_c) \propto 1 \tag{4}$$

The covariance matrix of the residuals is also drawn from a common distribution (in this case inverse-Wishart) with a common scale parameter $\bar{S}$:

$$p(\Sigma_{c,u}|\bar{S}, \kappa) = iW(\bar{S}, \kappa) \tag{5}$$

The purpose of this prior is to formalise the existence of a cross-country average covariance matrix, alongside $\bar{\beta}$, for use in calculating the impulse responses of the cross-country average model. This prior implies that the posterior of $\bar{S}$ can be used to estimate a cross-country covariance matrix centered around the harmonic mean of the individual country estimates. The degrees of freedom parameter, κ, which is defined on the positive real line, determines the degree of shrinkage of the estimated country specific covariance matrices towards said common mean as described below.

3.1.2 The second level of the hierarchy

The role of the second level of the hierarchy is to determine the common cross-country elements, specifically the the prior distributions of the hyper-parameters in the country models. As it is desirable to let the data determine the common means a diffuse prior is used for both $\bar{\beta}$ and $\bar{S}$:

$$p(\bar{\beta}) \propto 1 \tag{6}$$

$$p(\bar{S}) \propto |\bar{S}|^{-0.5(N+1)} \tag{7}$$

The degree of shrinkage applied to slope coefficients β_c is governed by the country specific covariance matrices Λ_c. It is assumed that this covariance matrix decomposed into a country specific positive definite matrix (L_c) and a common scale parameter contained in the set of positive real numbers (λ_1):

$$\Lambda_{1c} = \lambda_1 L_{1c}$$

The matrix L_c is considered to be deterministic and is constructed from the ratios of the variances of the residuals from univariate autoregressive estimates of endogenous country variables as described in Jarocinski (2010). The form of specification for L_c follows similar intuition to that behind the variance of the Minnesota prior. The idea being that the relative variance of a coefficient is determined by the relative size of the unexpected movements of the variables in question. See Litterman (1986) for a fuller discussion.

The parameter L_c only helps determine the relative variances of the coefficient estimates. What matters for the tightness of the parameter estimates about the common mean is λ_1. This hyper-parameter acts as a scale parameter for the overall variance of the slope parameters across countries and determines the degree of shrinkage. To understand the impact of λ_1 it is useful to consider the two extreme cases. An estimate of $\lambda_1 = 0$ is equivalent to saying there is no variance about $\bar{\beta}$ - that the slope coefficients are identical across countries. This results in posterior means of β_c equivalent to a pooled panel VAR. Conversely, as $\lambda_1 \to \infty$ the distribution about $\bar{\beta}$ is sufficiently diffuse such that there is no information contained within the common mean. As a result the posterior means of β_c are equivalent to those as if each country has been estimated separately[3]. Hence, any increase in λ_1 is equivalent to a reduction in shrinkage and the estimated country models are allowed to become increasingly different. While it is possible to give an interpretation to what changes in λ_1 imply for the model, the absolute level, as with the variance of VAR coefficients more generally, is harder to interpret. As a result an informative prior for λ_1 is difficult to justify. However, it is desirable to let the data itself speak for how much shrinkage is needed and therefore a non-informative prior is not problematic conceptually. The inverse-Gamma distribution has the appropriate support and delivers conditional conjugacy; this implies a prior of:

$$p(\lambda_1|s,v) = IG_2 \propto \lambda_1^{\frac{-v+2}{2}} exp\{-\frac{1}{2}\frac{s}{\lambda_1}\} \tag{8}$$

with hyper-parameters s and v. The hyper-parameters are specified as in Gelman (2006): $v = -1$ and $s = 0$, i.e. $p(\lambda_1) \propto \lambda_1^{-1/2}$. Which is equivalent to the standard deviations for the individual coefficients having uniformly distributed prior over the positive portion of the real line[4].

[3] Note that due to the prior assumption on common covariance matrices in equation 5 the degree of shrinkage on the estimated impulse responses also depends on κ.

[4] An alternative is to set $s = \varepsilon$, $v = \varepsilon$ with ε small - i.e. approximate $p(\lambda) \propto 1$. This means the variance rather than the standard

The parameter κ plays a similar but inverted role to λ_1 for the covariance matrices. As $\kappa \to \infty$ the distribution in 5 becomes degenerate with all the mass concentrated upon a point corresponding to the common mean covariance matrix (determined by $\bar{S}$); hence, the posterior means of $\Sigma_{c,u}$ would be identical. And as κ decreases, the country covariances are allowed to become increasingly different, to the extent that $\kappa = 0$ implies there is no shrinkage in terms of covariances.[5]

Due to the multivariate nature of the model, there is no classical distribution that can serve as a prior on κ that is conditionally conjugate. For now, κ is treated as deterministic. This is a less problematic assumption than with λ_1 as the interpretation of κ is clearer: it serves to determine the weight attached to $\bar{S}$ in posterior distribution of $\Sigma_{c,u}$. Prior hyper-parameters of this form are common in literature, the value of $N+2$ is used here (as suggested in Giannone et al (2012)) as it guarantees the existence of a prior mean for $\Sigma_{c,u}$ while imposing the minimum shrinkage. It is worth noting that since the time series dimension in the model is reasonably large, large values of κ are also needed before the model places more weight on $\bar{S}$ than the country specific sum of squares in determining the posterior mean of $\Sigma_{c,u}$. Therefore, the results are not too sensitive to the choice of this hyper-parameter.

3.2 Identification

Assume the existence of a set of unidentified structural shocks, denote them ϵ_t (for notational convenience the c subscript is dropped in this section). Following the standard SVAR assumptions, the reduced form shocks are a linear combination of structural shocks $\mathcal{A}u_t = \epsilon_t$, where $\mathcal{A}$ is an $N \times N$ identification matrix scaled such that the structural shocks have unit variance. The vector of structural shocks can be partitioned into a shock of interest, i.e. the systemic shock, and other structural shocks $\epsilon_t = (\varepsilon_t', \tilde{\varepsilon}_t')'$. Thus the identification matrix can be partitioned such that $\mathcal{A} = [a_1^{'}, a_2^{'}]^{'}$ with $\varepsilon_t = a_1 u_t$. The identification matrix is non-singular with $\mathcal{A}^{-1} = \mathcal{B}$ and $u_t = \mathcal{B}\epsilon_t$. The matrix $\mathcal{B}$ can be partitioned such that that $\mathcal{B} = [b_1, b_2]$ and $u_t = b_1\varepsilon_t + b_2\tilde{\varepsilon}_t$. Note, a_1 and b_1 are respectively $1 \times N$ and $N \times 1$ vectors.

In contrast to Mertens and Ravn (2013b), the approach taken here works with the $\mathcal{A}$ rather than the $\mathcal{B}$ matrix (i.e. an A-type model in the SVAR parlance of Amisano and Giannini (1997)). This is of technical importance as it means one can specify the relationship between the proxy and the shocks in a single equation. The vector a_1 is sufficient to estimate a time series of structural shocks ε_t and thus can be used for the purposes of counter-factual analysis. In order to construct impulse responses and variance decompositions, an estimate of b_1 is needed. With an estimate of the covariance matrix one can easily switch between the vectors a_1 and b_1. Using the relationship $\Sigma_u \mathcal{A}' = \mathcal{B}$, it follows that $b_1 = \Sigma_u a_1^{'}$.

deviation approaches the uniform prior. However, Gelman (2006) shows that this can have an unforeseen impact on the posterior as the prior density has a fat right tail which places less weight on cases where the models are very similar (and λ is small).

[5]In practice, it is not possible to apply the no-shrinkage case and estimate a common mean covariance matrix that conforms with equation 5; for the posterior distribution of $\bar{S}$ to be proper it is necessary to have $\kappa > (N-1)/C > 0$. However, setting $\kappa \approx (N-1)/C$, the minimum permissible value, leaves the common covariance matrix with an almost negligible role in determining the country specific estimates; in the model used here less than a 2% weight would be placed on the common covariance matrix for this value of κ in the posterior mean of the covariance for each country. Thus, this restriction upon the hyperparameter does not play a meaningful role.

3.2.1 Narrative identification using proxy variables.

I assume there exists a proxy variable m_t that has the properties $E(m_t\varepsilon_t) = \phi$, $E(m_t\tilde{\varepsilon}_t) = 0$ where ϕ is a scalar. This zero covariance is the critical identifying assumption. Since the reduced form of shocks of the model are observable one can estimate $E(m_t u_t)$; which can be used to construct an estimate (up to a signing convention) the coefficients of a_1. Following Mertens and Ravn (2013a,b), the proxy is could be considered a scaled version of the true shock measured with some error, for example $m_t = \phi\varepsilon_t + v_t$, where v_t is measurement error uncorrelated with the shock ($NID \sim (0, \sigma_v^2)$) and ϕ is a scalar coefficient. One can estimate:

$$m_t = \phi a_1 u_t + v_t$$

And the unit variance restriction on the structural shock implies the quadratic form $a_1 \Sigma_u a_1^{'} = 1$, this provides the additional restriction to identify ϕ. [6]

3.2.2 Aggregating high frequency events into the proxy.

The simple linear relationship above between the proxy and the reduced form model may be construed as an overly strong assumption in the context of the strategy used here. The construction of the proxy is discussed in great detail below but to summarise: the proxy is built by using news summaries to isolate country specific events. These are events are then timed using a news wire. Then the high frequency bond market reaction is calculated in a window containing that event. The proxy for a particular country is the sum of the bond market reactions around the events that happened abroad in a particular month.

This presents several issues from an econometric perspective. The proxy variable used is an aggregation of high frequency bond market reactions which are themselves stochastic. Events are not continuously observed and the event inclusion criteria means certain types of news are omitted. Thus the proxy is effectively the aggregation of censored observations.

The market reaction may also be an imperfect gauge of the informational content of an event. Market specific factors such as liquidity or large transactions can result in a noisy signal. The informational content of an event can be difficult to process quickly; there is the possibility that markets treat events differently once information is digested and change their initial reaction. Furthermore, the market reaction may be slowed by a lags in the decision making of institutional investors and the time taken for order books to be processed. The immediate response to a shock may propagate as the agents in financial markets adjust their positions accordingly over a longer horizon. This means that there will be measurement error contained in the observed reaction to each event. But it suggests there may be also scaling effects: the true informational content of the event may be different from the initial

[6] While assuming unit variances is standard practice in this setup it is not completely innocuous. There are two costs to this assumption. First, it is impossible to make any statement about the relative variances of shocks across countries. Second, a one standard deviation stock has no interpretation so one needs to make a scaling assumption when computing impulse responses.

reaction in a regular fashion.

Despite the empirical issues with this high frequency set up it is possible with a few assumptions to return to the simple linear model above and use it for the purposes of identification without any explicit estimation of the high frequency statistical process driving the bond yield or event occurrence. For the purposes of exposition consider events at a daily frequency, this can be extended done to a higher frequency by simply narrowing the time window. Assume that there exists is a daily series of scalar structural shocks that sum perfectly to create a monthly shock of interest:

$$\varepsilon_t = \sum_{d=1}^{M} \varepsilon_{dt} \tag{9}$$

Where, d denotes a day in month t and M is the total days in the month. The daily structural shocks have the property: $\varepsilon_{dt} \sim NID(0, 1/M)$ such that the monthly shock is Gaussian with unit variance. The series ε_{dt} is unobservable as is ε_t; however, a daily proxy series for the shocks exists and can be summed to give a monthly proxy in a similar fashion:

$$m_t = \sum_{d=1}^{M} m_{dt} \tag{10}$$

The proxy has the properties $E(m_t\varepsilon_t) = \phi$, $E(m_t\tilde{\varepsilon}_t) = 0$ where ϕ is a scalar; furthermore, it is assumed that $E(m_{dt}m_{st}) = 0$ and by extension $E(m_t m_s) = 0\ \forall t, s$. Evidence to back both these assumptions is offered in section 4.

As discussed, in reality the proxy is censored such that on some days it is not observed. A depiction of the process is:

$$m_{dt} = D_{dt}(\psi\varepsilon_{dt} + v_{dt}) \tag{11}$$

where v_{dt} is the measurement error associated with the proxy, ψ is a parameter to pick up scaling and D_{dt} is added as an indicator variable for censoring. It may be desirable to extend the model away from purely random censoring and allow for $E(\varepsilon_{dt}D_{dt}) \neq 0$ but this presents many technical challenges. For simplicity, it is assumed that D_t is an independent variable takes a value of 1 with probability p and zero otherwise (i.e. $1-p$ is the probability that an observation is censored). The density of m_{dt} can be expressed as:

$$P(m_d|\psi, \varepsilon_d, p) = \left((2\pi\sigma_v^2)^{-\frac{1}{2}} exp\left\{ -\frac{1}{2}\left(\frac{m_d - \psi\varepsilon_d}{\sigma_v}\right)^2 \right\} p \right)^{1-I(m_d=0)} (1-p)^{I(m_d=0)} \tag{12}$$

Where $I(m_d = 0)$ is an indicator variable which returns 1 if $m_d = 0$ and zero otherwise. As one would expect the maximum likelihood estimate for p is simply the proportion of censored daily observations.

Under random censoring it is possible to show that the conditional relationship between m_t and u_t is:

$$m_t = \Upsilon' u_t + \omega_t \tag{13}$$

where $\omega_t | u_t \sim iid(0, \sigma_\omega^2)$, $\Upsilon' = p\psi a_1$ and $\sigma_\omega^2 = pM\sigma_v^2$; the proof is provided in appendix A.

This returns us to the linear specification at a monthly frequency. There is a last hitch, the distribution of ω is non-Gaussian due to the censoring process. It has the same support as the normal and retains symmetry but there is a difference in the fourth moment. As censoring takes more mass from the shoulders of the normal distribution than from the tails it means that the final censored distribution is leptokurtic. A standard approximation for symmetric, leptokurtic distribution is a Student-t.

Bayesian estimation of linear models such as equation 13 extended to include t-errors has been covered extensively the literature (see Geweke (1993) for the classic treatment). Letting $\mathcal{M}$ be the matrix of proxy observations stacked over time; using this approach one can approximate the conditional distribution of the proxy as:

$$p(\mathcal{M} | U, \Upsilon, \sigma_w^2; \nu) \sim t(U\Upsilon, \sigma_\omega^2 I_T; \nu)$$

Where t denotes a multivariate scaled Student's-t distribution with mean and variance given by $U\Upsilon$ and $\sigma_\omega^2 I_T$ and a scalar degrees of freedom parameter ν. Once Υ is known, using the quadratic form $a_1 \Sigma_u a_1' = 1$ it is again possible to use the estimate of Σ_u to remove the effect of ψ and the censoring process and obtain an estimate of a_1.

3.2.3 Priors on the identification parameters.

As much of the data entering the model is at a monthly frequency, jointly estimating the daily censoring process adds an unnecessary layer of complexity. Instead, p is set deterministically and calibrated to the proportion of the days in the sample where m_{dt} is observed - which is equivalent to the maximum likelihood estimator of p. From the narrative series described below this results in a $p = 0.04$. Similarly M is deterministic and is set to 30 for the purposes of the estimation. The extent of the excess kurtosis that arises from the censoring process is a function of only of p and M; since the degrees of freedom parameter ν determines the excess kurtosis in the t distribution that is used to approximate the combined censored observations, this parameter is necessarily also deterministic. Given p and M it is possible to calculate ν via simulated method of moments; plugging in the numbers from above this results in $\nu = 13$ to the nearest integer. For these values of p and M the Student-t approximation is good and certainly improves on a normality assumption for ω. Figure 2 illustrates this by comparing the kernel density estimator of a simulated ω for arbitrary σ_v^2 and ψ with equivalent densities from a normal and Student-t with matched moments.

The parameters in equation 13 are assumed to be country specific with a cross-sectional relationships along the same lines as the parameters in the reduced for VAR. The country slope coefficients Υ_c are assumed to have prior normal distribution with common mean $\bar{\Upsilon}$ and variance Λ_{2c} which is country specific:

$$p(\Upsilon_c|\bar{\Upsilon}, \Lambda_{2c}) = N(\bar{\Upsilon}, \Lambda_{2c}) \tag{14}$$

As previously, the variance matrix can be decomposed into a country specific deterministic component and a common parameter determining the degree of cross-country shrinkage across Υ_c:

$$\Lambda_{2c} = \lambda_2 L_{2c}$$

With L_{2c} set along the same lines as L_{1c}. The parameter λ_2 has the same prior as λ_1 as described in equation 8, and plays the same role. The parameters $\bar{\Upsilon}$ and $\sigma^2_{c\omega}$ have a diffuse priors: $p(\bar{\Upsilon}) \propto 1$ and $p(\sigma^2_{c\omega}) \propto \sigma^{-1}_{c\omega}$.

3.3 Discussion of the impact of the hierarchical model

For interpreting the results and making cross country comparisons it is useful to elaborate on the impact of the assumed prior structure on the estimated parameters. As shown in appendix C the exchangeable prior on the coefficients leads to estimation that takes a form of partial pooling: the estimated parameters have a posterior mean that is a weighted average of the coefficients of a pooled model and the parameters as if every country model had been estimated separately (an unpooled model). For the model as a whole what determines how close the estimation is to each extreme is the parameter λ_1 in the case of the reduced form slope coefficients and λ_2 in the case of the identification model. However, the extent of the pooling also varies from country to country depending on well each country's model fits. This happens through two channels: first, in the posterior mean of each country's coefficient the weight attached to the unpooled estimates is increasing in precision of that country's model. So countries that are estimated less precisely are closer to the pooled mean. Second, the posterior mean of the pooled coefficients ($\bar{\beta}$,$\bar{\Upsilon}$) are a weighted average of the country specific estimates; and these weights are also increasing in the precision of each country's model. Hence, the pooled model is closer to the countries that are estimated more precisely. Given the multivariate nature of the model it is not possible to disentangle these relative weights in a single measure as they are parameter specific.

However, this partial pooling has implications for the identification strategy. It implies that in a country where the proxy variable used is not leading to a precise estimate of Υ_c, information from other countries is used to tighten the confidence bands and pin down the estimate. Thus, the proxy variable does not need to have a very strong correlation with the true structural shock in every country in the sample so long as it works for some countries and the countries in the sample are sufficiently similar.

3.4 Estimation

As one would expect the unconditional densities of the parameters cannot be determined analytically, hence they are computed numerically using Markov Chain Monte Carlo methods. The functional forms of the priors, as well as having an interpretation regarding a common average model, are motivated by computational convenience as they are, with one exception, conditionally conjugate. That is to say they lead to a set of conditional posterior distributions that are standard and of the same family as the prior. This motivates the use of a Gibbs Sampler to construct the posteriors. The departure from conditional conjugacy arises due to the presence of the proxy variable which alters conditional density of the coefficient estimates in the reduced form VAR. At first glance this may be viewed as counter intuitive: in most VAR models the reduced form coefficients are independent from the identification strategy. However, the proxy variable is informative about the residuals: if an element of $\mathcal{M}_c$ is large in absolute terms, depending on Υ_c, the individual reduced form residuals should also be so. In effect, this weights certain observations in the likelihood and leads to a non-standard conditional density. To deal with this a Metropolis-Hastings step is included within the sampler to approximate the non-standard density of the reduced form coefficients. This additional step aided is by recognising that the reduce form coefficients have a Gaussian conditional density when the proxy variable is not conditioned upon; these densities then serve as convenient candidate distributions for use in the Metropolis-Hasting algorithm.[7] The full form of the Metropolis-within-Gibbs sampling algorithm is laid out in appendix C.

4 Constructing the narrative series

This section describes the construction of the narrative series that serves as a proxy for systemic shocks. To construct the instrument three pieces of information are required: *(i)* a set of important events related to the crisis that are country specific; *(ii)* the time that each event occurred and *(iii)* the high frequency bond reaction around each event occurring. I describe how each piece is obtained in turn. The narrative analysis conducted from July 2009 to March 2013.[8]

4.1 Constructing the proxy variable

The Euro crisis is well documented and made up of a vast set of largely idiosyncratic events that make tracking the evolution of the crisis from a narrative perspective methodologically challenging. The flow of information related to the crisis has to be processed in an objective fashion to prevent systematic errors that may bias the results. A filter

[7] An alternative to this approach is to depart from a fully Bayesian set up and estimate the reduced form VAR without conditioning on the proxy variable. This leads a complete set of conditional conjugate distributions and can be estimated using a standard Gibbs Sampler. Experiments with this approach yielded little difference in the results and is computationally more straightforward; although a Bayesian purist may find it unappealing.

[8] The dataset is constructed in the order described. A mapping back from major bond market moves to the time series of events is undesirable as it means small or anticipated events are missed but it is also problematic from econometric perspective as it is almost equivalent to building a series of dummy variables for large movements.

for the information flow is, by definition, media outlets and these have been relied upon for narrative studies of the crisis elsewhere in the literature (e.g. Beetsma et al (2013) and Brutti and Saure (2013)). As the pan-Eurozone bond market reaction is the eventual variable of interest, this motivates a focus on events in terms of financial news.

The approach taken here is to rely on news summaries to isolate events.[9] The financial news sources *Bloomberg* and *EuroIntelligence* both compile a daily news briefing for the European economic news, with the former released in the afternoon and the latter in the morning. Both contain somewhere between 10-12 discrete news stories that are presented as digestible paragraph long summaries. The selection of stories by *EuroIntelligence* appears to be at the judgment of their editorial staff and includes a "headline" story which the staff consider the main the event for the day. The *Bloomberg* summary represents the most read (presumably by market participants as they are the main users) European news stories during the day.[10] As the objective of this narrative is to assess the impact of foreign, country-specific, events on local borrowing costs for use as an identification strategy, this reliance on pan-European news summaries serves as a filter as the country specific event must be of sufficient international interest to make the briefing. As discussed below, this is not to say the market reacts strongly to every news story within the summaries.

While there is a large overlap between these two sources, the timing of the news summary turns out to be of importance. Twenty-four hours is a long time at certain points in the crisis; stories that occur overnight can be overtaken by events the next day such that they do not make the afternoon briefing, similarly events that occur early in the day are out of date once a briefing is out the next morning. Combined these two briefings provide half day snapshots of the key news stories that should be affecting Eurozone bond markets.[11]

Given the set of stories that appear within the summaries, the next step is to determine whether any of them constitute an "event" that is of interest for the narrative. The news briefings are read manually and to be classified as an event and included in the narrative, a news story must satisfy the following criteria:

1. The story must relate to a single crisis hit country; specifically, one of Greece, Cyprus, Portugal, Ireland, Italy or Spain. Large pan-Eurozone policy interventions are not included as the identifying assumptions are harder to justify particularly when one considers the involvement of the ECB may imply a monetary aspect. Experiments with political events in non-crisis countries revealed that bond markets do not react strongly to this form a news and thus these countries are omitted for the sack of parsimony. Stories that relate to foreign or European policymakers intervening in specific crisis country are not omitted in the benchmark specification although they are in a robustness analysis.

[9]This is the approach taken in Beetsma et al (2013). Rather than rely on news briefings to build a narrative an alternative methodology is to use ready made crisis time lines such as those compiled by the ECB or by private media outlets (as in Brutti and Saure (2013)) . However, the former only contains events that involve the the ECB, the EU or the G20 in some capacity. The latter have richer coverage but cover inconsistent periods and have no clear criteria for event selection.

[10]The *Bloomberg* news coverage for the Eurozone is available from any *Bloomberg* terminal by entering TOP EUROPE; in the early evening a news item appears with the top stories for the day and a complete history of previous briefings is available. The details on how EuroIntelligence operates can be found in Beetsma et al (2013) who rely on the source to construct their narrative series.

[11]There are alternative sources available: for example from the *Reuters* news agency. However, the two sources used here are chosen due to the difference in release times. Experiments with alternative sources does not improve coverage as the stories have almost complete overlap with the two already considered.

2. The event must be timeable in the sense that it is possible to isolate when it occurs as to determine the market's reaction. The focus therefore is mainly (although not exclusively) on official announcements and on the record statements. It is important to emphasise that an event is considered to be something that happened at a particular time rather than any news story that is country-specific. This criterion is discussed in more detail below.

3. Certain sorts of news are not considered:

   (a) Anonymously sourced rumours/news that may make headlines.

   (b) Reports by private companies or about private companies. News about or statements by individuals are also not included unless that individual has an official policy making or political capacity.

   (c) Similarly, editorial statements are not included.

   (d) Data releases are also not included as while surprises in these indicators are strongly correlated across countries they are often reflective of real shocks such as a common Euro Area business cycle. An exception to this are official revisions to past and future projections of annual fiscal numbers which were of key importance during the early stages of the crisis in Greece. The relatively low frequency of these numbers and the lag in their release prevents the market reaction to them being related to cyclical news.

The next step is to time when these events occur and gauge the market reaction. Events are timed to the minute when the first headline related to the story appears on the *Bloomberg* newswire. This need for an initial headline is less restrictive than one may think. While many news stories are ongoing over several days or even weeks, most are a combination of discrete events that break at certain times. The bulk of events considered in the dataset are essentially announcements, speeches or statements to the press from an official source and therefore, the timing is not subjective. As a caveat, for this approach to be workable, news stories as they appear in the summary often have to be broken up into discrete announcements. For example, stories often include comments from several individual policymakers. In such circumstances the time of each statement would be used as an event - or combined into a longer window if the statements are close together.

However, there are exceptions where this chain of discrete announcements does not apply and, therefore, an untimeable story is one where it is impossible to identify an initial headline in an objective fashion. News stories where it is impossible to determine a time are often ongoing events and not breaking news, for example a strike which lasts all day has no specific time with which on can assess the market's reaction.[12] Alternatively, it could be story that mutates rapidly with either conflicting reports or more and more details emerging over a sustained period of time moving markets in a variety of ways. While it is possible to analyse such an event *ex-post* it is impossible to judge the appropriate time to asses the market reaction in real time.[13]

[12] On the other hand strikes are announced in advance and such announcements are considered a timeable event.

[13] As an example of this, consider the case of 20th October 2011 when a Greek protestor tragically died in violent demonstrations on

It is worth noting that some of the events may have little informational content or be entirely anticipated by the time of the announcement. However, the market's reaction filters out events which do not matter and this means it is optimal to consider as broad a class of events as possible, so long as they can be considered exogenous. What drives the variation in the narrative series are the small number of events that provoke large market reactions.

For events that can be timed, the market reaction is considered over a 20 minute window on either side of the initial headline. [14]

A further issue with timing is how to deal with events that occur when markets closed. European policymaker's penchant for late night meetings means that omitting these events altogether risks throwing out critical information. On the other hand, the long time window between close and open means there is more chance for another piece of important information to be released and distort the market's reaction. As a compromise, in the benchmark specification, events that occur outside trading hours are included if they are the "headline" story in the following morning news briefing. This implies they should be viewed as the most important European event that occurred overnight and thus, hopefully, represent what the market is reacting to at the open. A sensitivity analysis excluding all events outside the market opening period is presented in section 6. The market reaction to included events that occur outside of normal trading hours are calculated as the change from the previous close to 8:30am London time the morning of the first trading day after the event.

One may also be concerned with simultaneous events. Given the high frequency of the dataset this is a relatively low probability outcome in a trading day. Nonetheless, the following steps are taken to ensure markets are actually reacting to the event in question rather than other simultaneous news. The structure of the dataset means it is straightforward to single out any foreign event that overlaps with a local event and these would not be included in the proxy. Furthermore, if any local event occurs in a period when markets are closed, no foreign event that occurred in the same closed period would be included regardless if either event was a "headline" story in the following morning news briefing. The time of local and pan-European data releases are obtained from the *Bloomberg* economic calendar and events that would overlap with 20 minute windows around these releases are similarly omitted. Events that overlap with ECB decisions and press-conferences (i.e. anything that happens between 12:30pm and 14:50pm London time on the first Thursday of the month) are not included. Last, and most important, any country-specific event that overlaps with the announcement pan-European policy intervention, for example the commencement of the SMP programme in May 2010, is omitted. Such events are isolated using the ECB's time line of the crisis[15]

the 20th October 2011, Markets appeared to react as they have done to other episodes of violence in Greece but the news broke only gradually and the cause was revealed to be as a result of a heart attack only after some time.

[14]This window starts slightly earlier than suggested by Gurkaynak and Wright (2013) who recommend, in their primer on high frequency identification, that the window starts 5 minutes prior to the announcement (and 15 minutes after). A slightly wider window is used here because the events considered do not necessarily have a release fixed time (in contrast to a data release); thus there is no guarantee that the news wire has immediately picked up the announcement and a more conservative timing strategy seems appropriate. Some events, such as speeches and budget announcements last more than 20 minutes in which case the market reaction is considered to 20 minutes after the announcement ends (timed as the last relevant headline on the Bloomberg newswire). If a public event last more than 90 minutes it would not be considered timeable; however, this does not happen in the present version of the dataset. With closed door events/meetings, such as conferences and summits, the relevant time is taken as the start of the post-event press conferences which normally corresponds to the release of the press communique.

[15]See http://www.ecb.int/ecb/html/crisis.en.html

and are timed in an identical fashion to the country-specific events as described above.

Despite these steps to remove overlaps it is important to emphasise that it is impossible to be completely certain what drives the market move at any point in time. There may be news from outside the Euro-zone driving yields, as well as private information or rumours that cannot be picked up using the approach taken here. Furthermore, yields are driven by technical factors such as large transactions and variations in liquidity. That said, once one strips out coordinated policy actions and data releases, there little reason to think that any movement unrelated to the event is systematically correlated with other macroeconomic information. Therefore, unexplained market moves are unlikely to introduce bias but they will introduce measurement error which motivates the specification described in section 3.

The market reaction is defined as the change in the mid-yield to maturity on the benchmark 10 year sovereign bond for the country of interest. Note, this not the country where the event occurred; so if the country of interest is Italy and the event is in Greece, then the reaction would be the change in the Italian bond yield in the interval around the Greek event. The raw intra-day bond data is sourced at tick frequency from ICAP, a brokerage firm which gathers the data while intermediating wholesale trading between major commercial and investment banks. The tick data is converted into one minute 90% trimmed averages to remove any spikes at a very high frequency, the market reaction is calculated as changes in the averaged minutely series. The intra-day data covers the time when the London market is open.

The final step is to combine all the included high frequency events into a proxy variable that can be used for empirical analysis at a monthly frequency. This is done by aggregating all the market reactions around foreign events that occur over the course of the month. This is the version of the proxy, denoted m_t, that is plugged into the second stage of the model to identify the monthly structural shocks from the reduced form VAR.

4.2 Properties of the proxies

Completing the procedure described above leads to an amalgamation of policy announcements and political events relevant to the six countries over the course of the crisis period. The depth of coverage of events is encouraging, all the country specific events included in the ECB's time line of the crisis are captured; the same applies for Brutti and Saure (2013)'s narrative analysis of the crisis in Greece using a combination of crisis time lines compiled by private media outlets.

The depth of coverage and variety of events also means it is difficult to describe completely the narrative in a concise manner in the main text of this paper. Therefore, readers are referred to the online appendix for an exhaustive list of narrative events. The appendix of this document contains the relevant weblinks. However, some clarity over the type of events included in the narrative series can be seen by placing the events in loose classifications with accompanying examples:

1. *Domestic political events:* This is the broadest category and includes policy announcements by officials,

changes in government, votes in parliament, elections and important polls. Relevant news regarding scandals involving government officials, for example the donations scandal involving the Spanish Prime Minister in February 2013, are also included.

2. *Foreign interventions:* These refer to statements by foreign policymakers, particularly Troika members, about activities relating only to the specific crisis country. The various bailout agreements are obvious examples, as well as the approval of critical disbursements. Other examples included are the release of Troika reviews, decisions by the ECB regarding the acceptability of bonds as collateral and statements following Eurogroup meetings on specific countries.

3. *Technical events:* These refer to technical market news directly related to the sovereign bond market. This includes the results from important bond auctions (either from a liquidity perspective or due to their signaling value), pronouncements by credit rating agencies and decisions from the ISDA over whether certain policy actions (such as the bond buy back programme) constitute technical default.

4. *Domestic fiscal data:* These events relate to revisions in past fiscal numbers and future fiscal projections. News stories regarding statements from European and local authorities about the quality of data collection are also included. Note that events that relate to the standard monthly/quarterly data releases are not included.

5. *Domestic Instability:* Due to difficulty in timing the events, strikes and protests are generally not included as events unto themselves. However, what is included are the announcement about when strikes and protests will take place. Also included are violent events that occur during a protest; for example, in Greece, ministry buildings are stormed on several occasions. However, events of violence are included only if it is possible to find an objective time to assess the market's reaction.

Table 1 provides details of the number events identified in each crisis country. As one would expect given the country's role as the main instigator and protagonist of the crisis, Greece has the most events by raw number, followed by Spain and Italy reflecting their large size. The thirty-two Cypriot events are largely concentrated in March 2013 (Eighteen events occurred that month) indicative of the uncertainty surrounding the small country's bailout. The breakdown of events into the five classifications are similar across countries with the exception of Italy where foreign interventions are less prevalent but this is reasonable as the country is the only one not to receive a European bailout in some fashion over the time period.

This classification raises a point on the exogeneity assumptions underpinning the identification strategy. While it is plausible to argue that the market's reaction to domestic news in Greece, for example, not is driven by shocks in other Eurozone countries; it is harder to say the same for international policymakers who may be internalising the entire currency union when making their decisions. In the benchmark case, these events are included but a simple robustness check is to remove events falling into category two from the dataset and rerun the model. This

Table 1: Breakdown of event classification by country

Event Country:		Greece	Italy	Portugal	Spain	Ireland	Cyprus	Total
Number of Events:		287	148	120	163	105	32	855
Event Type (%)	Domestic Political	47.7	73.6	53.3	51.5	53.3	50.0	54.5
	Foreign Interventions	25.1	1.4	13.3	11.7	20.0	31.3	16.4
	Technical News	17.1	23.6	28.3	34.4	22.9	18.8	23.9
	Domestic Instability	7.0	1.4	3.3	1.8	1.0	0.0	3.5
	Fiscal Data	3.1	0.0	1.7	0.6	1.9	0.0	1.6

Notes: Total number of timeable, country specific events over the period July 2009-March 2013 as isolated using *EuroIntelligence* and *Bloomberg* European news summaries. Six crisis hit countries are considered and the total column represents the sum over the 6 countries. Events that overlap with data releases or other events as well as events that happen outside the trading hours are not excluded at this stage. Events are classified as in the main text. Percentages may not sum due to rounding.

is carried out in section 6 but turns out not to have any meaningful impact on the results as much of the variation in the proxy variable is driven by domestic news.

Greece and Cyprus are not included in the final empirical analysis, the latter due to its small size and the former due to issues of both data quality and the fact that debt was actually restructured. The final proxy variable is constructed only for Italy, Spain, Ireland and Portugal; Greek and Cypriot events are included in the proxies however.

Table 2 offers some descriptive statistics for the proxy variables in the four countries of interest. The first point to note is that the number of events that enter the proxy is substantially less than the total number of identified events across the six countries. This is because domestic events, events that overlap with other news and events outside the trading hours which are not "headline" news are excluded at this stage. The intra-day market reactions to event display similar statistical properties across countries which is somewhat surprising as on a daily basis Portuguese and Irish yields are more volatile. Greece is the largest contributor to the variation in the proxy; approximately 40% of the total absolute market movement around events is due to Greek news. This is share is roughly line with the relative number of events that are Greek origin - it is not the case that markets are reacting more strongly to Greek news on average but just that there are more Greek events to react to.

Figure 3 shows the proxy variable for the Italy (i.e. the aggregated change in the Italian yield around non-Italian events) plotted against the actual monthly change in the Italian yield; the graph is annotated with a selection of the events that correspond to major moves in the proxy variable. The correlation coefficient between these two series is given in the third row in table 2 and is strong at 0.75. From the point of view of the empirical strategy what matters is not the correlation with the actual change in the bond yield but with the residuals in the reduced form VAR; however, it is an encouraging sign nonetheless as the large movements in yields that show an alignment with proxy likely reflect the systemic shocks that the proxy is designed to identify. For the other included countries the correlation is not as strong; figure 4 presents the graphs for Spain (correlation: 0.65), Ireland[16] (correlation:0.55) and Portugal (correlation: 0.38). These correlation coefficients are all statistically significant at the 1% level.

[16] Irish tick data is not available between May 2011 and October 2011. In this period the Irish proxy is constructed using the daily

Table 2: Descriptive Statistics for the Proxy Variable

	Italy	Spain	Portugal	Ireland
Total Number of Included Events	452	391	479	393
Share outside trading hours (%)	23.7	13.0	21.1	30.8
Correl. with Actual Chg. in Bond Yield	0.76	0.65	0.38	0.55
Average Market Move (bp)	0.3	0.3	0.3	0.3
Average Absolute Market Move (bp)	2.0	1.8	2.0	1.6
Std. Dev. Market Move (bp)	3.4	3.1	3.5	3.1
Maximum Market Move (bp)	20.8	18.5	21.5	15.2
Minimum Market Move (bp)	-17.1	-23.1	-35.6	-33.2
Percentage of absolute change due to:				
Greece (%)	46.1	41.2	42.2	38.6
Italy (%)	0.0	24.3	17.9	20.1
Portugal (%)	11.6	12.7	0.0	14.7
Spain (%)	25.7	0.0	25.4	25.1
Ireland (%)	9.5	13.9	10.6	0.0
Cyprus (%)	7.0	7.9	4.0	1.5

Notes: Events included in the proxy variable satisfying the criteria in section4. Data period is July 2009 - March 2013. Irish proxy excludes May-October 2011 due to a break in intra-day data. Overlapping events, non-"headline" events outside the market open and domestic events are not included. The correlation is between the actual change in the bond yield in the month and the sum of market moves about events in that month. Market moves refer to change in local 10 year bond yields in a 20 minute window about an event. The percentage shares refer to the share total the absolute market move around events that can be attributed to events in a particular country. Percentages may not sum due to rounding.

A further important feature of the proxy data that is not apparent from table 2 is the relative importance of events. The variation in the proxy series is driven by large market reactions to a small share of the events rather than more moderate reactions to every piece of news. Figure 5 illustrates this by ordering events included proxies according to the square of the market reaction and then plotting the cumulative contribution of each ordered event to the total sum of squares to produce a graph analogous to a Lorenz curve. Reading off the chart it becomes apparent that the top 10% of events by absolute market move contribute somewhere between 80-90%, depending on the country, of the variance of the proxy. This implies that there are approximately 50 or so systemically important events in each proxy; still a large number but much less than the total number of identified events.

4.3 Predictability of the proxy variable

The critical identifying assumption underlying the narrative approach in this paper is that the proxy variable is only correlated with the structural shock of interest at time t. While it is not possible to conclusively rule out endogeneity from a statistical perspective, this section sets out to provide some evidence backing this assumption. The identification strategy used here relies on the local market move surrounding foreign events reflecting a "surprise" component of the announcement. Any systematic reaction to local macroeconomic shocks by foreign agents should already be anticipated by market participants at the time of the announcement. Furthermore, since foreign agents

change in yields during foreign events that were "headline" news, i.e. only major events.

do not have additional information private about local macroeconomic shocks than market participants the surprise component should not be correlated with those shocks. Thus the market move should be thought of as exogenous.

One corollary of this assumption is that the proxy variable should not be predictable and therefore should be unrelated to the market reaction to past events, both domestically and in other crisis countries, or the past realisation of local macroeconomic aggregates. This can be verified empirically. To do this, I set up a suite of univariate predictive regression models. As dependent variables I use the aggregate of market reaction to foreign events at weekly, biweekly and monthly frequencies; the weekly series is the sum of events all contained within the proxy that occur within a particular week etc. Exploring higher frequencies than monthly is necessary as market reactions within the month should also be unpredictable for the identification strategy to be justified. The predictive variables are lags of the dependent variable, lags of aggregated market reactions local events[17] at the same frequency and in the case of the monthly model, lags of the macroeconomic time series included in the VAR in section 5.1. The lag orders for the various models are determined automatically by selecting the order that maximises Bayesian Information Criterion up to a maximum of order of four months or equivalent. Estimation is conducted using least squares.

Table 3 presents F statistics and the adjusted R^2 as the output of interest from this analysis. In general, the message is as one would expect with rational and efficient markets: historical market moves and macroeconomic data have no predictive power over the market's reaction to current news. There is some evidence of predictive power in Ireland at a monthly frequency once macroeconomic aggregates are included but this is borderline only. The overall result is encouraging as it supports the identifying assumptions made. It also implies that the proxy behaves like a shock as opposed to being a predictable variable from a time series perspective.

A second empirical test is to gauge the extent to which the foreign events included in the proxy are a reaction to changes in local economic conditions, i.e. data releases, or ECB announcements. This could be thought of as a test of whether foreign events are uncorrelated to local macroeconomic shocks that are being captured by a data surprise. However, this is an imperfect test as the causality could run in the other direction; for example, events that raise yields may lower confidence and cause negative survey releases or provoke an ECB reaction. Nonetheless, the market should be anticipating these channels and we should not see any correlation between the market reaction to data or ECB meetings and events.

The timing of release of local economic data in each country is obtained from the *Bloomberg* economic calendar; the market reaction in terms of the local yield is considered in a twenty minute window about the release for the sake of consistency. For comparative purposes data releases are grouped in three categories to distinguish between their content: *(1)* Output releases correspond to industrial production, various confidence surveys and unemployment data; *(2)* Inflation releases: are the consumer and producer price releases; *(3)* Fiscal releases: correspond to monthly data on government finances from a cash accounting basis. Data releases and ECB meetings are at monthly frequency

[17] Local events that overlap with data releases, ECB meetings or pan-European policy interventions are omitted.

Table 3: Predictability of the proxy series

	Weekly Event Data	Biweekly Event Data	Monthly Event Data	Monthly Event Data + VAR data
Country	F-stat from Regression			
Italy	0.92	0.65	0.93	0.70
	(0.45)	(0.69)	(0.49)	(0.77)
Portugal	0.71	0.35	1.45	1.45
	(0.59)	(0.91)	(0.22)	(0.21)
Spain	0.70	0.97	0.50	1.68
	(0.59)	(0.45)	(0.80)	(0.13)
Ireland	0.50	0.94	0.45	1.91*
	(0.73)	(0.47)	(0.84)	(0.08)
	Adj R^2			
Italy	0.01	0.00	0.00	0.00
Portugal	0.02	0.00	0.06	0.16
Spain	0.00	0.00	0.00	0.22
Ireland	0.03	0.00	0.00	0.28*

Notes: Regression statistics from univariate predictive regressions of the proxy variable built from foreign events aggregated at different frequencies: weekly, biweekly, monthly. Sample Period: 1st July 2009 - 31st March 2013. Explanatory variables include a constant, lags of proxy variable and aggregated market reactions to local events. Lag order selected by the Bayesian Information Criterion on a country and frequency specific basis. Maximum lag order set to 16 weeks/8 fortnights/4 months. Monthly model with VAR date includes explanatory variables from the reduced form VAR model. *** Denotes significance at 1% level, ** 5%, *10%. Negative adjusted R^2 are normalised to zero.

so I do not consider weekly and biweekly aggregations in this case. Furthermore, a monthly measure of data surprises implies that releases on a quarterly basis are not considered for consistency reasons. Releases are considered by the month they refer to, so the industrial production release in August is data for June and therefore corresponds to the latter month; the results are not sensitive to this approach, carrying out the analysis on data by month of release leads to similar outcomes. For ECB meetings the market reaction is are considered from 12:30pm-14:50pm on the day of the meeting to capture both the interest rate announcement and the press conference. The market reactions are then aggregated for the month and the correlation with the foreign proxy and local events is presented in table 4. See appendix B.2 for detailed discussion of the included data releases.

Encouragingly, there is no general pattern of market reactions to local data being correlated with the proxy series, this is true as well for market reactions to ECB announcements. This suggests little systematic correlation between market reactions and that participants are internalising information when forming expectations as one would hope given the identification strategy. The same is almost true for the correlation between local events and local releases with the exception of the fiscal releases in Spain and Italy. This is an interesting result. In both countries the fiscal data is released with a two month lag; therefore, essentially what the correlations suggests is that the market reaction to local events in June, for example, will predict the reaction to the data when the fiscal numbers are released in August.

As this analysis is not the main focus of this paper, I have not investigated in depth how robust this correlation is. However, there is a potential interpretation. With output and inflation data governments typically receive little or no advanced information about the release; the data takes time to compile, statistics offices are independent

Table 4: Correlations between market reaction to events

	All Data	Output	Inflation	Fiscal	ECB Meetings
	Correlation with Foreign Events				
Italy	0.20	0.18	0.13	0.00	0.16
Portugal	-0.15	-0.10	-0.13	-	-0.22
Spain	-0.03	-0.04	0.01	0.03	0.21
Ireland	0.11	0.15	-0.07	-	-0.09
	Correlation with Local Events				
Italy	0.00	0.05	-0.03	-0.36***	0.00
Portugal	-0.16	-0.23	0.15	-	-0.02
Spain	-0.10	0.00	-0.26*	-0.44***	0.06
Ireland	0.20	0.19	0.07	-	0.23

Notes Sample correlation coefficients between market reactions in 20 minute windows aggregated into monthly series about events, data and ECB meetings. Foreign events refer to the proxy series as described in the main text; market reactions to local events are aggregated in a similar fashion excluding events that overlap with data releases, ECB meetings or pan-European policy interventions are omitted. Data releases organised by relevant month. Output: IP, Confidence Surveys, PMI's, Unemployment. Inflation: CPI and PPI, Fiscal: Monthly fiscal data and Government Debt (where applicable). Sample Period: July 2009 - March 2013. *** Denotes significance at 1% level, ** 5%, *10%.

and surveys often come from non-government sources. However, with cash government finance data policymaker's potentially have up-to-date information on the numbers long before the official release. Thus, the behaviour of policymaker's takes into account this data before the market is aware of it. The negative correlation is consistent with this line of reasoning. Events that lower the bond yield precede fiscal data that raises it; which suggests government's behave more conservatively when the fiscal data release will disappoint markets. A fully efficient market should take this into account and this perhaps provides a little evidence of the limits to market foresight. It also illustrates that point raised previously about private information. If there is an information asymmetry about the state of local macroeconomic conditions between market participants and the agent responsible for an event at the time the announcement, markets could potentially learn about these fundamentals from the announcement thus the identifying strategy used here would breakdown.

5 Empirical results

5.1 Specification

The panel VAR described above is run using a sample of four crisis hit Euro Area countries: Ireland, Italy, Portugal and Spain. The panel is balanced and covers a time period from 2007m1:2013m03. The sample begins before the start of the crisis to make sure that the time series of observations is not too short and that different macro environments are captured. However, this pre-crisis period also helps with the identification. The systemic shocks will be small (if even present) in the earlier period and the proxy variable takes a zero value prior to July-2009. Including this pre-crisis sample makes the identification equation consider the combinations of residuals seen before the crisis as less likely to be a systemic shock and thus allows for a cleaner identification of the shocks once the crisis commences. Including only the crisis sample period does not affect the results in a meaningful fashion.

In terms of the included variables, the starting point is the standard monetary policy VAR. Output is proxied on a monthly basis using the unemployment rate and a broad index of industrial production including the manufacturing, energy, utilities and construction sector. Prices are taken as the headline HICP reading. These datasets are provided by Eurostat. The monetary policy stance is captured using the 3-mth Eurepo rate. The series is computed as monthly average of the European Banking Federation's daily fixing and is not country specific. As a collateralised lending rate Eurepo ameliorates the heightened level of counterparty risk that has disrupted the interbank market since August 2007.

Given the context of this paper, it is natural to also include the borrowing cost of local sovereign. This is captured by the monthly average yield of the benchmark 10 year bond in each country. A measure to pick up the long-run risk free rate is also required. This is a difficult series to capture during the crisis. The German bond yield is inappropriate as it may have a negative converability premium embedded. Instead, the 10 year overnight interest rate swap (OIS), with EONIA as the floating leg, is used. This long term measure of risk free nominal interest rates also has the advantage of capturing, to an extent, the impact of the ECB's non-standard measures on the monetary policy stance.[18]

To assess the impact on elevated sovereign borrowing costs on the private sector cost of finance in a concise manner a composite measure of the interest rate on debt securities, bank lending and equities is used. This is computed internally by the ECB aggregates various sources in accordance with flows of new lending. Last, as a measure of the fiscal stance, the monthly general government primary balance is included as an annualised percentage of nominal GDP. The raw fiscal data is available on a quarterly basis from the flow of funds dataset available in Eurostat (net/lending or borrowing by the general government sector plus interest payments). Monthly fiscal data is constructed from the quarterly series using the regression based interpolation methodology of Mitchell et al (2005). This sets $N = 8$. Details of the data sources are laid out in appendix B.3. The set of deterministic variables, Z, is set to include only a constant for all countries. The trended series enter the VAR in log year-on-year difference; other series are included in levels. The lag length L is set to 4, capturing a complete quarter of data and an additional month.[19]

The posterior is simulated using 600000 draws from the MCMC sampler in the appendix; the first 100000 are discarded as a burn-in and the remaining chain is thinned by a factor of 50 leaving 10000 draws for inference. Results presented are the median of the 10000 retained draws and 95% uncertainty bands are computed using standard Bayesian Monte-Carlo methods.[20]

[18] As with Eurepo, the EONIA rate is less distorted by concerns by counterparty risk due to the short maturity of the loan. However, the OIS is still an imperfect measure. It was not clear how the contracts would be honoured in the event of Euro break-up for example. But it is less clear that the OIS has some embedded risk of redomination that

[19] Due to the short sample period and the medium scale of the model a parsimonious lag selection procedure is appropriate. The lag order is determined by testing up: starting by setting $L = 1$ and add more lags until the median estimated residuals display no serial correlation. The lag selected matches that selected using the Schwarz-Bayesian criterion assessed on the panel version of the VAR with homogenous slope coefficients and covariance matrix; equivalent to the model with maximum shrinkage ($\kappa \to \infty, \lambda_1 = 0$). So it is robust to alternative lag selection criteria.

[20] The independent Metropolis step has observed acceptance rates between 60-80%, depending on the particular step and the country in question. The algorithm appears to mix well; the standard diagnostic tests in Geweke (1992) are passed comfortably.

As with most VARs, the model can be used to produce three main results of interest; impulse response analysis provide an assessment of the propagation of systemic shocks to the included variables and variance decompositions give an indication of the relative importance in the shocks in explaining the fluctuations of the included variables. Of most interest is the counterfactual analysis, by identifying a time series of systemic shocks one can reconstruct the dataset omitting the impact these shocks have had on the included variables. This allows for an estimate of the contribution of the shocks to the borrowing costs of crisis hit countries.

5.2 Impulse responses and variance decompositions

Figure 6 presents the impulse responses to a systemic shock scaled to be consistent with a 100bps increase in 10 year bond yield on impact to the mean country model (constructed from the estimates of $\bar{\beta}$, $\bar{\Sigma}$ and $\bar{\Upsilon}$). Several features are apparent. systemic shocks propagate a little with regards to government bond yields with a peak of 1.2ppt after one month before declining steadily such that after 9 months the impact has dissipated. Part of the explanation for this correction likely lies in the soothing impact of policy: monetary easing following the shock, albeit with a lag, with a peak response of a 40bp decline in the Eurepo and 10 year OIS rates after 4 months. The unemployment response is statistically insignificant on impact but the shock propagates and leads to a peak of a 0.9% after 7 months. The response is also persistent, taking 18 months to return to zero. The industrial production (output) does not respond significantly on impact but the growth rate declines by 2ppt after 4 months. Inflation does not react on a statistically significant basis.

The pass through onto private borrowing costs is less than one-for-one: yields on private finance increase by only 40bp on impact and the effect is shorted lived as with the sovereign yield. An explanation for this lies with the composition of the data. The Eurozone private sector largely finances itself using bank loans and as a result these rates have a high weight in the composite cost of finance. However, bank lending rates are also sticky and slow to respond to market conditions. In section 6 a sensitivity analysis is conducted decomposing index into its components; the response of market based financing sources - equity and debt securities - are much greater.

Figure 7 presents the results of country specific models. The first thing to note is that the close similarity of the dynamics and response on impact suggests data is returning a model which is close the mean estimator; supporting a model which is close to a slope homogeneity assumption. In a model with improper priors it is not possible to construct standard Bayesian likelihood ratio tests on the pooled versus partially pooled model and a switch to a model with weakly informative priors can lead to unforeseen consequences - indeed it can bias the results away from the fully pooled case (see Gelman (2006)). An alternative suggested by Jarocinski (2010) is to rely on the deviance information criterion (DIC) of Spiegelhalter et al. (2002) which summarises the trade-off between the improvements in fit from not imposing homogenous parameters against the over parameterisation that may arise from partial pooling. The DIC is simply a sum of the expected deviance, a measure of fit related to the mean square error, and the effective number of parameters, which in the context of the hierarchical model with flat priors is close to the actual number of parameters the fully pooled model. See Spiegelhalter et al. (2002) for a full discussion of this criterion and how to calculate it. The smaller it is the better the model; in the case of the panel used here partial pooling returns a DIC of -14,503 for the reduced form VAR, while the fully pooled model returns a DIC of -14,477 and the country by country estimate returns -14,463. This suggests that partial pooling is effective even if inspection of the impulse responses confirms that the optimal degree of shrinkage is quite large and is close to a slope homogeneity assumption.

A couple of points of country heterogeneity are worth pointing out. Portugal has a positive inflation response. This could be interpreted as a potential working capital effect or may be analagous to the price puzzle. In terms of the fiscal response, on impact the systemic shock reduces the balance (an increase in the deficit) in all countries. This result is not distinguishable from zero in the mean the country model or in Italy and Portugal. But it is in Spain and Ireland. Since this is the primary balance this is not an automatic response to a higher interest burden. Instead, it is likely a reflection of lower revenues due to a weakening economy. Policy seems to quickly correct for this, the balance is back to zero after 5-6 months in both countries. There is not an over correction though; the increasing borrowing costs does not lead to a primary surplus. Is it is worth considering, therefore, if this is evidence of a lack of austerity in response to higher borrowing costs? Note that the primary balance is back to zero at the point where the unemployment response peaks. So if one defines an austerity package as an adjustment in the *cyclically adjusted* primary balance then austerity is taking place on that basis. It is just insufficient to overcome the negative fiscal consequences of a weaker economy that stemming from the shock. This chimes with the experience of Eurozone countries during the crisis who have struggled to consolidate their borrowing.

Figure 8 presents the variance decompositions, i.e. the portion of forecast error in each variable that is explained by the systemic shock at various horizons, for both the mean and country specific models. The decomposition reveals that on impact around 80% of the variation in the bond yield is caused by the shock which is substantial but not unsurprising given the crisis. At longer term horizons the importance of shock for the variance of yields seems to dissipate as is apparent with the impulse responses. This is perhaps reflective of patterns of intensification followed by periods of relative calm that have marked the Euro crisis. As a caveat, the shocks identified are designed to be orthogonal to local macroeconomic fundamentals; it may be that changes in the sovereign risk premia that stem from deteriorating local conditions are more persistent and thus more damaging for the local economy.

A second finding of note is that 45% forecast error of unemployment is explained at a forecast horizon of 6 months. This suggests there are more persistent consequences of the shock. And that this form of shock has contributed heavily to the variation in unemployment over the crisis period. Implying that exogenous variation in the bond yield via this systemic channel has had a meaningful effect on driving macroeconomic conditions in crisis hit countries..

In terms of other variables, systemic shocks explaining around 35% of the variation the private cost of finance on impact but the effect disappears surprisingly quickly at longer horizons. Little of the variation in the remaining series can be explained by the shock. Given, the unemployment response this is a little surprising. Indeed, the output (industrial production) variance decomposition is quite small in relation to the response of unemployment when one considers that they are both a proxy for cyclical conditions in the economy. An explanation could simply be that industrial production is a noisy series thus a greater a proportion of its forecast error is explained by its own volatility.

5.3 Counterfactual systemic premia

In order to gauge the extent of the contribution of the systemic component to changes in the bond yields in the crisis hit countries, a simple counterfactual analysis is carried out. For each draw from the posterior distribution of the parameter space, a time series of systemic shocks for each country is extracted. From there the corresponding draws of the slope coefficients, covariance matrix and identification equation coefficients can be used to remove the impact of these shocks from the data. This is equivalent to a counterfactual dataset where no systemic shocks occurred over the course of the whole sample. The relatively short-lived nature of the risk shocks means the analysis is not sensitive to exactly when counterfactual exercise commences and the start of the sample serves as a reasonable benchmark. A systemic premium can then by calculated as the difference between the true sovereign bond yield and its counterfactual equivalent. As this is exercise carried out for every draw from the posterior, the model produces a simulated distribution of the premia which enables the calculation of Bayesian confidence intervals.

Figure 9 presents the results of this analysis for the four countries in the sample; on the top panel is the actual versus the median counterfactual for the 10 year bond, the bottom has the implied premia and accompanying confidence intervals.

Several points standout. First, as one would expect, there is little evidence of a sustained systemic premium in any of the countries prior to the start of the crisis in 2009; the estimated premium fluctuates around zero and for the most part not statistically significant. However, once the crisis intensifies significant positive premia are apparent with peaks of the median estimate at 97bp for Spain, 127bp for Italy, 381bp for Portugal and 383bp in Ireland. Taking this into consideration between 40-60% of the peak to trough move in yields across the four country's can be explained by these shocks. This is order of magnitude is about what one would expect given the variance decomposition result and the model is consistent in this respect.

The pattern varies across countries; Italy suffers from two periods of high systemic premia first over the Autumn of 2011 and then in the Spring of 2012, both periods are contemporaneous with political instability Greece with the fall of the country's government followed by an indeterminate election. Spanish premia also peak around the Greek election and in November 2011 aren't significant at any other point. Portugal and Ireland suffer an extended run of elevated premia peaking around the summer of 2011 before declining relatively steadily towards the end of the sample.

By the end of the sample (March 2013) there are no positive statistically significant systemic premia in any of the countries considered. It was a reduction in the actual yield that achieved this reduction in the premia rather than a rise in the counterfactual; indeed, counterfactual yields appear to fall towards the end of the sample. The decline in the premia coincides, particularly in Italy and Spain, with a period of ECB action over the summer culminating in the announcement of Outright Monetary Transactions in September. The fall the counterfactual yield may reflect a monetary component to the shock that accompany those announcements.

In Portugal, the counterfactual suggests that the yield should be slightly higher than observed and this is to the extent that the systemic premia is statistically significant and negative near the end of the sample. This effect lasts only a month so it may be spurious. However it is worthwhile making the point that while this may seem counter-intuitive at first glance, a negative risk premia is not inexplicable. If one interprets the systemic premia as investors beliefs about the strength of multilateral cooperation and commitment to the Eurozone as an entity; financial markets can just as easily believe that strong policy interventions on a European level justify yields less than local fundamentals suggest. Indeed, that may be an interpretation for the origin story for the crisis.

In general what the results suggest is that the ECB's intervention was effective in the Autumn of 2012 in bringing yields back towards a more neutral setting and soothing the crisis.

6 Sensitivity analysis

This section presents three sensitivity checks for the empirical benchmark empirical results in the previous section. First, how different assumptions over the construction of the proxy can affect the results is explored. Next, the composite cost of private sector finance is decomposed into its elements: bank lending, debt securities and equity securities, all of which are included separately in the VAR. Last, a placebo is study is carried out to show that the results are not a feature of simply oversampling of the bond yield at a high frequency. For the sake of brevity the results of these analyses are presented just as impulse responses; variance decompositions and counterfactuals are a function of the impulse responses so it is sufficient to focus our attention on this aspect of the model. None of these checks alter the message presented above.

6.1 Alternative proxies

Construction of the proxy in section 4 involved several assumptions that should tested for robustness. The first alternative proxy is constructed one hour windows rather than 20 minute windows. This alternative gives markets longer to react to events and the cost of potentially capturing moves unrelated to the event in question. The second alternative proxy dispenses intraday day data altogether and just considers the daily change in the yield on days where then is a "headline" event; i.e. the ones that make the top of the morning news briefing. Essentially, this looks at the daily reaction to what journalists perceive as the most important events. The third alternative proxy excludes all events that happen outside of traded hours. The fourth drops all events involving foreign interventions for the reasons discussed in section 4.

The last alternative proxy just looks at events that happen in the first week of the month. This is designed to strengthen identification assumptions further. If the market reaction to an event is partly a function of local structural economic shocks, by looking just at events early in the month there is less opportunity for foreign agents to react the local shocks that happen that month and act upon them. This is a similar line of reasoning that one would use regarding a causal ordering for SVAR identification. By looking at the first week, one can make a stronger case for the proxy "moving first" as it were.

All the alternative proxies remove events that overlap with local data, local events, ECB meetings and pan-European events as in the benchmark case. The specification of underlying reduced form model is held constant.

Figure 10 present the median responses of the mean country model under the alternative proxy definitions overlaid on the benchmark specification with corresponding confidence intervals. The top set of impulses contain the first two alternative proxies; the bottom set the last three. The alternative proxies are closely correlated with the benchmark, so unsurprisingly the results turn out much the same regardless of the proxy used. There are quantitative differences but qualitatively the message is the same. Furthermore, all the median alternative impulse responses are within confidence set of the benchmark.

6.2 Alternative financing sources

Figure 11 presents the mean country impulse response functions for the decomposed elements of the private cost of finance. As the equity yield is not consistency available over the sample this is substituted for using the year-on-year change in the headline equity price for each country sourced from Eurostat. The average interest rate on loans and corporate bonds are sourced from the ECB composite cost of finance. The responses to the non-financial variables are almost indistinguishable from the benchmark case so are not shown for the sake of compactness. In terms of the different financing sources, there is some heterogeneity across countries but in general: there is strong pass through to corporate yields with a response over 1%. Equity prices also fall with a peak decline of 7%. Loan rates barely react in contrast. Rising by only 0.25ppt.

6.3 Placebo study

Given there are approximately 400 events included in the proxy a valid concern could be that the approach taken is simply sampling the actual changes in the yield and producing a noisy measure of the overall change in the month rather than picking up any particular shock. This would then be equivalent to just treating the yield as contemporaneously exogenous but badly measured. With 400 events in the proxy approximately 4% of the trading time during the sample period is covered by the event windows chosen. However, the question of whether the 4% coverage is sufficiently small to rule out this potential sampling problem is not possible to answer from a theoretical basis. Hence I attempt to verify it empirically using a placebo study.

The solution taken is to simply recreate the benchmark proxy in an identical fashion using the same event time but one trading day previously. This retains the something close to the same distribution of events across the months in the sample but gives a different set of placebo market reactions not in the vicinity of the original event. These new windows may overlap with other windows in the benchmark proxy, if events happen on sequential days, but this is exactly the sort of sampling problem we wish to account for so no adjustment is made. For the same reason, and unlike in the benchmark case, no attempt to drop any other overlaps, for example with local data or local events, are made.

Figure 12 presents a comparison of the mean country impulses based on the placebo proxy against that with the benchmark proxy holding the reduced form specification constant. The first point to note is that none of the responses are distinguishable from zero (excluding the response of the sovereign yield which has to be due to the scaling assumption). Second, although the pattern of median impulses looks somewhat similar the scale of the error bands iscompletely different to the benchmark. This reflects the inaccuracy of the identification stage of the model when using the placebo. The response on impact depends on the *relative* size of the parameters in $\bar{\Upsilon}$. These estimates are close to zero as the reduced form residuals have almost no explanatory power over the proxy. However, when they are rescaled to be consistent with a 100bp increase in yields the size of the response can become overly large as this involves dividing through by a parameter which is itself close to zero.

The poor performance of the placebo suggests that the approach taken is not equivalent to just generating a noisy measure of the bond yield at time t and there is real information contained within the benchmark proxy.

7 Conclusion

This paper attempts to identify and empirically quantify the size of systemic shocks that struck Eurozone countries over the course of the recent financial crisis. The identification is structured around a narrative procedure where the high frequency market reaction around foreign events is used to proxy changes in the systemic premium. This methodology is then coupled with a panel reduced form VAR model for the purposes of impulse response analysis and counterfactuals.

The headline results based around counterfactual analysis was that there was substantial systemic premia apparent in Euro-zone countries over the course of the crisis with peak levels in the order of magnitude in the hundreds of basis points but recent events have soothed this risk premia bringing it back to almost neutral levels coinciding with a string of interventions by the ECB. The impulse response analysis suggests that innovations systemic was an important driver of financial markets over the crisis period but the shocks are relatively short-lived. The innovations to the yield also lead to a relatively long-lived change in unemployment and appear to contribute substantially to the forecast error in that series. The suggests that systemic shocks were an important driver in macroeconomic conditions during the crisis.

References

[1] G. Amisano and C. Giannini. *Topics in Structural VAR Econometrics.* Springer, 1997.

[2] Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Clara Vega. Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, 93(1):38–62, March 2003.

[3] Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Clara Vega. Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics*, 73(2):251–277, November 2007.

[4] Magnus Andersson. Using intraday data to gauge financial market responses to fed and ecb monetary policy decisions. Working Paper Series 726, European Central Bank, February 2007.

[5] Stern H. Andrew Gelman, Carlin J. and Rubin D. *Bayesian Data Analysis.* Chapman & Hall, 2 edition, 2003.

[6] Andrew Ang and Francis A. Longstaff. Systemic sovereign credit risk: Lessons from the u.s. and europe. NBER Working Papers 16982, National Bureau of Economic Research, Inc, April 2011.

[7] Patrick Bolton and Olivier Jeanne. Sovereign default risk and bank fragility in financially integrated economies. *IMF Economic Review*, 59(2):162–194, June 2011.

[8] Filippo Brutti and Philip U. Saure. Transmission of sovereign risk in the euro crisis. Mimeo, March 2012.

[9] Willem Hendrik Buiter, Giancarlo Corsetti, and Paolo A. Pesenti. Interpreting the erm crisis: Country-specific and systemic issues. Princeton studies in international economics, International Economics Section, Departement of Economics Princeton University,, 1998.

[10] Craig Burnside, Martin Eichenbaum, and Jonas D. M. Fisher. Fiscal shocks and their consequences. *Journal of Economic Theory*, 115(1):89–117, March 2004.

[11] Fabio Canova and Matteo Ciccarelli. Panel vector autoregressive models a survey. Technical Report 1507, European Central Bank Working Paper Series, January 2013.

[12] James Cloyne. What are the effects of tax changes in the united kingdom? new evidence from a narrative evaluation. CESifo Working Paper Series 3433, CESifo Group Munich, April 2011.

[13] Vitor Constancio. Contagion and the european debt crisis. Financial Stability Review 16, Banque de France, April 2012.

[14] Barry Eichengreen, Andrew K. Rose, and Charles Wyplosz. Contagious currency crises. NBER Working Papers 5681, National Bureau of Economic Research, Inc, Jul 1996.

[15] Barry Eichengreen and Charles Wyplosz. The unstable ems. *Brookings Papers on Economic Activity*, 24(1):51–144, 1993.

[16] Carlo Favero and Francesco Giavazzi. Measuring tax multipliers: The narrative method in fiscal vars. *American Economic Journal: Economic Policy*, 4(2):69–94, May 2012.

[17] Kristin Forbes. The "big c": Identifying and mitigating contagion. Working Paper 4970-12, MIT Sloan School of Management, August 2012.

[18] Andrew Gelman. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, 1(3):515–533, 2006.

[19] Stefan Gerlach and Frank Smets. Contagious speculative attacks. *European Journal of Political Economy*, 11(1):45–63, March 1995.

[20] J Geweke. Bayesian treatment of the independent student- t linear model. *Journal of Applied Econometrics*, 8(S):S19–40, Suppl. De 1993.

[21] Morris Goldstein. *The Asian Financial Crisis: Causes, Cures, and Systematic Implications*. Institute for International Economics, Washington D.C., USA, 1998.

[22] Shafik Hebous. The effects of discretionary fiscal policy on macroeconomic aggregates: A reappraisal. *Journal of Economic Surveys*, 25(4):674–707, 09 2011.

[23] Marek Jarocinski. Responses to monetary policy shocks in the east and the west of europe: a comparison. *Journal of Applied Econometrics*, 25(5):833–868, 2010.

[24] Olivier Jeanne and Paul Masson. Currency crises, sunspots and markov-switching regimes. *Journal of International Economics*, 50(2):327–350, April 2000.

[25] Graciela L. Kaminsky and Sergio L. Schmukler. What triggers market jitters?: A chronicle of the asian crisis. *Journal of International Money and Finance*, 18(4):537–560, August 1999.

[26] Delroy Hunter Karan Bhanot, Natasha Burns and Michael Williams. Was there contagion in eurozone sovereign bond markets during the greek debt crisis? Working Paper 006FIN-73-2012, The University of Texas at San Antonio, College of Business, March 2012.

[27] Robert B Litterman. Forecasting with bayesian vector autoregressions-five years of experience. *Journal of Business & Economic Statistics*, 4(1):25–38, January 1986.

[28] Karel Mertens and Morten O Ravn. The dynamic effects of personal and corporate income tax changes in the united states. *American Economic Review*, page Forthcoming, 2013.

[29] Valerie A. Ramey. Identifying government spending shocks: It's all in the timing. *The Quarterly Journal of Economics*, 126(1):1–50, 2011.

[30] Valerie A. Ramey and Matthew D. Shapiro. Costly capital reallocation and the effects of government spending. *Carnegie-Rochester Conference Series on Public Policy*, 48(1):145–194, June 1998.

[31] David Lando Rene Kallestrup and Agatha Murgoci. Financial sector linkages and the dynamics of bank and sovereign credit spreads. Mimeo, March 2012.

[32] Frank de Jong Roel Beetsma, Massimo Giuliodori and Daniel Widijanto. Spread the news: How the crisis affected the impact of news on the sovereign bond markets. mimeo, April 2012.

[33] Christina D. Romer and David H. Romer. Does monetary policy matter? a new test in the spirit of friedman and schwartz. In *NBER Macroeconomics Annual 1989, Volume 4*, NBER Chapters, pages 121–184. National Bureau of Economic Research, Inc, December 1989.

[34] Christina D. Romer and David H. Romer. The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review*, 100(3):763–801, June 2010.

[35] Roberto A. De Santis. The euro area sovereign debt crisis: safe haven, credit rating agencies and the spread of the fever from greece, ireland and portugal. Working Paper Series 1419, European Central Bank, February 2012.

[36] David J. Spiegelhalter, Nicola G. Best, Bradley P. Carlin, and Angelika Van Der Linde. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639, 2002.

[37] James H. Stock and Mark W. Watson. Disentangling the channels of the 2007-2009 recession. NBER Working Papers 18094, National Bureau of Economic Research, Inc, May 2012.

A The conditional distribution of the instrument.

This appendix details the distribution of the proxy variable under random censoring and derives the linear relationship with the reduce form errors. Recall that the density of m_{dt}, given the true structural shock, can be expressed as:

$$P(m_d|\psi, \varepsilon_d, p) = \left((2\pi\sigma_v^2)^{-\frac{1}{2}} exp \left\{ -\frac{1}{2} \left(\frac{m_d - \psi\varepsilon_d}{\sigma_v} \right)^2 \right\} p \right)^{1-I(m_d=0)} (1-p)^{I(m_d=0)} \tag{15}$$

Dropping the time subscripts, this density implies the moment generating function is given by:

$$M_{m_d}(t) = \int_{-\infty}^{\infty} exp \{tm_d\} \left(p(2\pi\sigma^2)^{-\frac{1}{2}} exp \left\{ -\frac{1}{2} \left(\frac{m_d - \psi\varepsilon_d}{\sigma_v} \right)^2 \right\} \right)^{1-I(m_d=0)} (1-p)^{I(m_d=0)} dm_d \tag{16}$$

Solving the integral yields:

$$M_{m_d}(t) = (1-p) + pexp\left\{t\psi\varepsilon_d + \frac{t^2}{2}\sigma_v^2\right\} \tag{17}$$

Using the independence of observations m_d; the moment generating function of m is simply the product of the M_{m_d} over d:

$$M_m(t) = \prod_{d=1}^{M}((1-p) + p.exp\left\{t\psi\varepsilon_d + \frac{t^2}{2}\sigma_v^2\right\}) \tag{18}$$

The moments of m follow $E(m^n) = \frac{\partial^n M_m(0)}{\partial t^n}$. For exposition it is useful to define $exp\left\{t\psi\varepsilon_d + \frac{t^2}{2}\sigma_v^2\right\} = x_d$, noting that $x_d = 1$ if $t = 0$. Using the product rule:

$$\frac{\partial M_m(t)}{\partial t} = \sum_{d}^{M}(p(\psi\varepsilon_d + t\sigma_v^2)x_d \prod_{s\neq d}^{M}((1-p) + px_s) \tag{19}$$

which implies the first moment is given by:

$$E(m|\varepsilon_d) = p\psi\sum_{d}^{M}\varepsilon_d = p\psi a_1 u = E(m|u) \tag{20}$$

To calculate the second moment, the second differential is:

$$\frac{\partial^2 M_m(t)}{\partial t^2} = \sum_{d}^{M}\left[p(\sigma_v^2 + (\psi\varepsilon_d + t\sigma_v^2)^2)x_d \prod_{s\neq d}^{M}((1-p) + px_s) + p^2(\psi\varepsilon_d + t\sigma_v^2)x_d \sum_{s\neq d}^{M}\left\{((\psi\varepsilon_s + t\sigma_v^2)x_s \prod_{r\neq s,d}^{M}((1-p) + px_r)\right\}\right]$$

hence:

$$E(m^2|\varepsilon_d) = p\sum_{d}^{M}\left[(\sigma_v^2 + (\psi^2\varepsilon_d^2)) + p^2\psi^2\varepsilon_d \sum_{s\neq d}^{M}\{\varepsilon_s\}\right]$$

Since the individual daily observations are unobservable, what we care about is $E(m^2|u)$, which is simply:

$$E(m^2|u) = pM\sigma_v^2 + p^2\psi^2\sum_{d}^{M}\left[E(\varepsilon_d^2|u) + \sum_{s\neq d}^{M}E(\varepsilon_s\varepsilon_d|u)\right]$$

Noting that ε_d and u are Gaussian, with $Cov(\varepsilon_d, u) = \frac{a_1\Sigma_u}{M}$. Hence, one can write $E(\varepsilon_d|u) = \frac{a_1 u}{M}$, $Var(\varepsilon_d|u) = 1 - \frac{a_1\Sigma_u\Sigma_u^{-1}\Sigma_u a_1^{'}}{M^2} = (M-1)/M^2$ and $Cov(\varepsilon_s\varepsilon_d|u) = -1/M^2$. Thus $E(\varepsilon_d^2|u) = \frac{a_1uu'a_1^{'}+(M-1)}{M^2}$ and $E(\varepsilon_d\varepsilon_s|u) = \frac{a_1uu'a_1^{'}-1}{M^2}$. From here it is obvious that $Var(m|u) = pM\sigma_v^2$.

B Data and Sources

B.1 The narrative proxy series

Due to the size of the narrative dataset it is more straightforward to communicate the included events and the size of the market reaction to them in a spreadsheet format rather than via a document. Therefore, I provide the following links to online datasets that describe the narrative series of events and the market reaction to them. Unfortunately, due to licensing issues the underlying tick data cannot be made available.

1. This **spreadsheet of events** provides the following information:

   (a) of all the identified events in Cyprus, Greece, Ireland, Italy, Portugal and Spain.

   (b) the source news summary corresponding to each event.

   (c) the time that event occured, identified as the first headline relevant to the event on the Bloomberg news wire. If applicable, an end time is included related to the last headline relevant to the release. This is relevant to events that are extended announcements over several minutes such as speeches.

   (d) The classification of the events into the categories as laid out in the main text.

   (e) Whether or not the event was the top story (headline) in the morning news briefing.

   (f) For the case of Ireland, Italy, Portugal and Spain the high frequency bond market reaction in the relevant window around the event. It is worth noting that these are raw market reactions. No attempt has been made here to check for overlapping events that may provoke large unexplained reactions.

2. The calculations of the proxies themselves can be the following spreadsheets. These sheets detail the high frequency bond market reaction to every included event and calculate whether there is an overlap with another event. The various proxies used (including those for robustness checks) are all included.

   (a) The **Irish proxy**

   (b) The **Italian proxy**

   (c) The **Portuguese proxy**

   (d) The **Spanish proxy**

B.2 Data releases

Data releases serve two purposes in the main analysis of this paper. First, events which overlap with a twenty minute window about local data releases are excluded from the proxy. Second, the reaction of the market to data releases is aggregated for each month and compared to the reaction about events as a robust check (see table 5 in the main paper). Here, data releases considered are listed. For the purposes of table 5, those marked with a

1 are used as output releases, those as 2 are inflation releases and those marked as 3 are fiscal releases - the sum of three corresponds to the all data column. Note that only series released monthly are included in this analysis (which is why GDP is not used for example). The first release is always used rather than the final revised number. Descriptions here correspond to those listed on the Bloomberg Economic Calendar.

- **Italian Data Releases:** Budget Balance (3), Business Confidence (1), Consumer Confidence (1), CPI Final, CPI Preliminary (2), Current Account, Deficit to GDP, GDP final, GDP Preliminary, General Government Debt, Hourly Wages, Industrial Orders (1), Industrial Production (1), Industrial Sales (1), Labor Costs, New Car Registrations, PMI Manufacturing (1), PMI Services (1), PPI (2), Retail Sales (1), Trade Balance , Unemployment Rate (1).

- **Spanish Data Releases:** CPI Final, CPI Preliminary (2), Current Account, GDP final, GDP Preliminary, House Price Index, House transactions (1), Industrial Output (1), Labour Costs, Mortgages on Houses, Producer Prices (2), Retail Sales Volumes (1), Spain Budget Balance (3), Spain Business Confidence, Spain Consumer Confidence (1), Spain Manufacturing PMI (1), Spain Services PMI (1), Total Housing Permits, Trade Balance, Unemployment (1), Unemployment.

- **Portuguese Data Releases:** Construction Works Index, Consumer Confidence, Consumer Price Index, Current Account, Economic Climate Indicator, GDP (YoY) final, GDP Preliminary, Industrial Production (1), Industrial sales (1), Labour Costs, Producer Prices, Retail Sales, Trade Balance, Unemployment Rate.

- **Irish Data Releases:** Consumer Confidence (1), CPI (2), Current Account Balance, GDP, Industrial Production (1), Live Register Level, Manufacturing PMI (1), New Vehicle Licences (1), PPI (2), Property Prices, Retail Sales Volumes (1), Services PMI (1), Trade Balance, Unemployment Rate.

The following important international and European data releases are also used to exclude overlapping events from the proxy (admittedly this an abritrary selection):

- **International Releases:** Eurozone Services PMI, Eurozone Manufacturing PMI, German IFO, European Union Fiscal Data, US Labour Market (non-Farm Payrolls), European Comission Confidence Surveys, Eurozone GDP Final, Eurozone GDP Preliminary.

B.3 VAR Data sources:

10 year sovereign bonds: The intraday data only extends back to July-2009. For the complete VAR sample, from 2007 to 2013 , the monthly average of the daily yield on the 10 year benchmark sovereign bond on Bloomberg is used instead. The correlation between this series and the intraday yield at close is greater than 0.95 for all four countries on a daily basis from July 2009 to March 2013. The relevant Bloomberg codes are: *Italy: GTITL10Y; Spain: GTESP10Y; Ireland: GTIEP10Y; Portugal: GTPTE10Y.*

Industrial Production: The industrial production index is sourced from Eurostat. The broadest index possible is used, including the manufacturing, energy and construction sectors (*Eurostat code: sts_inpr_m*). The underlying data is presented as an index with 2005 as a base year.

Consumer Prices: The harmonised index of consumer prices (HICP) is sourced from Eurostat. The headline index is used - all items including the food and energy (*Eurostat code: prc_hicp_midx*). The underlying data is presented as an index with 2005 as a base year.

Unemployment: The harmonised unemployment rates are sourced from Eurostat and expressed as a percent of the labour force (*Eurostat code: une_rt_m*).

3 month Eurepo Rate: The 3 month Eurepo rate is measured as the monthly average of the daily Eurepo fixing by the European Banking Federation *(http://www.euribor-ebf.eu/eurepo-org/about-eurepo.html).*

10 year overnight index swap (OIS) rate: The 10 year OIS rate is measured as the monthly average of the daily series compiled by Bloomberg from over-the-counter brokers in the OIS market *(Bloomberg code: EUSA10 CMPN)*

Private Sector Cost of Finance: This is computed internally by the Capital Markets/Financial Structure division of the ECB for each country in the Euro Area. It is the amalgmation of the cost of loans to the non-financial private sector, the cost of corporate bonds and the cost of equity (the latter two apply to non-financial corporations only). The cost of the three sources of finance are weighted using flows of new liability acquisition by non-financial private sector. This creates an average cost of finance faced by the private sector analagous to an overall interest rate on financial liabilities. The cost can be decomposed into its constitutiuent components as is shown in the robustness analysis. The cost equity is not available consistently throughout the sample so equity prices are used instead.

Equity Prices: The main equity price index for each country is sourced from Eurostat as a monthly average. The indices are rebased such that 2005=100 (*Eurostat code: mny_stk_spy_m*). The country indices are better known as: *Italian FTSE MIB Index, Portuguese Stock Index 20, Irish Stock Exchange Equity Overall Index, Spanish Association of Stock Exchanges Index.*

Primary Fiscal Balances: This is the most complex input into the VAR. As no official monthly data for fiscal balances exists on an accurals basis, one is constructed using interpolation methods. Since fiscal numbers are available on a cash accounting basis at monthly frequency, these series serve as natural interpolands. The quarterly primary fiscal balance is defined as the net lending/borrowing of the general government sector plus interest payments. This is sourced from the Eurostat flow of funds database; the fiscal balance is created using the non-financial accounts (*Eurostat code: nasq_nf_tr*). Flow of funds data are in millions of nominal euros and are not seasonally adjusted. The unadjusted balance as a percentage of GDP is calculated by dividing through by quarterly, nominal GDP from Eurostat in millions of Euros (*Eurostat code: namq_gdp_c*). The adjusted quarterly balance is created by placing this data through an X.12 filter. Monthly nominal GDP is constructed by placing a cubic spline through the quarterly series in each country; since monthly GDP is the relatively stable denominator in the monthly fiscal series this choice of interpolation technique is little importance. The interpolation procedure for the fiscal balance is conducted in percentage of GDP terms using the regression based procedure in Mitchell et al (2005). The interpolation regression estimated using maximum likelihood; it is assumed the underlying fiscal balance is an ARX(1,1) on a monthly basis restricted such that the sum of the monthly balances equal the quarterly figure. Experiments with alternative lag structures revealled little sensitivity to alternative specifications. The differences across countries in the availability of monthly fiscal data across countries mean that the interpolands and sample periods are country specific:

- **Italy:** The first interpoland is monthly the central government balance less central government interest payments (both millions of Euros, calculated on a cash accounting basis and non-seasonally adjusted). The second interpoland is the change in general government debt (millions of Euros, non-seasonally adjusted). Both interpolands are divided through by monthly nominal GDP and seasonally adjusted using an X.12 procedure. Both series are sourced from the Italian Finance Ministry. The sample period for the estimation is January 2000 to March 2013. The model is extended beyond the sample for the VAR to improve the quality of the fit.

- **Spain:** The first interpoland is monthly the central primary government balance (in millions of Euros, calculated on a accurals basis and non-seasonally adjusted). The second interpoland is the monthly change in central government gross debt outstanding (millions of Euros, non-seasonally adjusted). Both interpolands are divided through by monthly nominal GDP and seasonally adjusted using an X.12 procedure. Both series are sourced from the Spanish Finance Ministry. The sample period for the estimation is January 1999 to March 2013.

- **Portugal:** The first interpoland is monthly the central government balance (bin millions of Euros, calculated on a cash accounting basis and non-seasonally adjusted). The second interpoland is the change in general government debt (millions of Euros, non-seasonally adjusted). Both interpolands are divided through by monthly nominal GDP and seasonally adjusted using an X.12 procedure. Both series are sourced from the Portuguese Finance Ministry. The sample period for the estimation is January 2000 to March 2013.

- **Ireland:** There is a single interpoland which is monthly the Exchequer surplus, equivalent to the central government balance, (in millions of Euros, calculated on a cash accounting basis and non-seasonally adjusted).

The interpolands are divided through by monthly nominal GDP and seasonally adjusted using an X.12 procedure. The series is sourced from the Irish Finance Ministry. The sample period for the estimation is January 2000 to March 2013.

The interpolation procedure appears to work well, there are no unusually large spikes in the monthly series and they interpolated figures do not resemble the output from a deterministic interpolation procedure, suggesting the monthly interpolands are informative.

C MCMC Sampler

Define the parameter space in the model as:

$$\Theta = \{\beta_1, \ldots, \beta_C, \Sigma_{1,u}, \ldots, \Sigma_{C,u}, \gamma_1, \ldots, \gamma_C, \bar{\beta}, \lambda_1, \bar{S}, \Upsilon_1, \ldots, \Upsilon_C, \sigma_{1\omega}, \ldots, \sigma_{C\omega}, \bar{\Upsilon}, \lambda_2\}.$$

It is useful to also divide the parameter space into those referring to the reduce form model and those referring to the identification equations:

$$\Theta_1 = \{\beta_1, \ldots, \beta_C, \Sigma_{1,u}, \ldots, \Sigma_{C,u}, \gamma_1, \ldots, \gamma_C, \bar{\beta}, \lambda_1\}$$

$$\Theta_2 = \{\Upsilon_1, \ldots, \Upsilon_C, \sigma_{1\omega}, \ldots, \sigma_{C\omega}, \bar{\Upsilon}, \lambda_2\}.$$

To simplify the notion define the set of data used in the reduced form VAR as $Y = \{Y_1, \ldots, Y_C, X_1, \ldots, X_C, Z_1, \ldots, Z_C\}$ and the proxy variables as $\mathcal{M} = \{\mathcal{M}_1, \ldots, \mathcal{M}_C\}$. Define the data matrix of reduced form VAR residuals,U_c, as: $U_c = Y_c - X_c B_c - Z_c \Gamma_c$. By Bayes rules the likelihood of the data is proportional to the product of the likelihood of the proxy variables conditional on both the reduced form model and the data and the likelihood of the reduced form model given the data: $p(\mathcal{M}, Y|\Theta) = p(\mathcal{M}|Y, \Theta)p(Y|\Theta)$.

In terms of the former, as is standard with linear models with Student-t errors, by expanding the parameter space it is possible to rewrite the conditional density as a Gaussian regression model with heteroskedastic errors:

$$\mathcal{M}_c|Y_c \sim N(U_c \Upsilon_c, \sigma^2_{\omega c} \Xi_c)$$

Where the matrix Ξ_c is a diagonal vector of unknown parameters equal to $diag\{\xi_{c1}, \ldots, \xi_{cT})$. With the prior assumption $\nu/\xi_{c1} \sim \chi^2(\nu)$, where ν are the degrees of freedom on the student-t errors, Geweke(1993) shows this is equivalent to a linear model with t-errors as described in the main text. The intuition follows from the definition of the t-distribution as a ratio between a normal and a χ^2. Define the residuals from the proxy model as:

$V_c = (\mathcal{M}_c - U_c\Upsilon_c)$.

The likelihood of the reduced form VAR model, $p(Y|\Theta)$, is the product of the country specific Gaussian distributions as defined in equation 2. Combining these two densities with the priors gives a joint posterior density, $p(\mathcal{M}, Y|\Theta)$, proportional to:

$$|\bar{S}|^{\frac{C\kappa-(N+1)}{2}}(\lambda_2\lambda_1)^{-\frac{v+2+C}{2}}\prod_c\left(\sigma_{c\omega}^{-2}|\Sigma_{c,u}|^{-\frac{T+\kappa+N+1}{2}}\right) \quad \ldots$$

$$exp\left\{-\frac{1}{2}\left(\sum_c\left\{\frac{s}{\lambda_1}+\frac{s}{\lambda_2}+tr\left[(U_c^{'}U_c\Sigma_c^{-1})+\bar{S}\Sigma_c^{-1}\right]+(\beta_c-\bar{\beta})'(\lambda_1L_{1c})^{-1}(\beta_c-\bar{\beta})+(\Upsilon_c-\bar{\Upsilon})'(\lambda_2L_{2c})^{-1}(\Upsilon_c-\bar{\Upsilon})+\sigma_{c\omega}^{-2}V_c'\Xi_c^{-1}V_c-\nu tr(\Xi_c^{-1})\right\}\right)\right\}$$

For most parameters in the model the conditional densities used in the Gibbs Sampler are in the form of classical distributions. However, as discussed in the main text the presence of the proxy variable alters the slope coefficient estimates in the reduced form VAR.

$$p(\beta_c|Y,\mathcal{M},\Theta\setminus\beta_c) = exp\left\{-\frac{1}{2}\left(\left\{tr\left[(U_c^{'}U_c\Sigma_c^{-1})\right]+(\beta_c-\bar{\beta})'(\lambda_1L_{1c})^{-1}(\beta_c-\bar{\beta})+\sigma_{c\omega}^{-2}tr(V_c'V_c\Xi_c^{-1})\right\}\right)\right\}$$

This conditional density is not proportional to a standard distribution. However, the density of the slope coefficients conditional only on the reduced form parameters in Gaussian:

$$p(\beta_c|Y,\Theta_1\setminus\beta_c) \propto N(D_c^{-1}d_c, D_c^{-1}) \tag{21}$$

where

$$D_c = \Sigma_{c,u}^{-1}\otimes X_c'X_c + \lambda_1^{-1}L_{1c}^{-1}$$

$$d_c = (\Sigma_{c,u}^{-1}\otimes X_c')vec(Y_c - Z_c\Gamma_c) + \lambda_1^{-1}L_{1c}^{-1}\bar{\beta}$$

The density represented in equation 21 is that of the coefficient estimates ignoring the additional information contained in $\mathcal{M}$; this is used as a candidate distribution for an independent Metropolis-Hastings step within the Gibbs-Sampler.

The coefficients on the deterministic terms in the reduced form VAR have the same problem, the conditional density is given by:

$$p(\gamma_c|Y,\mathcal{M},\Theta\setminus\gamma_c) = exp\left\{-\frac{1}{2}\left(\left\{tr\left[(U_c^{'}U_c\Sigma_c^{-1})\right]+\sigma_{c\omega}^{-2}tr(V_c'V_c\Xi_c^{-1})\right\}\right)\right\}$$

Which is a non-standard distribution. However, the same solution exists as with the slope coefficients is possible.

The density of the deterministic coefficients conditional only on the reduce for parameters is of the form:

$$p(\gamma_c|Y,\Theta_1 \setminus \gamma_c) \propto N(F_c^{-1}f_c, F_c^{-1}) \tag{22}$$

where

$$F_c = \Sigma_c^{-1} \otimes Z_c'Z_c$$

$$f_c = (\Sigma_c^{-1} \otimes Z_c')vec(Y_c - X_cB_c)$$

Again this density serves as a candidate distribution in a second independent Metropolis-Hastings step within the Gibbs sampler. The remainder of the sampler uses conditionals with well-known distributions. The conditional posterior of Σ_c is proportional to:

$$p(\Sigma_c|Y,\Theta_1 \setminus \Sigma_{c,u}) \propto |\Sigma_{c,u}|^{-\frac{T+\kappa+N+1}{2}} exp\{-\frac{1}{2}tr\left[(U_c^{'}U_c)+\bar{S}\right]\Sigma_{c,u}^{-1}\}$$

which is consistent with an inverse-Wishart distribution:

$$p(\Sigma_c|Y,\Theta_1 \setminus \Sigma_{c,u}) \propto iW((U_c^{'}U_c)+\bar{S}, T+\kappa) \tag{23}$$

In terms of the cross-country hyper-parameters, $\bar{\beta}$ has a conditional posterior proportional to a Normal:

$$p(\bar{\beta}|Y,\Theta_1 \setminus \bar{\beta}) \propto N(\left[\sum_c G_c\right]^{-1}\left[\sum_c g_c\right], \left[\sum_c G_c\right]^{-1}) \tag{24}$$

$$G_c = (\lambda_1 L_{1c})^{-1}$$

$$g_c = (\lambda_1 L_{1c})^{-1}\beta_c$$

The conditional posterior of $\bar{S}$ is proportional to:

$$p(\bar{S}|Y,\Theta_1 \setminus \bar{S}) \propto |\bar{S}|^{\frac{C\kappa-N-1}{2}} exp\{-\frac{1}{2}tr\bar{S}\left[\sum_c \Sigma_{c,u}^{-1}\right]\}$$

which corresponds to a Wishart distribution:

$$p(\bar{S}|Y,\Theta\setminus\bar{S}) \propto W\left(\left[\sum_c \Sigma_{c,u}^{-1}\right]^{-1}, C\kappa\right) \tag{25}$$

Note that $E(\bar{S}) = C\kappa(\sum_c \left[\Sigma_{c,u}^{-1}\right])^{-1}$. This implies that the expected value of $\bar{S}$ is the harmonic mean of the individual country covariance matrices scaled by the degrees of freedom parameter κ. This is used to determine the covariance of the cross-country model,$\bar{\Sigma}$. By setting $\bar{\Sigma} = \bar{S}/\kappa$, one obtains a matrix that is analogous to a covariance matrix and in expectation is equivalent to the harmonic mean of the estimated country covariances.

The conditional posterior for the shrinkage parameter, λ_1, is proportional to:

$$p(\lambda_1|Y,\Theta\setminus\lambda_1) \propto \lambda_1^{-\frac{CN^2L+v+2}{2}} exp\left\{-\frac{1}{2}\left(\frac{s}{\lambda}+\sum_c\left[(\beta_c-\bar{\beta})'\lambda_1^{-1}L_{1c}^{-1}(\beta_c-\bar{\beta})\right]\right)\right\}$$

or

$$p(\lambda_1|Y,\Theta\setminus\lambda_1) = iG_2\left(s+\sum_c\left[(\beta_c-\bar{\beta})'L_{1c}^{-1}(\beta_c-\bar{\beta})\right], CN^2L+v\right) \tag{26}$$

Where iG_2 refers to an inverted Gamma-2 distribution. For computational convenience, it is easier to draw from the posterior distribution of the inverse of λ_1 which is easily shown to be proportional to a standard Gamma distribution.

In terms of the identification parameters, the slope terms have the following conditional densities:

$$p(\Upsilon_c|Y,\Xi_c,\Theta\setminus\Upsilon_c) \propto N(K_c^{-1}k_c, K_c^{-1}) \tag{27}$$

where:

$$K_c = \sigma_{\omega c}^2 U_c'\Xi_c^{-1}U_c + \lambda_2^{-1}L_{2c}^{-1}$$

$$k_c = \sigma_{\omega wc}^{-2} U_c'\Xi_c^{-1}\mathcal{M}_c + \lambda_2^{-1}L_{2c}^{-1}\bar{\Upsilon}$$

And the conditional posterior of $\sigma_{\omega c}^2$ is proportional to:

$$p(\sigma_{\omega c}^2|Y,\Xi_c,\Theta\setminus\sigma_{\omega c}^2) \propto \sigma_{c\omega}^{-T+1} exp\{-\frac{1}{2}\left[(V_c^{'}\Xi_c^{-1}V_c)\right]\sigma_{\omega c}^{-2}\}$$

which is consistent with an inverse-Gamma distribution:

$$p(\sigma_{\omega c}^2|Y,\Xi_c,\Theta\setminus\sigma_{\omega c}^2) \propto iG((V_c^{'}\Xi_c^{-1}V_c), T) \tag{28}$$

The conditional posterior of ξ_{ct} can be expressed as:

$$p(\xi_{ct}|Y,\Theta) \propto \xi_{ct}^{-(\nu+3)/2} exp\left[-(\sigma_{\omega c}^{-2}(m_{ct} - \Upsilon^{'} u_{ct}) + \nu)/2\xi_{ct}\right]$$

Which is consistent with each diagonal element ξ_{ct} being related to the inverse of a χ^2, specifically:

$$p((\sigma_{\omega c}^{-2}(m_{ct} - \Upsilon^{'} u_{ct}) + \nu)/\xi_{ct}|Y,\Theta) \propto \chi^2(\nu+1)$$

In terms of the cross-country hyper-parameters, $\bar{\Upsilon}$ has a conditional posterior proportional to a Normal:

$$p(\bar{\beta}|Y,\Theta_1 \setminus \bar{\Upsilon}) \propto N(\left[\sum_c J_c\right]^{-1}\left[\sum_c j_c\right], \left[\sum_c J_c\right]^{-1}) \tag{29}$$

$$J_c = (\lambda_2 L_{2c})^{-1}$$

$$j_c = (\lambda_2 L_{2c})^{-1}\Upsilon_c$$

Last, the posterior of $\lambda_2$is proportional to:

$$p(\lambda_2|Y,\Theta \setminus \lambda_2) = iG_2\left(s + \sum_c \left[(\Upsilon_c - \bar{\Upsilon})'(\lambda_2 L_{2c})^{-1}(\Upsilon_c - \bar{\Upsilon})\right], CN + v\right)$$

Where iG_2 refers to an inverted Gamma-2 distribution.

Figure 2: Simulated ω compared with classical distributions

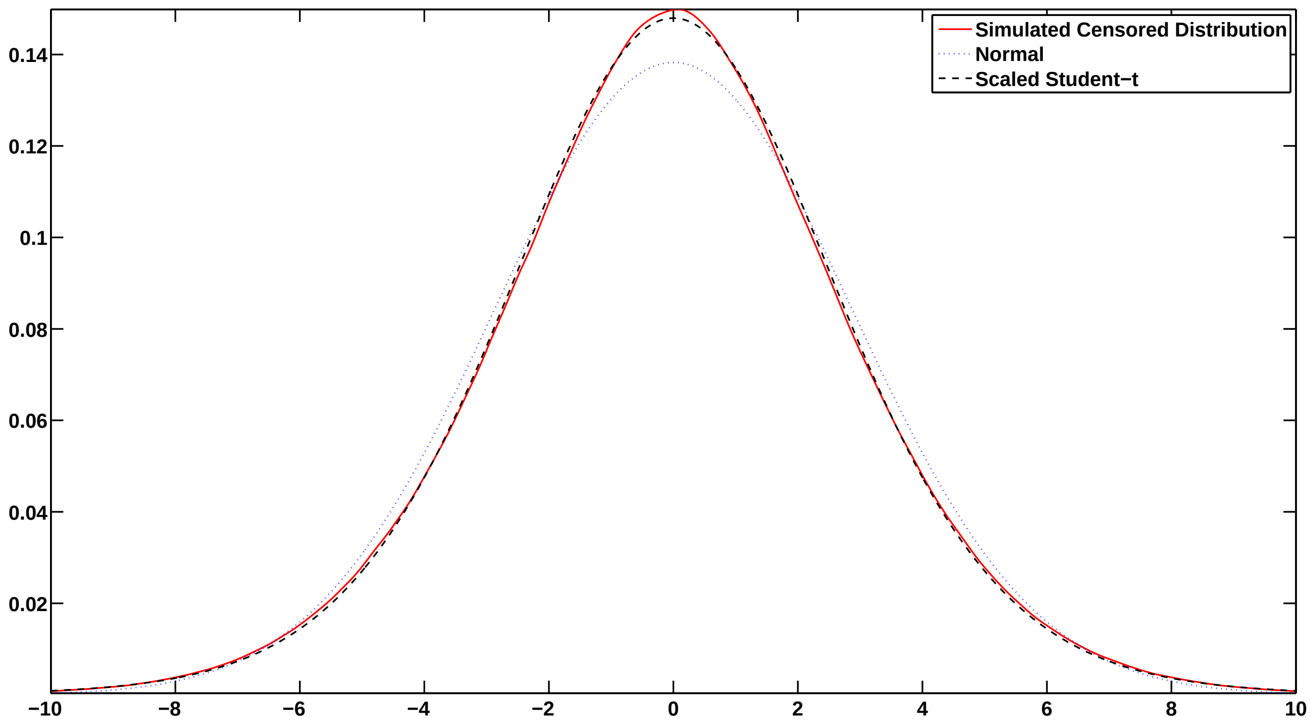

Notes: Censored distribution the kernel density of 1 million draws from the random censoring model with $p = 0.04$, $M = 30$ and $\psi = \sigma_v^2 = 1$. The normal distribution is selected to match the first two moments, the scaled t the first four moments.

Figure 3: Proxy variable and actual changes in the bond yield: Italy

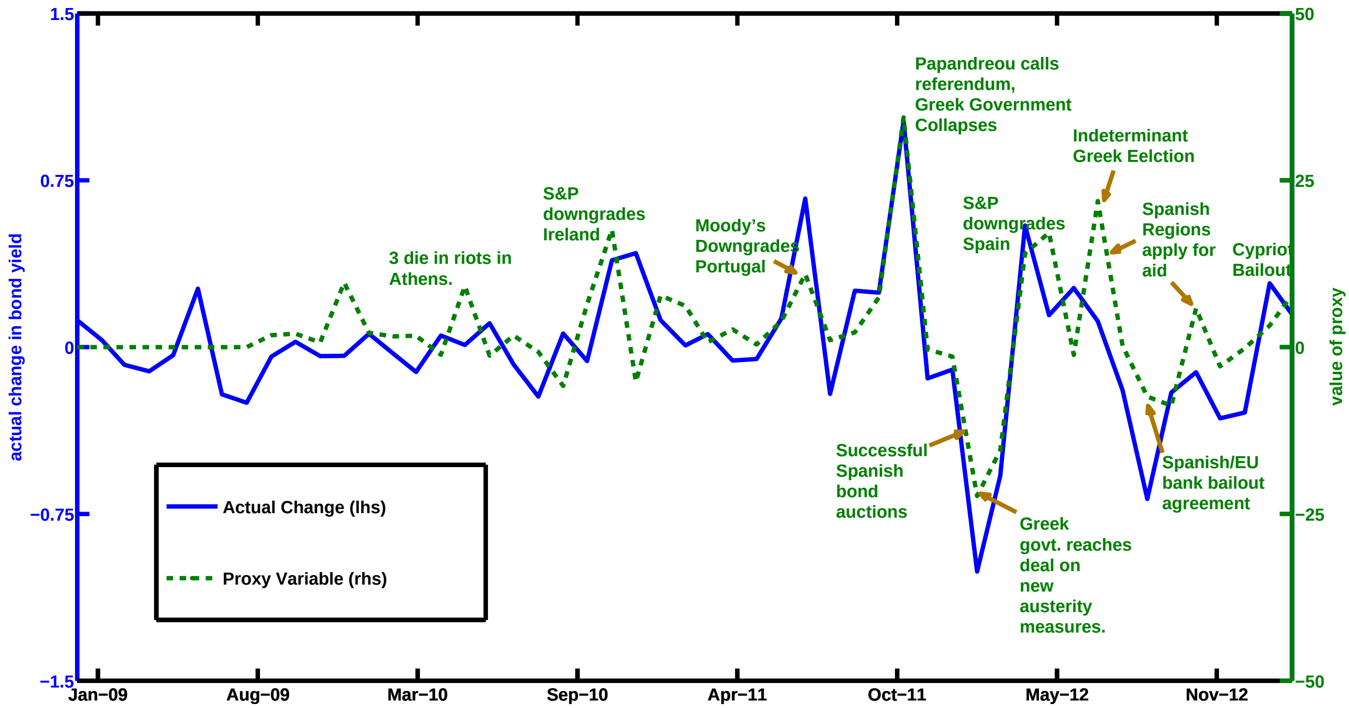

Notes: Y-axis denotes percentage moves in the month. Green dashed line is the proxy variable calculated as the summed value of changes Italian yields during a window around included foreign events in that month (right hand axis), blue undashed line is the the actual change in the 10-year bond yield over the course of the month (left hand axis). The graph is annotated with an illustrative set of important events.

Figure 4: Proxy variable and actual changes in the bond yield, other countries.

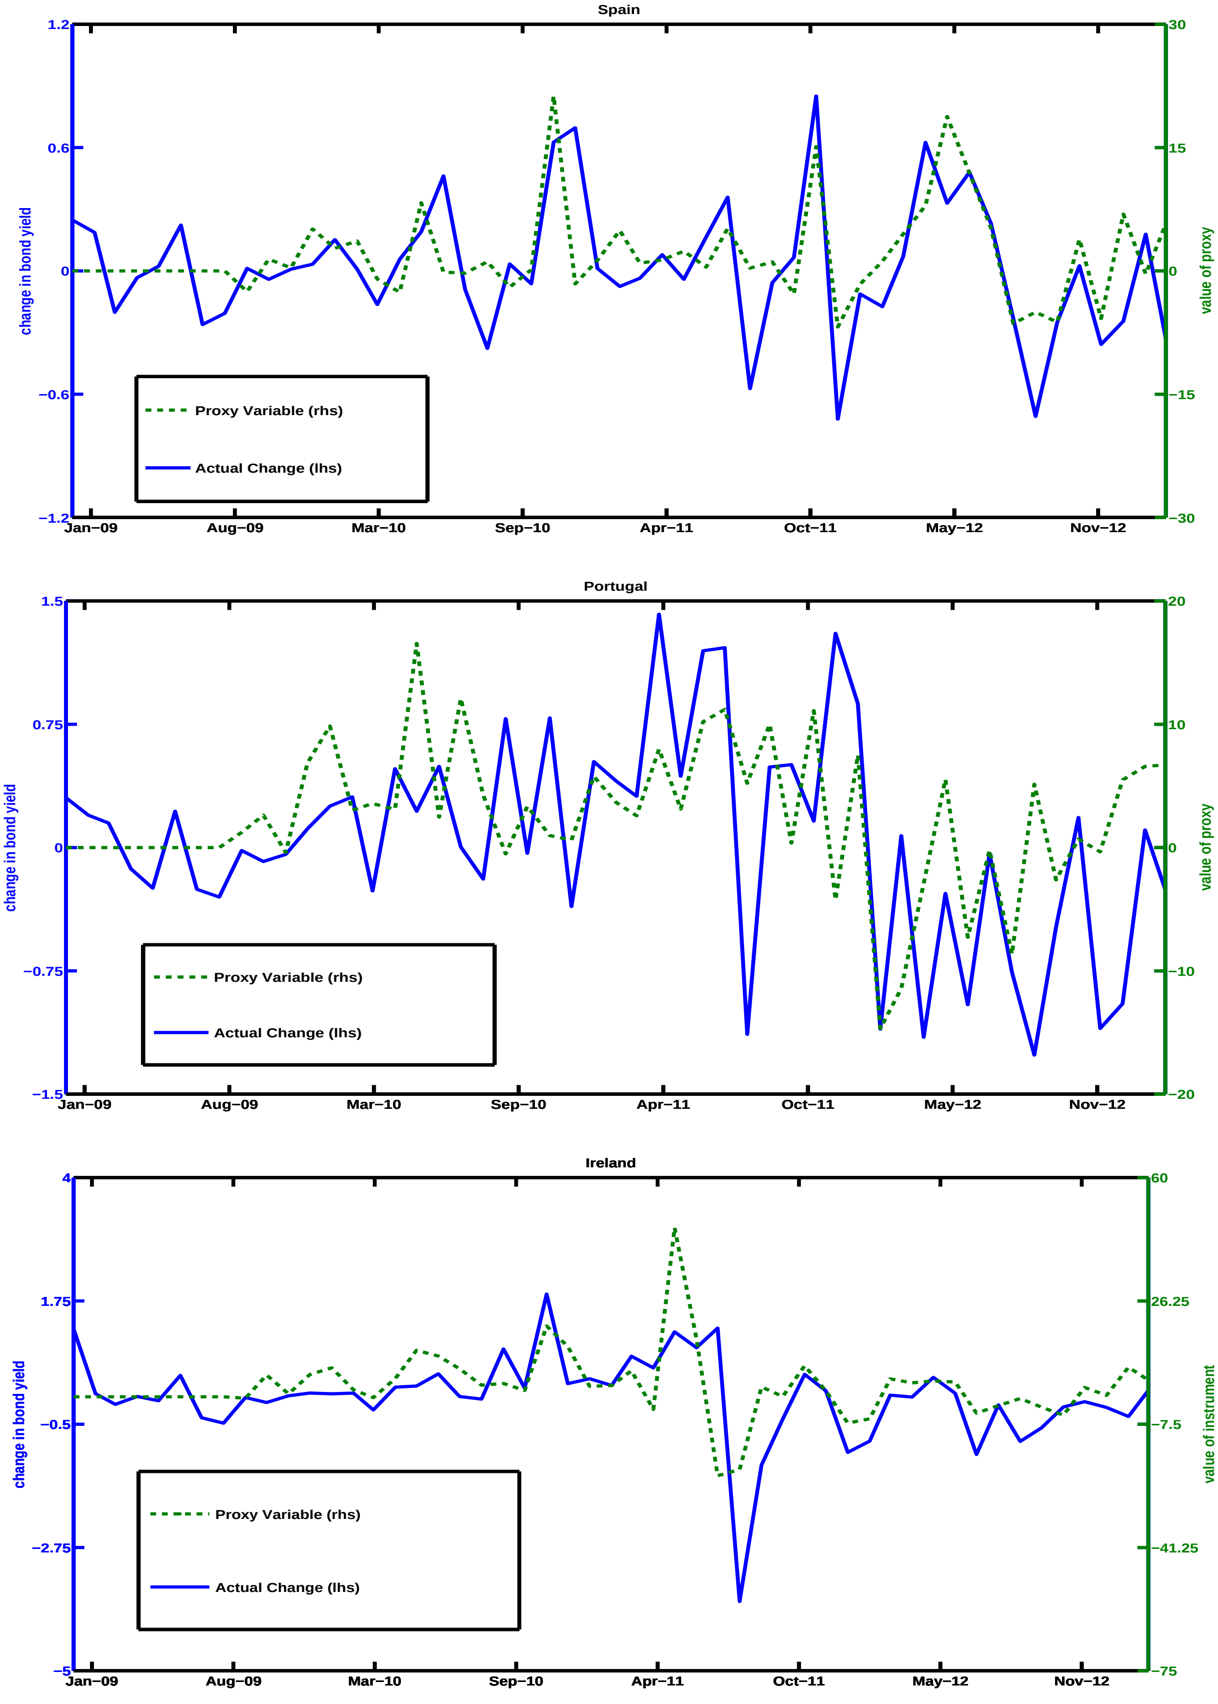

Notes: Y-axis denotes percentage moves in the month. Green line is the proxy variable calculated as the summed value of changes during events in that month, right hand axis, blue line is the the actual change in the 10-year bond yield, left hand axis.

Figure 5: Events ranked by their squared market reaction

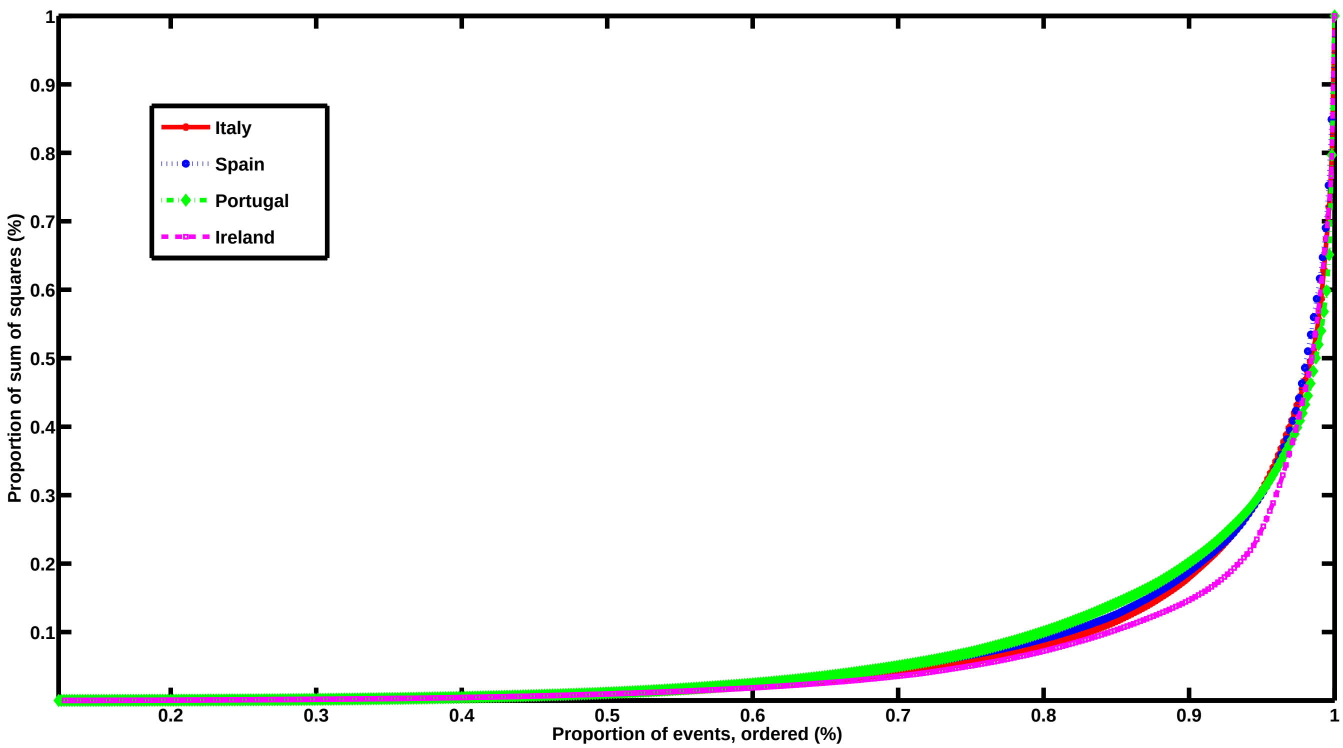

Notes: X-axis denotes the cumulative share of events in proxy ordered by size of market reaction. Y-axis denotes cumulative share total sum of squared market reactions. Period July 2009-March 2013. Irish proxy excludes events from May-October 2011 due to a break in intra-day data. Overlapping events, non-"headline" events outside the market open and domestic events are not included. Market moves refer to change in local 10 year bond yields in a 20 minute window about an event.

Figure 6: Mean Country Impulse Responses to a Systemic Shock

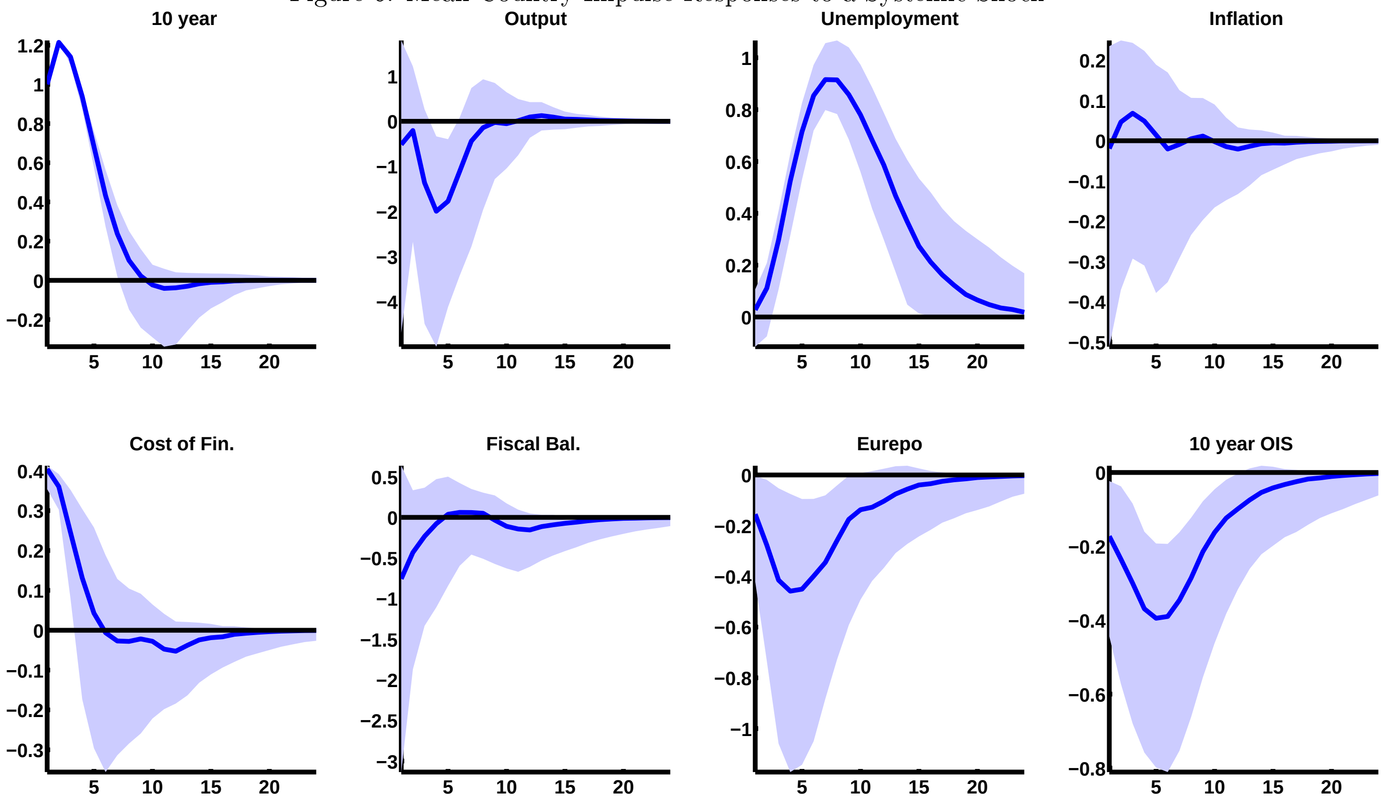

Notes: IRFs are scaled to be consistent with a 100bp increase in sovereign yield on impact and are computed over 24 months.Y-axis is percentage points in all cases. Mean country model refers to impulse responses estimated using $\bar{\beta}$, $\bar{\Upsilon}$ and $\bar{\Sigma}$. Centre line is the median of 10000 non-sequential draws from the simulated posterior. Error bands are 95% Bayesian credible intervals. 10 year refers to the 10 year soverign bond and output refers to industrial production.

Figure 7: Country Specific Impulse Responses to a Systemic Shock

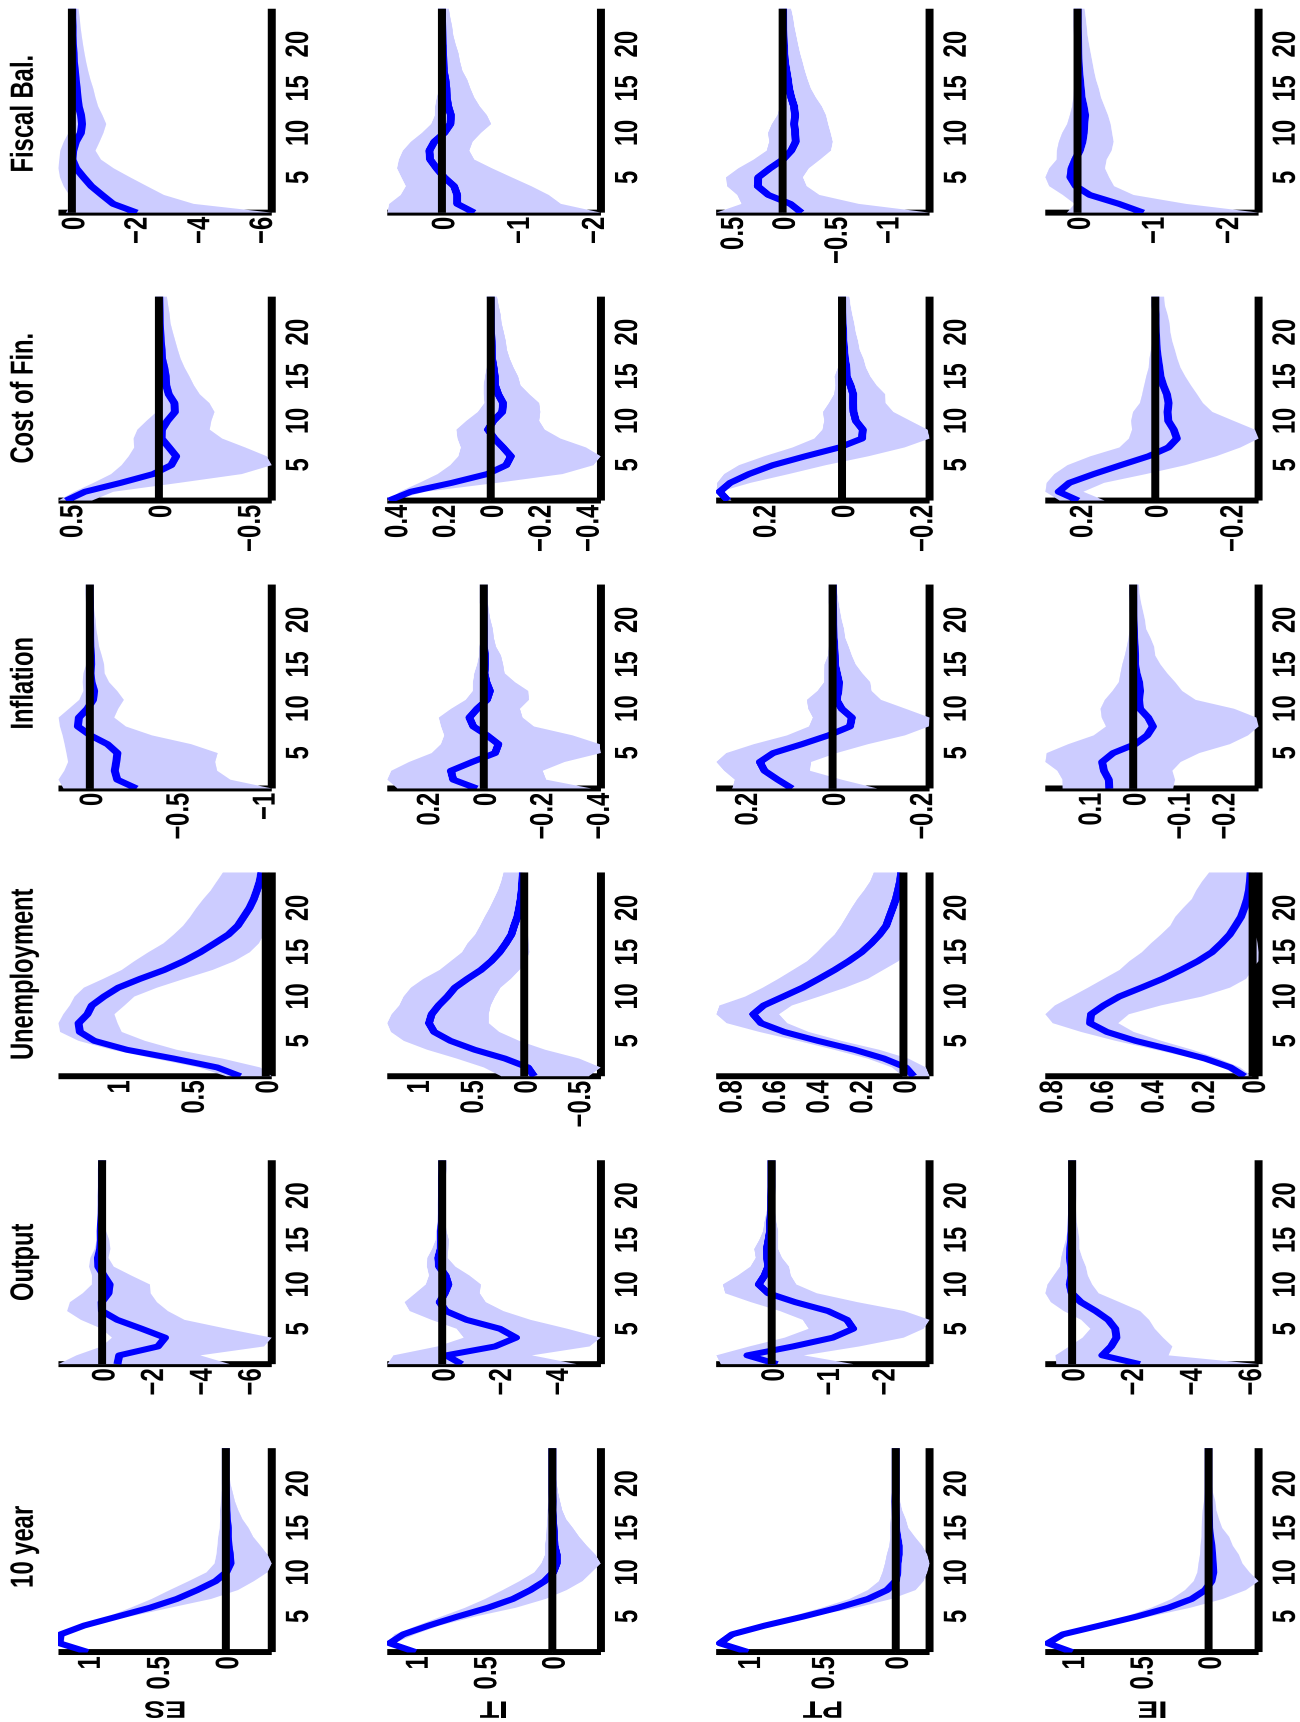

Notes: IRFs are scaled to be consistent with a 100bp increase in sovereign yield on impact and are computed over 24 months.Y-axis is percentage points in all cases. Centre line is the median of 10000 non-sequential draws from the simulated posterior. Error bands are 95% Bayesian credible intervals. Due to similarity with mean country models responses of the Eurepo and OIS rates are not presented. 10 year refers to the 10 year soverign bond and output refers to industrial production.

Figure 8: Forecast Error Variance Decomposition (Contribution of Systemic Shock)

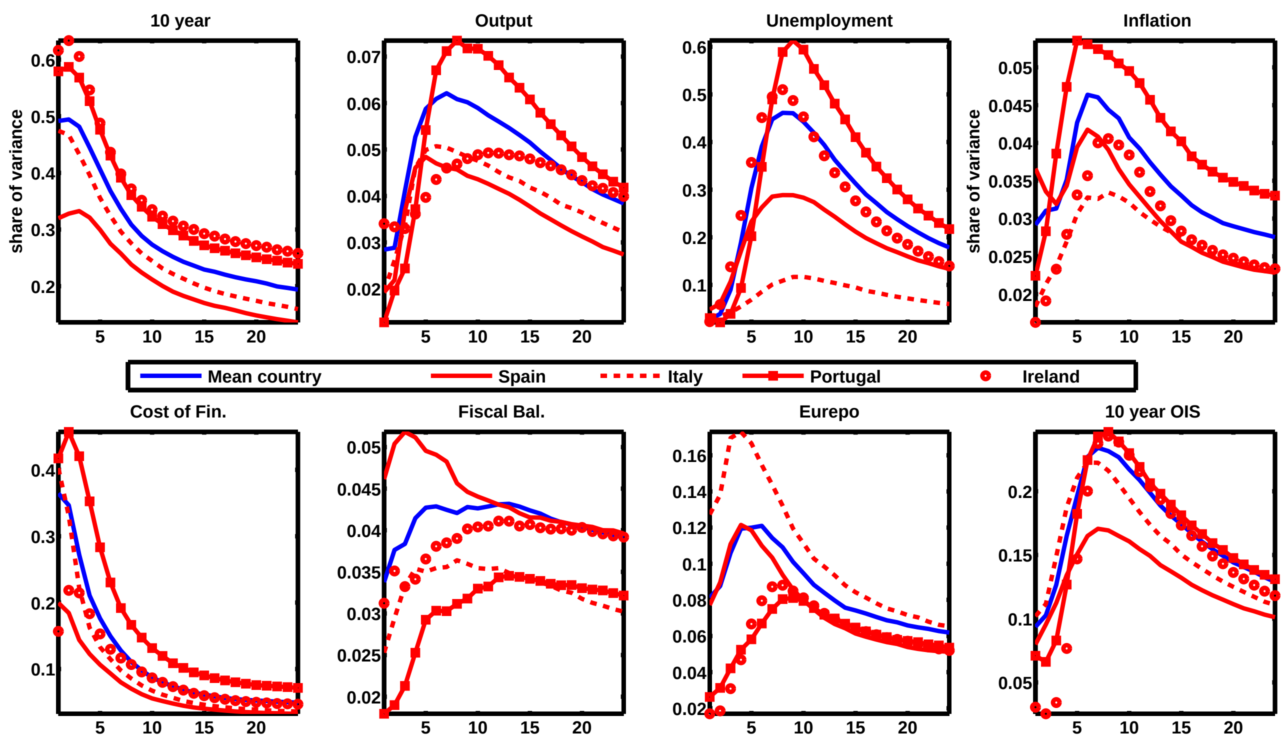

Notes: Forecast error variance decompositions to systemic shock. Computed over a 24 month horizon (x-axis). Line is the median of 10000 non-sequential draws from the simulated posterior. Mean country model (in blue) refers to decomposition estimated using $\bar{\beta}$, $\bar{\Upsilon}$ and $\bar{\Sigma}$. 10 year refers to the 10 year soverign bond and output refers to industrial production. 10 year refers to the 10 year soverign bond and output refers to industrial production.

Figure 9: Counterfactual Analysis

Actual versus counterfactual sovereign bond yields

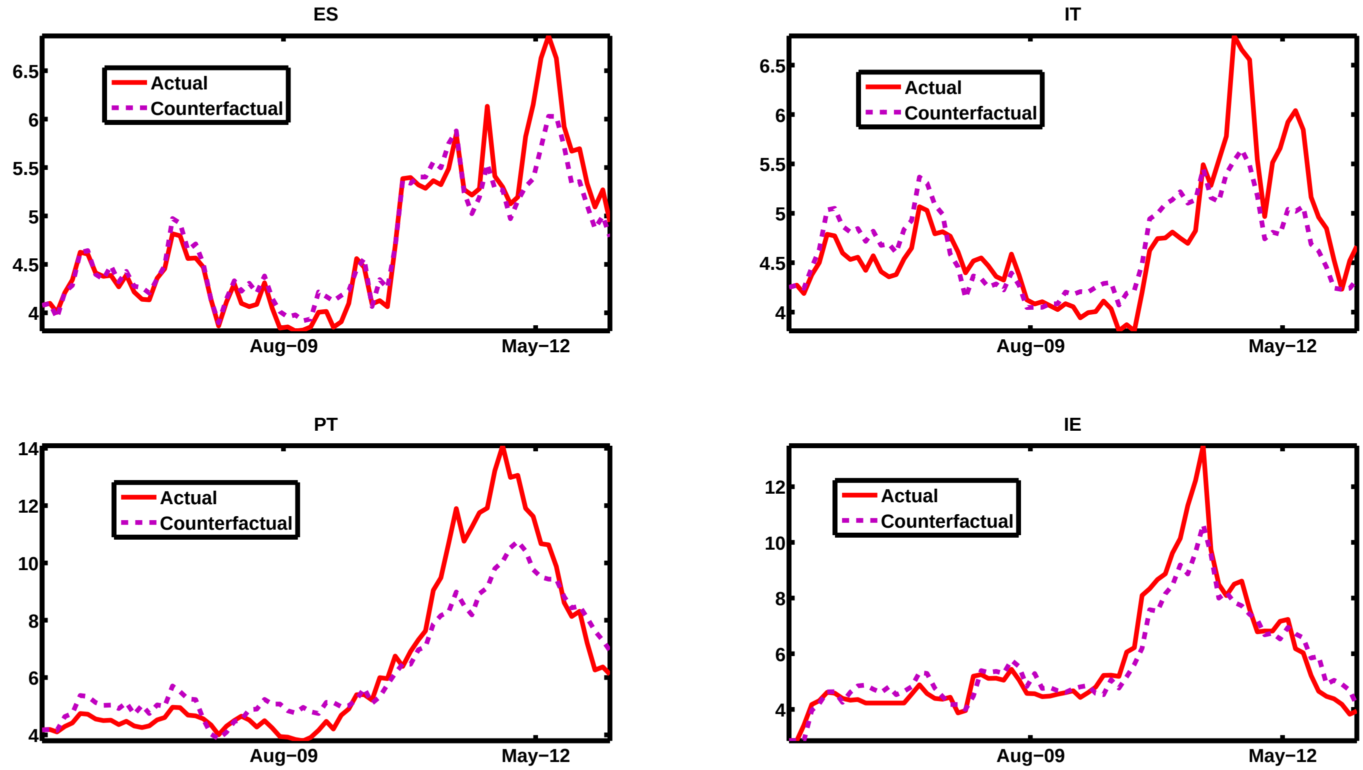

Implied premia

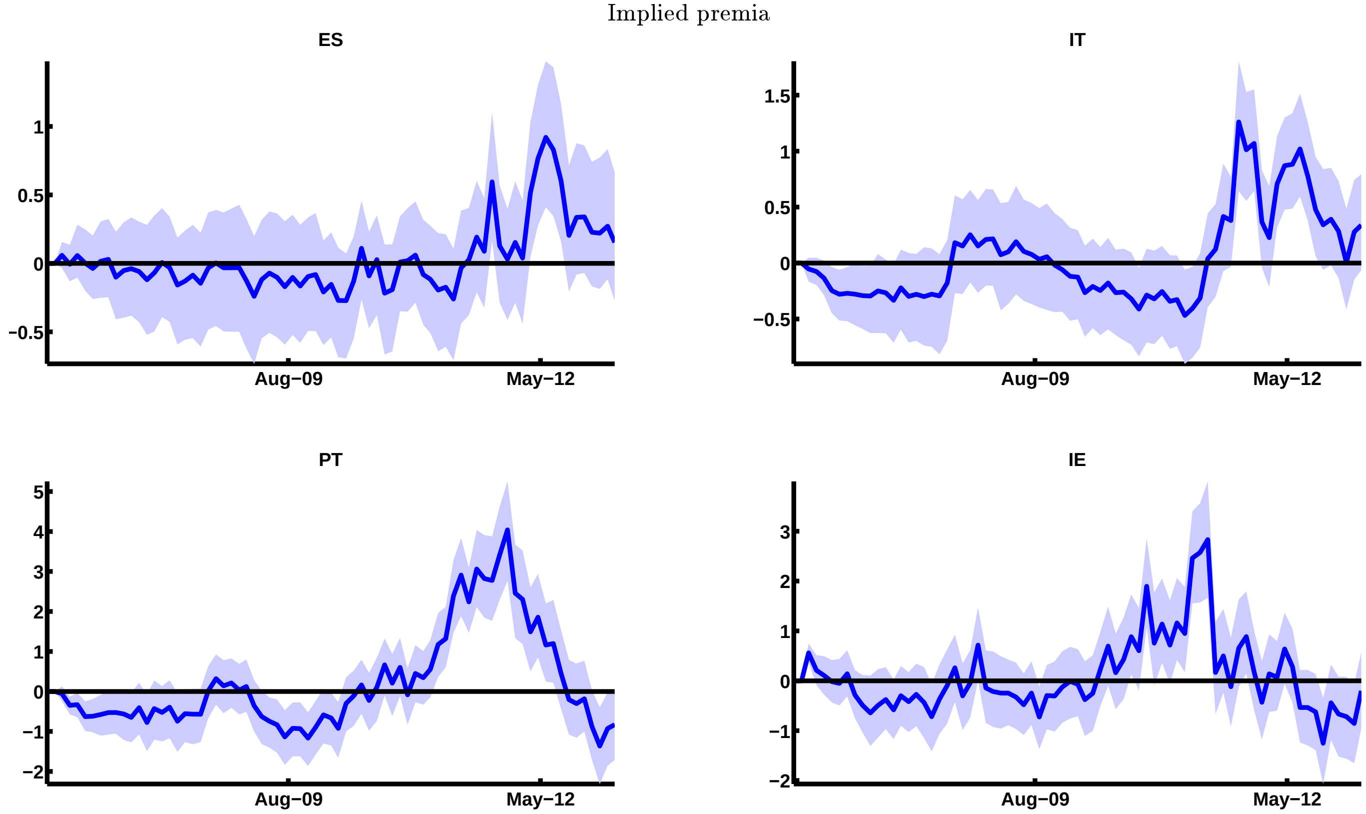

Notes: Counterfactuals are constructed by zeroing the systemic shocks and recreating the yield. Y-axis is percentage points. Centre line is the median of 10000 non-sequential draws from the simulated posterior. Risk premia (lower pane) is equivalent to Actual−Counterfactuals. Error bands are 95% confidence intervals.

Figure 10: Mean country impulse response under alternative proxy definitions

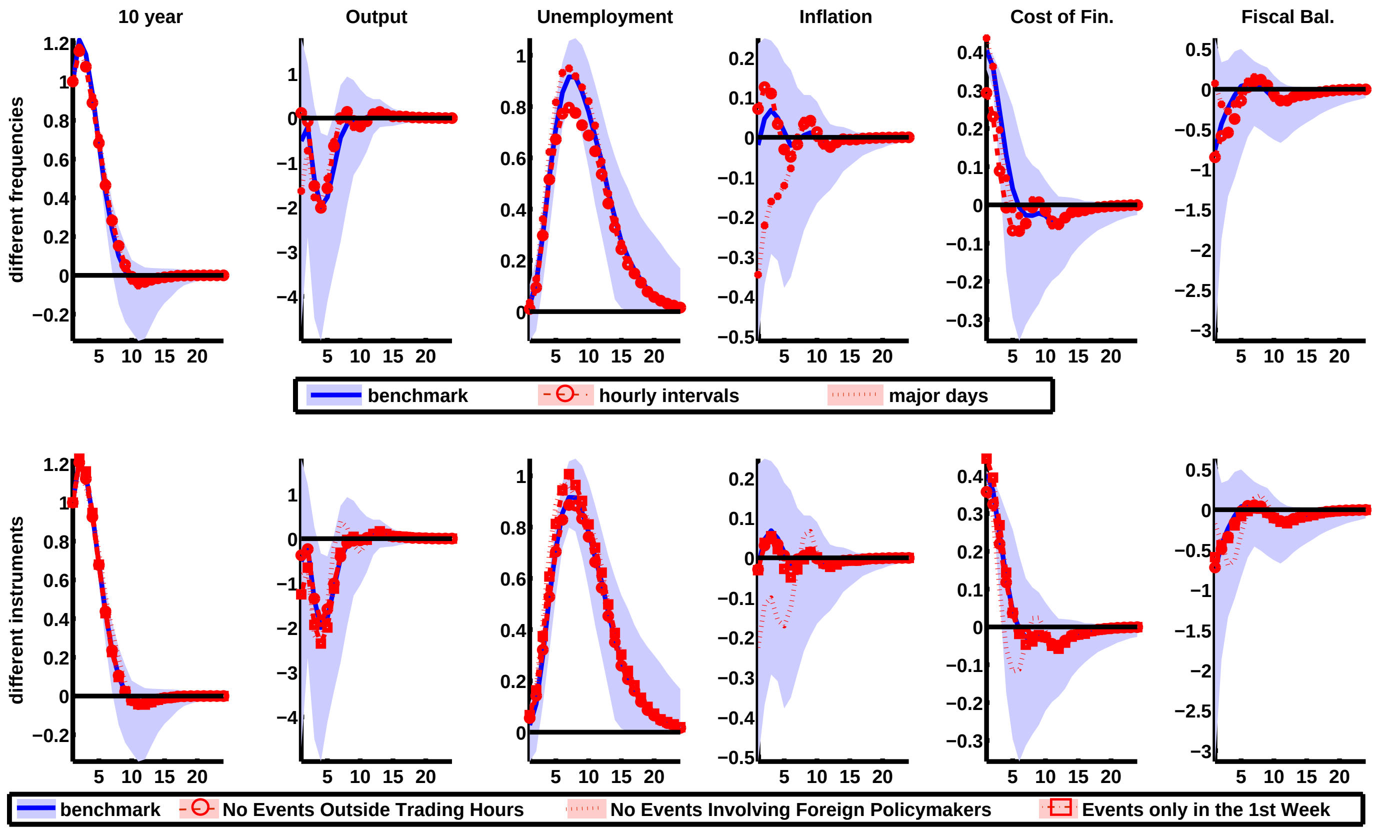

Notes: Y-axis is in percentage points, X-axis is months. Impulse responses scaled to be consistent with a 100bps increase in the bond yield.

Figure 11: Mean Country Impulse Responses using Decomposed Private Sector Financing Sources

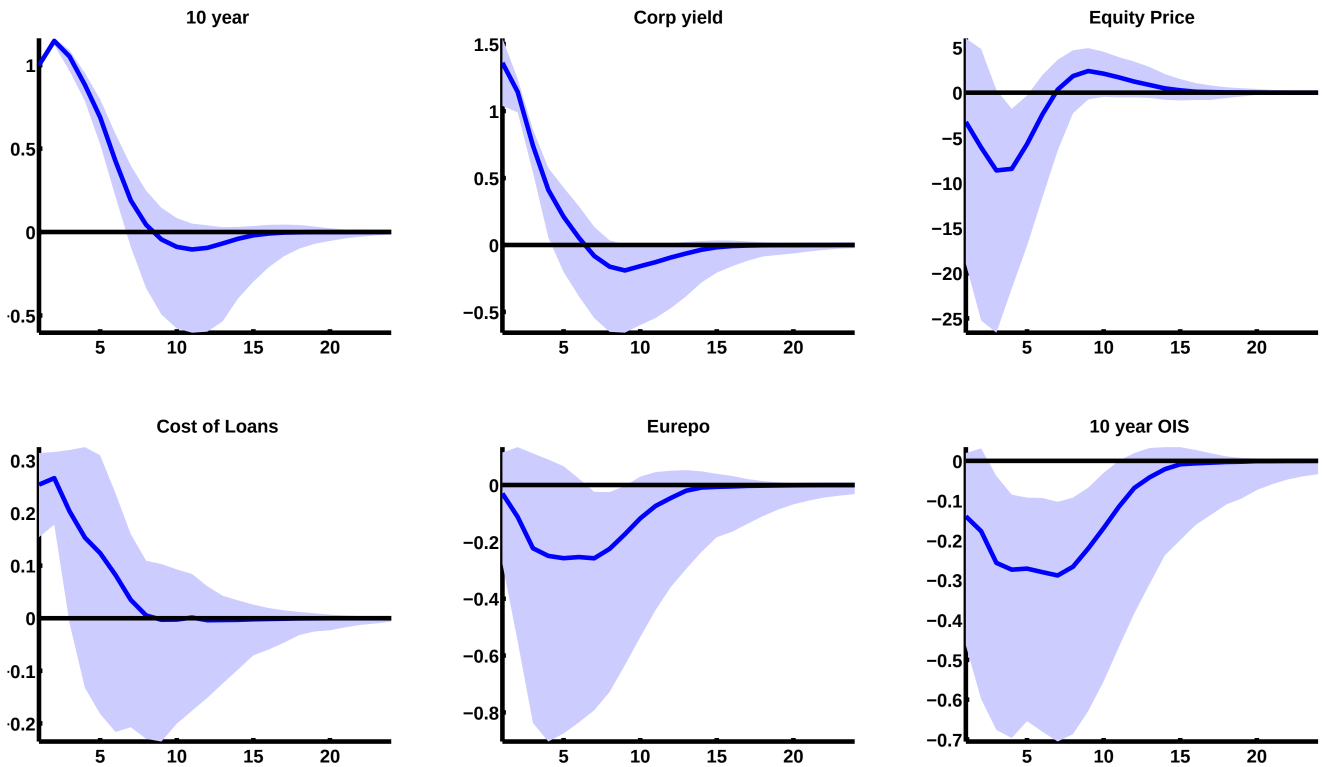

Notes: Y-axis is in percentage points, X-axis is months. Impulse responses scaled to be consistent with a 100bps increase in the bond yield. Other controls included bu not shown as described in the main text but not shown.

Figure 12: Placebo Study

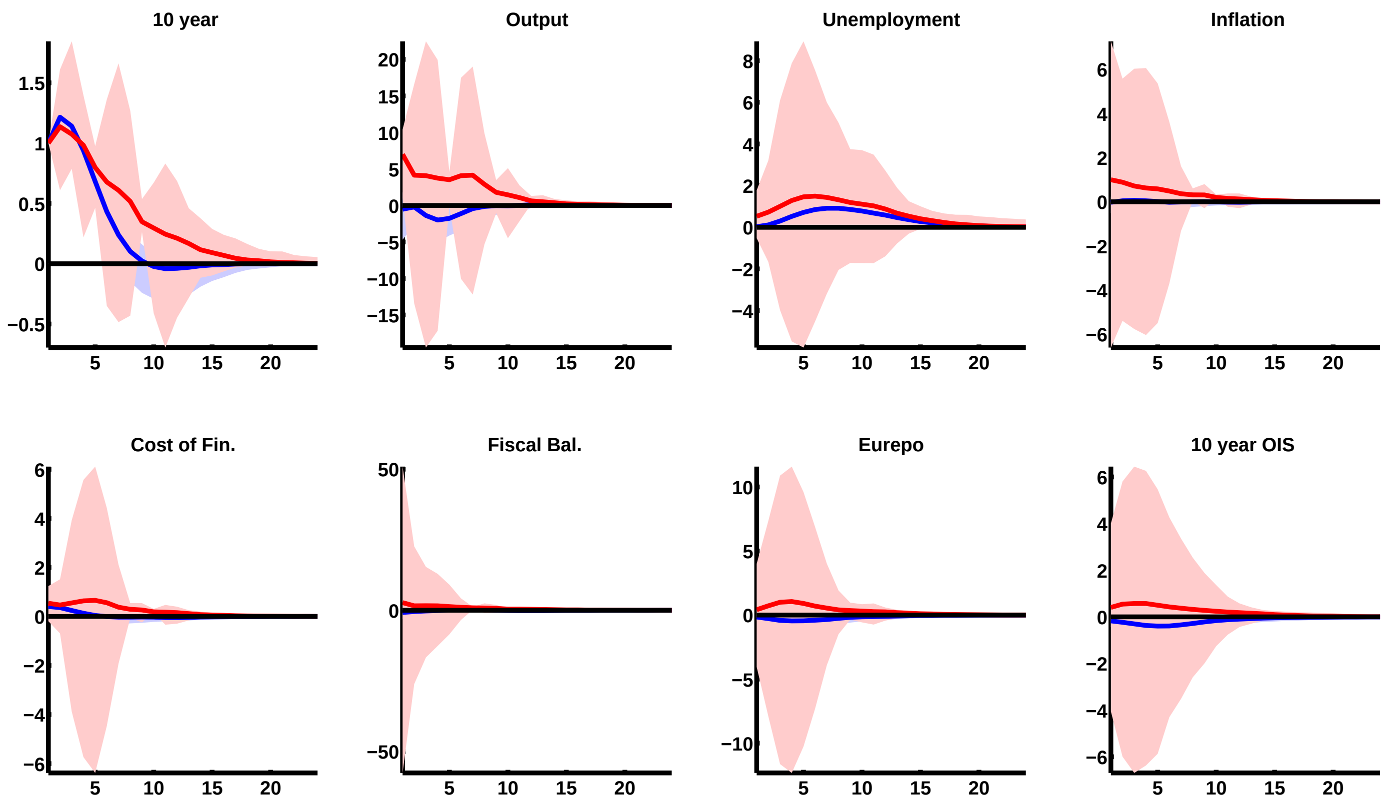

Notes: Y-axis is in percentage points, X-axis is months. Impulse responses scaled to be consistent with a 100bps increase in the bond yield. Placebo proxy constructed using same events timed to the previous trading day. Red line and shaded error refers to placebo study with 95% credible intervals. Blue line and shaded areas refer to the benchmark case.