Systemic Sovereign Risk:

Macroeconomic Implications in the Euro Area

Saleem A. Bahaj*

January 24, 2014

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

Abstract

What are the macroeconomic implications of changes in sovereign risk premia? In this paper, I use a novel identification strategy coupled with a new dataset for the Euro Area to answer this question. I show that exogenous innovations in sovereign risk premia were an important driver of the economic dynamics of crisis-hit countries, explaining 30-50% of the forecast error of unemployment. I also shed light on the mechanisms through which this occurs. Fluctuations in sovereign risk premia explain 20-40% of the variance of private borrowing costs. Increases in sovereign risk result in substantial capital flight, external adjustment and import compression. In contrast, governments appear not to increase their primary balances in response to increases in sovereign risk. Identifying these causal effects involves isolating a source of fluctuations in sovereign borrowing costs exogenous to the economy in question. I address this problem by relying upon the transmission of country-specific events during the crisis in Europe to the sovereign risk premia in the remainder of the union. I construct a new dataset of critical events in foreign crisis-hit countries and I measure the impact of these events on yields in the economy of interest at an intraday frequency. An aggregation of foreign events serves as a proxy variable for structural innovations to the yield to identify shocks in a proxy SVAR. I extend this methodology into a Bayesian setting to allow for flexible panel assumptions. A counterfactual analysis is used to remove the impact of foreign events from the bond yields of crisis hit countries: I find that 40-60% of the trough-to-peak moves in bond yields in crisis-hit countries are explained by foreign events, thereby suggesting that the crisis was not purely a function of weak local economic conditions.

Key words: High frequency identification, Narrative identification, Contagion, Bayesian VARs, Proxy SVARs, Panel VARs.

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

JEL Classification Numbers: E44, E65, F42

1 Introduction

What are the macroeconomic implications of high and variable sovereign risk premia? To what extent can fear of default in the sovereign debt market drive an economic downturn and destabilize the economy? While these questions have dominated recent policy debates, there is limited empirical evidence on the effects that fluctuations in sovereign risk may exert on macroeconomic dynamics. The challenge is simultaneity. On the one hand, a rise in premia may reflect fundamental economic weaknesses that foreshadow worsening primary balances and a rise in public debt. On the other hand, high interest rates on public debt may in turn exacerbate fiscal distortions and credit conditions, curbing output and thus undermining the budget. The main contribution of this paper is to separate the latter from the former and address this endogeneity problem.

There is a substantial literature describing the economic conditions and dynamics associated with sovereign debt crises and sovereign default.[1] However, the observation of a crisis does not determine its causality. Nor is how one should define a crisis, beyond default, clear. This paper focuses on sovereign borrowing costs. I utilise a novel combination of high frequency and narrative identification strategies to isolate an exogenous source of variation in sovereign risk premia. Specifically, I rely on two well-known observations about the recent financial crisis in Europe.[2] First, the prices of Euro Area sovereign bonds reacted strongly to specific events - be they policy announcements, speeches, riots, elections etc. Second, events that were specific to individual countries transmitted to sovereign borrowing costs across the rest of the union. The maintained assumption is that this transmission of foreign events reflects movements in sovereign bond yields that are plausibly orthogonal to innovations to the local economy.

To isolate bond market movements due to foreign events, I first build a new narrative dataset of the crisis period. I use news summaries to isolate key, country-specific events in economies suffering from elevated sovereign borrowing costs during the crisis period; specifcally, I consider events in Ireland, Italy, Portugal, Spain, Cyprus and Greece. I determine, using a news wire, the time at which an event occurred. I then measure the impact on sovereign yields in other crisis-hit countries by looking at the response of the relevant sovereign bond market in an immediate time window spanning an announcement.

Using this dataset, I show that, for countries that have experienced elevated sovereign risk premia, up to 50% of the forecast error variance of the sovereign bond yield was unrelated to the local economy. The macroeconomic consequences were sizeable. Innovations to sovereign risk were a critical driver of recent unemployment dynamics; they explain close to 40% of the variation of the unemployment rate and a substantial proportion of the observed increase in unemployment in crisis hit countries since 2010. In terms of relative magnitudes: on average a 100bps increase in the sovereign yield corresponds to a 2 percentage point reduction in industrial production growth and adds 0.9 percentage points to the unemployment rate (both are peak responses).

This paper's key methodological feature is integrating high frequency bond market reactions into the identification stage of a dynamic, macroeconomic time series model. I use the monthly aggregation of high frequency market reactions to foreign events as a proxy variable for a latent structural innovation to the yield. This follows the proxy SVAR approach of Mertens and Ravn (2013a,b). The proxy SVAR is critical in this

---

[1]See, for example, Reinhart and Rogoff (2009), De Paoli et al. (2009), Borensztein and Panizza (2009), Levy Yeyati and Panizza (2011), Furceri and Zdzienicka (2012) and Mendoza and Yue (2012).

[2]To highlight a few examples from this literature documenting these facts: Gade et al. (2013) show that statements by European politicians (both nationally and at a EU level) have had meaningful impact on sovereign borrowing costs. Brutti and Saure (2013) show that during the early stages of the crisis in Greece (2009-2011), critical events related to that country passed through to CDS spreads in the remainder of the EU. Afonso et al. (2012), De Santis (2012) and Arezki et al. (2011) highlight the importance and transmission of ratings decisions. Attinasi et al. (2009) and Acharya et al. (2011) conduct similar analysis for bank bailout decisions. In the broadest case, Beetsma et al. (2013) show that "news" in general (as isolated from a news summary) as opposed to specific events also move markets throughout the union.

context as it is designed to deal with the mismeasurement that is endemic to my identification strategy. Events are observed irregularly, intraday data is noisy and market reactions may be prone to over/undershooting. One cannot assume that the high frequency bond market reactions are a perfect measure of the informational content of an event. For example, they may be attenuated due to pre-announcement rumours. To account for these features of the data, I explicitly model the proxy as an aggregation of infrequently observed, poorly measured signals over the true structural innovation.

I extend the methodology of Mertens and Ravn into a panel setting. The crisis provides only a short time series of observations. In order to improve the precision of the estimates I use the information contained in the cross-section of crisis countries. I focus my attention only on Ireland, Italy, Portugal and Spain.[3] However, the crisis has struck countries in different ways with varying intensity - a homogenous parameter setup is an overly strong assumption. This motivates a Bayesian approach. I select priors that utilise cross-sectional shrinkage to exploit the information from multiple countries without imposing cross-country homogeneity (following Jarocinski (2010)).[4] This methodology allows for an estimate of heterogeneous country-specific models that make use of the information in the cross-section, as well as an average model around which the heterogeneous parameters are centred. How close the parameters are to this cross-country average, i.e. the degree of shrinkage, is allowed to be data dependent. The panel setup works in both stages of the estimation procedure. Information from the cross-section is used to inform the parameters in the reduced form VAR as well as aid the identification of shocks.

The adaptation of the proxy SVAR into a Bayesian setup represents an additional methodological contribution of this paper. I rewrite the model in a fashion that allows me to obtain the density of the proxy conditional on the reduced form model. This results in a hierarchical, joint posterior density which is straightforward to simulate numerically. This procedure uses all the information contained in the proxy and macroeconomic time series data when constructing estimates of the model parameters. Credible intervals reflect estimation uncertainty about both the reduced form and identification stages of the model.

To set the scope of the analysis, I do not attempt to disentangle the sources of the transmission between countries: the focus is purely on the macroeconomic implications. Potential channels have been described elsewhere in the literature. To highlight a few: countries in a currency union are vulnerable to belief-driven crises (De Grauwe (2011)); foreign events could transmit by coordinating investors onto bad equilibria. Alternatively, negative events in one country could alter the parameters of policy in the rest of the union either by committing limited bailout resources or altering political support for interventions.[5] Transmission could flow through common creditors (Arellano and Bai (2013)). Last, the sustainability of the currency union also came under question; alterations in market confidence in the political commitment to the union may result in a convertibility premium during times of stress.

These channels are not mutually exclusive, nor do I wish to pretend that the above is exhaustive. However, a strength of the analysis is that it is robust to the channels listed above and does not require making statements over their relative importance. Furthermore, these sources of transmission represent reasons why one can observe a change in the riskiness of the sovereign separate from economic conditions. They are also illustrative of the fragilities in the Euro Area that enable the identification strategy of relying on external events to isolate exogenous movements in yields. However, this implies that movements in yields are caused by an intensification of the crisis in Europe as a whole. This may have different implications from a purely

---

[3]Cyprus and Greece are omitted from the main analysis due to a lack of intraday data. Cyprus' small size and the fact that Greece experienced default suggest they may not be appropriate to include in the sample regardless. One can attempt to extend the analysis to non-crisis countries (Germany, France etc.); however, a proxy constructed with foreign events has only a weak relationship with yields in those countries and the analysis fails to produce meaningful results.

[4]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.

[5]Corsetti et al. (2006) show how the resources available to an international lender of last resort alters equilibria in a self-fulfilling crisis. Tirole (2012) offers a discussion of cross country insurance mechanisms in the context of the crisis.

domestic-driven crisis.

As a consequence, I use the terminology employed in Ang and Longstaff (2013) and refer to identified shocks as systemic sovereign risk shocks. Ang and Longstaff (2013) define this concept from an asset pricing framework as the country-specific vulnerability of sovereigns to common adverse events, as opposed to idiosyncratic risk related to sovereign specific factors. The idea being that a sovereign could default due to a purely local issue such as an isolated political crisis; alternatively, sovereign risk could be due to a shock that triggers a chain of defaults, which in this context, would reflect the Euro Area's fragility as a system. The external nature of the identification implies that I am identifying an increase in sovereign risk which affects several countries simultaneously and thus the identified shock refers to this systemic component of the risk premia.[6]

The critical identifying assumption underpinning the proxy is that local bond yield movements due to foreign events are independent of the economic situation of the local country. One may have several concerns over whether this is true. First, events in specific countries could be a reaction to prior macroeconomic shocks in other countries or to common shocks. My identification strategy is strengthened by its reliance upon intraday financial market reactions to events. Market participants are able to anticipate systematic reactions between events abroad and economic conditions at home. Therefore, the market reaction is the unanticipated component of an announcement, orthogonal to shocks that the market is already aware of (see Gürkaynak and Wright (2013)).

Second, the observed high frequency market move could be a function of multiple shocks to the local economy that are occurring simultaneously. Care is taken to only consider events where no other market-relevant news happened in the same time window. I exclude data releases from the analysis as they may be capturing innovations in a common business cycle component. Pan-European interventions are also omitted as they contain a role for the ECB and would be correlated with monetary shocks. A further concern may be that the reaction to the foreign event occurs because market participants are learning about local economic conditions. However, once one excludes events that may be informative about common shocks (data releases, ECB announcements etc.), there is little reason to think that foreign agents have additional private information about the local economy.[7]

Third, the market reaction could also be explained by real or financial inter-linkages between the event country and the local economy. This cannot be ruled out and may weaken the identifying assumptions. However, the main results can largely be replicated using the narrower set of events related only to Greece, a country representing a minimal share of Euro Area GDP, imports and external liabilities.

In a preliminary analysis, I provide evidence to support the identifying assumptions. By running predictive regressions at alternative time frequencies I show that there is no evidence that the market reactions of a country's bond yield to foreign events can be explained either by past market reactions to events, be they local or foreign, or by past macroeconomic data.[8] Aggregations of the market reaction to foreign events are

---

[6]The term systemic risk is loosely defined. It is both used to represent exposure common or propagated shocks; as well as vulnerabilities in the system unrelated to economic fundamentals such as a sudden liquidity crises or shifts in investors beliefs. For the crisis in Europe both interpretations fit, there is a sense that a portion of the crisis was self-fulfilling while at the same time the institutions governing the Euro were a shared fundamental. The term also fits well in this context given the external nature of the identification. Certain countries were clearly systemically important during the crisis; Greece, in particular, came perilously close to leaving the currency union and rendering the Euro no longer irreversible. Even the tiny country of Cyprus potentially took on systemic importance when losses were threatened on insured depositors.

[7]One may argue that foreign policymakers may reveal additional information over the state of policy on a pan-European level. For example, the extent of political support for interventions from international policymakers. However, this is exactly the sort of innovation to the yield we wish to identify.

[8]The fact that the narrative used here is based upon market reactions means one should be less concerned about anticipation affects. Any prior information should be priced. We should of think of the proxy series used here as capturing an unanticipated shock. Furthermore, it is a shock to a forward looking variable in the shape of the bond yield, which could capture private sector expectations about future economic conditions. This means the concerns about foresight in VAR identification as raised by, for example, Leeper et al. (2013) are less relevant. This is not to say that the true state of the world has an invertible VAR representation during the crisis. However, as Sims (2012) discusses SVARs, can still perform well even when with working data

also contemporaneously uncorrelated with market reactions to data releases, local events, or ECB decisions (capturing monetary shocks). This supports the assumption that the moves in local bond yields due to foreign events are not caused by a simultaneous change in local economic conditions.

Beyond the results described above, I also shed light on the channels by which increases in the sovereign borrowing costs feed through into the real economy. There is no direct evidence that the increase in yields provokes governments to improve their primary fiscal balance and reduce their rate of borrowing. Indeed, in some 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 response of the fiscal balance is at zero at the same point that the unemployment rate peaks. Therefore, if one defines fiscal tightening as a change in the cyclically-adjusted balance, then the relative co-movement of the fiscal balance and the unemployment rate is still evidence of subsequent contractionary fiscal policy.

The theoretical literature that embeds sovereign risk into a general equilibrium model of the macroeconomy highlights two further mechanisms. Corsetti et al. (2013) argue that a channel via which increases in the sovereign borrowing costs feed through into real economy is by causing a deterioration in private sector financial conditions, thereby acting as a negative financial shock. This effect is readily apparent in the results of this paper: a 100bp increase in the sovereign bond yield leads to a more than one to one increase in corporate yields, an increase in private loan rates and a fall in equity prices. I show that 20-40% of the forecast error variance in a composite measure of private borrowing costs was due to fluctuations in sovereign risk. Mendoza and Yue (2012) highlight the external implications of sovereign debt crises, arguing that capital flight leads to a loss of access to imported inputs that cannot be substituted for domestically, thereby harming productivity. My results also support this mechanism. I find that systemic shocks lead to capital flight; a 100bps increase in the bond yield is consistent with private capital outflows worth 3.7% of GDP. This leads to an external adjustment which results in an improvement in the trade balance of about 0.6% of GDP, which appears to result from import compression.

A second result of interest pertains to a counterfactual analysis: what would the sovereign bond yields have been if the identified systemic shocks had not occurred and borrowing costs were solely determined by local conditions? There has been some debate in the literature about the extent to which sovereign borrowing costs amongst Euro Area countries are explained by local macroeconomic variables. For example, De Grauwe and Ji (2013) and Aizenman et al. (2013) argue that movements in yield could not be explained by macroeconomic determinants of sovereign risk during the crisis period. While using an estimated DSGE embodying sovereign default risk, Bi and Traum (2013) suggest borrowing costs in Greece were consistent with macroeconomic fundamentals. The counterfactual analysis lies somewhere in between these two extremes. I find that 40-60% of the trough to peak move in borrowing costs in crisis-hit countries was as a result of systemic shocks. This effect peaked in July 2011: the model estimates that systemic shocks added 127bp to the 10 year Italian bonds with an equivalent figure of 381bp for Portugal. The systemic component spiked again in May 2012, around the time of the Greek election,the interpretation being 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 of factors 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 impact of systemic shocks had abated and yields appear to be at a neutral setting in line with local macroeconomic conditions.

The remainder of the paper is organised as follows; following a brief review of the relevant literature, section 2 lays out the concept behind the identification strategy. Section 3 gives context to this identification strategy by detailing specific occasions on which foreign country-specific events move local yields in other crisis hit countries, and also delineates the data collection procedure. Section 4 describes the Bayesian estimation

---

generated from non-invertible models; especially when augmented with forward looking market prices.

procedure. Section 5 presents the results of the analysis, section 6 discusses robustness and section 7 provides the conclusion.

1.1 Related Literature

This paper is closely related to the strand of literature which documents the impact of debt crises and default.[9] There is evidence that crises are associated with periods of economic weakness but opinion regarding the extent and dynamic impact still lacks broad consensus. Analysing data for default events, De Paoli et al. (2009) suggest debt crises correspond to output losses of around 5% of GDP lasting as long as 10 years. In an annual panel growth model Furceri and Zdzienicka (2012) show that crisis episodes correspond to output losses of around 10% after 8 years. In contrast, Borensztein and Panizza (2009) find output losses of 0.5-2.0% that do not last more than a year and Tomz and Wright (2007) find only a weak relationship between output and defaults. By studying the evolution of output on a quarterly basis Levy Yeyati and Panizza (2011) find that the default quarter is typically the trough in activity and recovery follows immediately afterwards. This suggests that the decline in activity appears to be due to the anticipation of, rather than being a product of, default. An acknowledged weakness in this literature is that debt crises are endogenous.[10] Another issue is that these analyses rely upon actually observing default. Taking the Levy Yeyati and Panizza (2011) results at face value suggests that it is the increase in the sovereign risk premia that does damage to the economy. Countries that experience elevated borrowing costs but honour their obligations are not considered.[11]

A paper more close to the analysis conducted here, and which also focuses on the crisis in Europe, is Neri and Ropele (2013). They use a factor model to isolate a common sovereign risk component during the crisis and use it as an endogenous variable in a FAVAR. The distinction with this work is that their identification strategy relies upon a Cholesky ordering, with the sovereign risk factor ordered after monetary policy but ahead of macroeconomic variables. This represents a strong identifying assumption as it requires that monetary policy did not respond to sovereign risk tensions contemporaneously and also that sovereign risk premia are not a contemporaneous reaction to macroeconomic or financial shocks. Uribe and Yue (2006) investigate the relationship between business cycles and borrowing costs in emerging markets; they also rely on a Cholesky ordering such that real variables do not respond contemporanously to innovations in country risk premia (i.e. the opposite of Neri and Ropele (2013)); although they add theoretical evidence to back their assumptions.

From a methodological standpoint, this paper is closely related to that strand of the empirical macroeconomic literature which uses narrative methods to identify macroeconomic shocks; that is to say, by carefully reading policy documents or by an analysis of historical events one can identify innovations in economic series that are used as exogenous regressors separate from the estimated disturbances in a time series model. This literature follows the seminal contribution of Romer and Romer (1989).[12] The approach taken by Mertens and Ravn (2013a,b) and Stock and Watson (2012), and which is also used here, is distinct in the sense that the outside information contained in the narrative is not treated as a separate variable in the model but is instead used as a "*proxy*" or "*external instrument*" for the structural shock of interest.[13]

---

[9]Reinhart and Rogoff (2009) offer a historical review of the aftermaths of financial crises in general.

[10]There have been attempts to improve identification: Borensztein and Panizza (2009) use a two step procedure where they estimate the probability of default separately based on a selection of contemporaneous determinants and include it as a regressor in the growth equation. Furceri and Zdzienicka (2012) try to obtain identification by restricting attention to debt crises which occur during periods of good economic performance. However, neither of these approaches provide a true source of exogenous variation in sovereign risk.

[11]Pescatori and Sy (2007) make this point and find the defining debt crises by high borrowing costs rather than default alters empirical findings.

[12]Some notable examples include Romer and Romer (2004) for monetary policy shocks, Ramey and Shapiro (1998) and Ramey (2011b) for government spending shocks and Burnside et al. (2004), Romer and Romer (2010), Cloyne (2013) and Favero and Giavazzi (2012) for tax shocks.

[13]I favour the use of the word proxy to describe the narrative time series used here as, unlike Stock and Watson (2012), the

A second strand of methodological literature relates to high frequency identification. High frequency market reactions to events are a rich source of exogenous variation, because common information about other macroeconomic shocks should already be embedded in market prices around the time of an announcement. High frequency market reactions to federal reserve decisions have been used to identify the impact of monetary policy shocks; for example, Cochrane and Piazzesi (2002) use the daily change in bond prices on FOMC meeting days as their preferred measure. The more recent works of Nakamura and Steinsson (2013) and Gorodnichenko and Weber (2013) drill down to an intraday frequency and assess market reactions to FOMC announcements in much tighter windows in order to identify monetary shocks. Faust et al. (2004) discuss how one can identify VARs using high frequency reactions in the futures market.Gürkaynak and Wright (2013) offer a detailed review of this literature.

To my knowledge, this is the first paper to look at intra-day market reactions with respect to the Euro-crisis, and to furthermore use those reactions in the narrative identification framework described. However, there is a substantial literature looking at the determinants of yields during the crisis. As described, Ang and Longstaff (2013) use an asset pricing model to seperate Eurozone country borrowing costs into those due to local idiosyncratic factors and those due to a vulnerability to common shocks, which they denote the systemic component. Aizenman et al. (2013), De Grauwe and Ji (2013), Giordano et al. (2013) and Manasse and Zavalloni (2013) all investigate how sovereign borrowing costs depend on macroeconomic conditions, either across countries or over time.

2 Proxy SVAR Identification

To fix ideas, I begin by briefly describing the identification methodology. I rely upon the proxy SVAR approach of Mertens and Ravn (2013a,b) and Stock and Watson (2012) whereby identification of the contemporaneous relationships between macroeconomic variables is obtained via an external proxy that is assumed to be correlated only with the structural shock of interest. Consider the following dynamic macroeconomic model:

$$\alpha(L)Y_t = \mathcal{B}\epsilon_t$$

Where Y_t is a $N \times 1$ vector of observed macroeconomic variables (including sovereign borrowing costs) and ϵ_t is a set of unidentified structural shocks satisfying: $\mathbb{E}(\epsilon_t) = 0$, $\mathbb{E}(\epsilon_t \epsilon_t^{'}) = I_N$ and $\mathbb{E}(\epsilon_t \epsilon_s^{'}) = 0 \, \forall s \neq t$. Following the standard SVAR assumptions, the reduced form residuals from a regression of Y_t on its lags are a linear combination of the structural shocks: $u_t = \mathcal{B}\epsilon_t$; where u_t are the reduced form residuals and $\mathcal{B}$ is a non-singular $N \times N$ matrix. The identification problem emerges because there is not enough infomation in the covariance matrix of the reduced form residuals, $\mathbb{E}(u_t u_t^{'}) = \Sigma_u$, to identify $\mathcal{B}$. Solutions to this problem in the literature often impose restrictions on the matrix $\mathcal{B}$ (long or short) or propose restrictions *a priori* on the impact certain shocks have on the included variables. In this context, short-run restrictions are not viable as we are interested in the response of variables derived from market prices which react almost instaneously to new information. Long-run or sign restrictions are potentially also questionable.

The solution taken here is not to make any assumptions about the contemporaneous relationships between macroeconomic series, but instead use an external variable to estimate those relationships. The vector of structural shocks can be partitioned into a shock of interest and other structural shocks $\epsilon_t = (\varepsilon_t, \tilde{\varepsilon}_t^{'})^{'}$. In this case the scalar structural shock of interest, ε_t, is the systemic sovereign risk shock - i.e. the structural shock to the sovereign's cost of borrowing unrelated to local economic conditions. I abuse terminology somewhat and describe the remaining structural shocks, $\tilde{\varepsilon}_t$, as other local economic shocks but the vector also contains

---

identification stage of the model in this paper is not an instrumental variable based regression.

shocks stemming from external sources - such as monetary shocks from the ECB.

I assume there exists a proxy variable m_t that is correlated with the structural shock of interest, $E(m_t\varepsilon_t) = \phi$, and uncorrelated with the other structural shocks, $\mathbb{E}(m_t\tilde{\varepsilon}_t) = 0$, where ϕ is a scalar. These are the critical identifying assumptions and are analogous to the exogeneity and relevance conditions for a valid instrument. Following Mertens and Ravn (2013a,b), the proxy could be considered a scaled version of the true shock measured with some error, for example $m_t = \phi\varepsilon_t + \omega_t$, where ω_t is measurement error uncorrelated with the shock ($\omega_t \sim IID(0, \sigma_\omega^2)$) and ϕ is a scalar coefficient.

From the non-singularity of $\mathcal{B}$, there exists $\mathcal{A}^{-1} = \mathcal{B}$ with the corresponding relationship $\mathcal{A}u_t = \epsilon_t$; where $\mathcal{A}$ is an $N \times N$ identification matrix scaled such that the structural shocks have unit variance. The identification matrix can be partitioned such that $\mathcal{A} = [a_1', a_2']'$ with $\varepsilon_t = a_1 u_t$. Given u_t, the vector a_1 is sufficient to identify the structural shock ε_t.[14] In contrast to Mertens and Ravn (2013b), the variation in the approach here is that I use this relationship to estimate the first row of the $\mathcal{A}$ matrix, a_1, as opposed to the first column of the $\mathcal{B}$ matrix; specifically:

$$m_t = \phi a_1 u_t + \omega_t \tag{1}$$

The unit variance restriction on the structural shock implies the quadratic form $a_1 \Sigma_u a_1' = 1$, which provides the additional restriction to identify ϕ.[15] Working 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)) is of technical importance as it means one has single equation regression between the proxy and the reduced form residuals about which one can make panel assumptions. Furthermore, once one has specified the distribution of ω_t, the conditional density of the proxy $p(m_t|u_t; \phi a_1)$ is known. This property is useful for embedding the identification strategy in a Bayesian framework. I return to the issue of implementing this identification strategy in a Bayesian VAR, including the distributional assumptions over ω_t in section 4. However, the challenge is constructing a proxy, m_t, which satisfies the identifying assumptions; i.e. a proxy which is correlated with changes in sovereign risk premia but not other local macroeconomic shocks. This is dealt with in the next section.

3 The proxy variable

3.1 Evidence from specific events

3.1.1 Catalonia Requests a Bailout, 28th of August 2012

At 13:01 London time 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 Spain's Regional Liquidity Fund. The report was confirmed officially by a Catalan government spokesman 5 minutes later. The move had been signalled by the Catalan authorities but the announcement still provoked a reaction in financial markets. At 13:00 London time Spanish 10 benchmark 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.[16]

---

[14]The matrix $\mathcal{B}$ can also be partitioned such that that $\mathcal{B} = [b_1, b_2]$ and $u_t = b_1\varepsilon_t + b_2\tilde{\varepsilon}_t$. 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 $b_1 = \Sigma_u a_1'$. Note, a_1 and b_1 are respectively $1 \times N$ and $N \times 1$ vectors.

[15]While assuming unit variances is standard practice in this setup it is not completely innocuous. First, it means that a_1 is only identified up to a signing and scaling convention. A one standard deviation stock has no interpretation so one needs to make a scaling assumption when computing impulse responses; variance decompositions and counter facturals are unaffected however. Second, it is impossible to make any statement about the relative variances of shocks across countries.

[16]See http://www.eurointelligence.com/briefings/2012-08-29 for a general description of events surrounding the bailout. A live summary of events during the day can be found here: http://www.telegraph.co.uk/finance/debt-crisis-live/9502734/Debt-

Figure 1: Catalonia Requests a Bailout, 28th August 2012: High Frequency Market Reaction
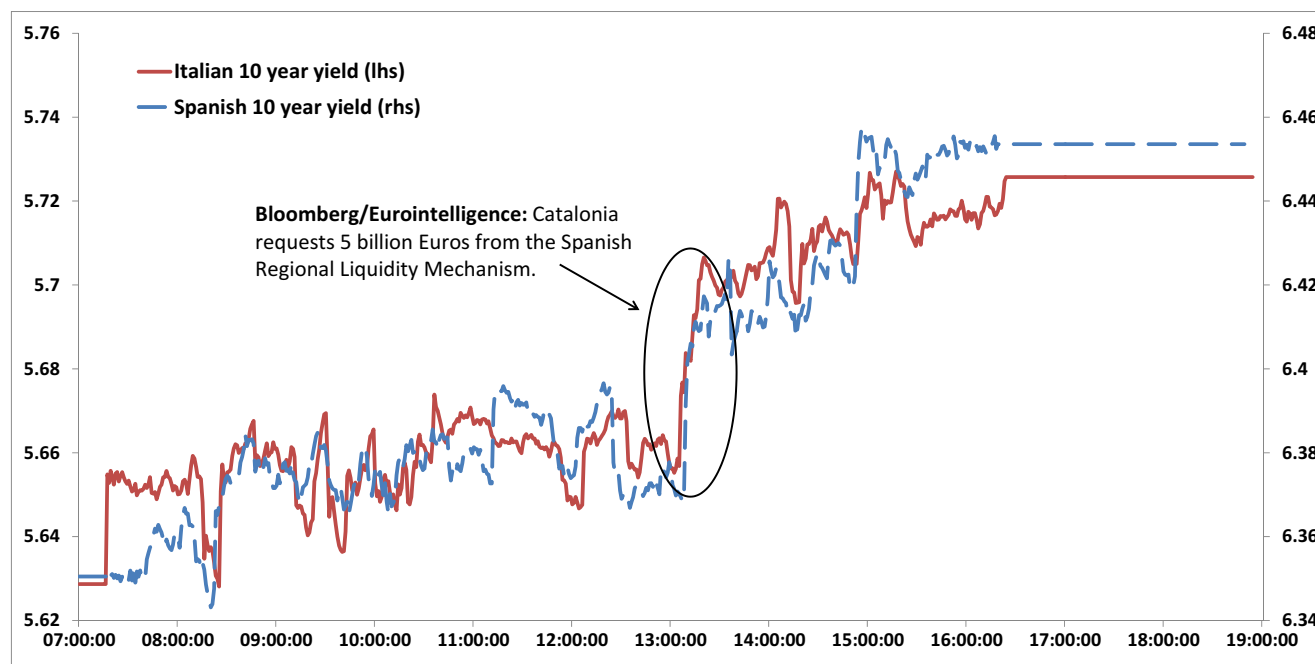
Notes: Intraday bond market moves in Spain and Italy around Catalonia's decision to request a bailout from the Spanish regional liquidity mechanism (28th of August 2012). The series are tick data aggregated into 1 minute averages (90% trimmed) for the mid-yield to maturity on 10 year benchmark bond. Italian yield is the red solid line and refers to the left hand y-axis. Spanish yield is the blue dashed line and refers to the right hand y-axis. The x-axis refers to London time.

The bailout decision was largely a 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 confined only to Spain. Italian 10-year bond yields increased by 5 basis points immediately subsequent to the announcement and yields in "core" countries declined. For example, German yields fell by 1.5bp in the few minutes following the announcement.[17] 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.

3.1.2 The Greek Parliament Approves Austerity, 12th of February 2012

By early 2012 it was clear to both policymakers and market participants alike that Greece would need a second bailout. An agreement in principle was reached between Greece and the Troika of official creditors. In exchange for official sector financing, Greece was required to embark on additional austerity measures, including 150,000 public sector job losses and 3.3 billion euros of spending cuts. The package was approved by the Greek Cabinet on the 10th of February 2012 but this led to the resignation of 6 Cabinet members. Late in the evening on Sunday the 12th the package would be voted on in the Greek parliament. The leadership of the two main Greek political parties backed the agreement. However, political support was not unanimous. A small nationalist party, Laos, withdrew its support from the governing coalition and the days leading up to the vote were marked by social unrest. Reports from the time suggest the package was deeply unpopular

crisis-as-it-happened-August-28-2012.html

[17] 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. Furthermore, Ireland was on its way to recovery at this point and and had recently started reissuing long term debt.

with the electorate.[18]

A no vote could have resulted in disorderly default. Greece had a 14.5 billion euro bond payment scheduled on March 20th; a payment that, in the absence of a bailout, it appeared that the country would be unable to meet. Official creditors refused to sanction a bailout unless austerity measures were approved. Furthermore, Greece's prime minister, Lucas Papademos, warned that a failure to pass the bill may result in the country's exit from the currency union.

Ultimately the legislation passed comfortably, by 199 votes in favour to 74 against, out of 300 lawmakers. The news was not all positive: Athens suffered from rioting with several buildings burned; 40 government MPs also rebelled against the measure leading to their being ousted from their parties and cutting the ruling coalition's majority. Nonetheless, financial markets reacted positively to the vote. Sovereign bond yields fell sharply in Greece when the market opened on the 13th and the response spread throughout the union. By 8:30am London time yields on benchmark 10 year sovereign bonds had fallen by 19.9bp in Portugal, 10.0bp in Italy and 4.0bp in Spain when compared to the close on the 10th (there was, however, little reaction in Ireland).

3.2   Discussion of the identification strategy

For clarity, the remainder of this paper adheres to the following terminology. The "event country" refers to the country where an event takes place. The "local country" refers to the country for which the bond market reaction is being recorded; the local country is the country of interest from the perspective of the analysis and usually not the event country. Tautologically, a "foreign country" is one external to the local country; so a "foreign event" is something that happens abroad from the perspective of the local country. In the Catalan example, Spain is the event country. If Italy was the country of interest, Catalonia's bailout is a foreign event which provoked a local bond market reaction of 5bp. If Spain was the country of interest, then the Catalonia decision is a "local event"; the event and local country are the same.

These two examples illustrate the transmission of foreign events to the borrowing costs of other Euro Area nations. More importantly, one can claim that there is little direct causality between the local bond market moves and other preceding local economic shocks. For example, it is difficult to think that the move in the Italian yield, in response to Catalonia's decision, as being a caused by a change in fundamental macroeconomic conditions in Italy in August 2012. It was not a reaction to a weakening in Italian total factor productivity, or a change in the Italian government's fiscal stance. Indeed, it is unlikely that the Catalans were thinking about Italy at all when they made their announcement. Similarly, the Greek MPs that voted for austerity should be concerned only with economic conditions in Greece and not those in Portugal, regardless of the impact their decision had on Portugal's borrowing costs. However, even if there was endogenous feedback between policy choices in one country from macroeconomic conditions in others, and such reactions may be reasonable in light of how the crisis developed, any systematic reaction should be anticipated by market participants, and should be priced. This also applies to common macroeconomic shocks. Therefore, in the Catalan example, we can think of the move in the Italian yield observed above as not being caused by changes in Italian economic conditions; rather, it may cause changes in those conditions, but was not caused *by* them.

This identification strategy breaks down if agents in the event country are reacting to changes in macroeconomic conditions in the local country that the market is not yet aware of, such 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 rather that markets are rational and there is no

---

[18] For an account of events and political mood surrounding the Greek vote I refer readers to Hewitt (2013) pages 238-246. For more journalistic accounts written in immediate aftermath of the vote see: http://www.eurointelligence.com/briefings/2012-02-13/ & http://www.telegraph.co.uk/finance/financialcrisis/9078221/Greece-passes-crucial-bailout-vote-as-country-burns.html.

information asymmetry between market participants and foreign agents regarding these shocks at the point of an event.

The information asymmetry problem demonstrates why a focus on foreign events is necessary. If, for example, we considered how local bond yields respond to a local announcement of an austerity package, then any market reaction would also be a response to the fiscal shock that the austerity package captures. Thus, local policymakers can reveal information about other local economic shocks when they make their policy announcements. For the same reason, events at a pan-European level cannot be considered. Pan-European policy interventions normally have some involvement from the ECB, which implies that they have monetary component; market reactions would be correlated with monetary shocks. Similarly, the market reaction to data releases cannot be used as they may be informative about common macroeconomic shocks.[19]

Once one rules out a direct causal relationship between the market reaction to foreign events and local/common economic shocks, the question remains: why does the market react? It could simply reflect the real or financial interlinkages between countries. However, these channels should be assessed in light of the following. First, the main results in this paper can be obtained solely by focusing on events in Greece; a small economy from a European perspective. For example, in the first quarter of 2012, Italy's banking system's exposure to Greek assets (private and public) was less than a tenth of a percent of GDP.[20] Yet Greece's vote in February 2012 caused a substantial decline in the Italian bond yield. To give context to the unusual size of the reaction: on the 15th of February 2012 (3 days after the Greek vote) the Italian statistical agency reported that Italy's GDP growth in Q4 2011 was 0.1% less than market participants expected,[21] more than the direct exposure to Greece, and yet the market's reaction to the data release was only to increase Italian borrowing costs by 1.1bp. Second, the relative response to events seems weakly related to cross-country exposures: Spain and Italy were much less exposed to Greece than was Germany in February 2012, yet German yields rose in response to the Greek vote.[22] Third the market reacted to events in Cyprus, a country an order of magnitude smaller than Greece. Fourth, the sharp reduction of the financial interlinkages between countries over the crisis period has not been met by a reduction in correlation between country borrowing costs.[23]

Real and financial interlinkages cannot be ruled out completely; however, there are other channels that can explain the extent of the reaction. Foreign events could trigger shifts in market sentiment or a reassessment of pre-existing information. This ties in with the idea that the crisis is self-fulfilling: foreign events act as the catalyst which shifts investors' expectations between equilibria.[24] Similarly, foreign events can influence local

---

[19]To see this, consider a specific example: the Italian bond market reaction to a Spanish industrial production release. Spanish industrial production is usually released before its Italian equivalent. Yet the two series are closely correlated, which likely reflects common shocks to their economies. This means that the Spanish data release is informative about Italian economic conditions due to the fact it contains information about those common shocks. In effect, the Spanish statistical office has private information about Italy which it reveals with the data release.

[20]This figure is available from the BIS locational banking statistics. Direct trade flows are also of a similar scale.

[21]Taken as the consensus forecast from the Bloomberg economic calendar.

[22]One can conduct an analysis of this nature more formally and regress the relative reaction to identified events on trade or financial interlinkages. The message from an analysis such as this is that trade interlinkages do not matter: the coefficient is insignificant and of the wrong sign. However, the market's reaction is weakly related to the size of interlinkages. Increasing the exposure of the domestic banking system to an event country by 1% of GDP increases the relative market reaction by 0.08bp. However, a 1% of GDP increase in exposure is substantial. For example, in the first quarter of 2012, when the Greek vote took place, the range between the most exposed country (Germany) and the least exposed (Finland) was 0.7% of GDP. If one restricts attention to just the four crisis countries in the sample the message is almost identical. However, an analysis of this nature is undermined by an endogeneity problem as the size of the exposure is a function of the events themselves.

[23]To be specific, of the four crisis hit countries considered in the main analysis of the paper: the average total financial exposure of each individual countries to all the others has fallen from 13.9% of GDP in Q1 2010 to 3.1% of GDP in Q1 2013; yields still remain highly correlated.

[24]The view that the crisis was in part self-fulfilling was one shared by ECB officials during the crisis. Mario Draghi specifically laid out these concerns when explaining the motivation for outright monetary transactions on September 6th 2012. The literature has emphasised that sovereign debt is vulnerable to self-fulfilling crisis either due to a lack of commitment on the part of policymakers (Calvo (1988)) or a sudden loss of liquidity (Cole and Kehoe (2000)). This feature of sovereign debt has motivated recent theoretical literature on the the Euro-crises, with the observations concerned with how self-fulfilling crises

borrowing costs through the impact which they have on the ability or willingness of unionwide authorities to intervene, either by eroding political support for interventions amongst taxpayers in creditor countries or by exhausting resources that have been precommited for interventions. The overidding concern during the crisis was the sustainability of the currency union as a whole. As the future of the monetary union became less certain, the cost of borrowing in any country would contain convertibility premia. Negative foreign events can shift the probability of the state of the world as being one in which the Euro is irreversible to one in which it is not.

I do not attempt to disentangle these channels. However, the above mechanisms are all means by which the sovereign risk premia are affected by factors external to the state of the local economy. This is the focus of this paper: I take an observed set of changes in the bond yield about which it is plausible to argue that the cause of the move was not an innovation to local economic conditions and therefore justifies an exogeneity assumption.

A final issue with this identification strategy is the nature of the systemic shocks. The innovations capture a change in bond yields caused by an intensification of the crisis as a whole. The shock is not country-specific. This means that there are additional channels through which systemic shocks may influence the economy. First, there is an additional uncertainty channel. Concerns over resolution of the institutional set up of the currency union will influence the economy above and beyond the uncertainty over the local government's ability to pay its debt. Second, even if direct interlinkages are not the primary reason for the transmission of events, the shocks hit neighbouring countries in the union as well (albeit asymmetrically); thus there may be additional effects via external demand or via a deterioration in pan-European market conditions. This is important to bear in mind when interpreting the results.

3.3 The proxy data

A key contribution of this paper is constructing a narrative of the crisis period in order to identify a set of events similar to those descibed above. Given these events I then time when they occur and measure the relavant bond market reaction to them. The proxy for the month for a local country is the sum of the local bond market reactions to foreign events. This is the version of the proxy, denoted m_t, that is used to identify the monthly structural shocks from the reduced form VAR. The proxy is constructed from July 2009 to March 2013.

The data appendix (appendix B.1) describes the methodology used in constructing the dataset and its summary statistics in detail. However, given the importance of the proxy to the main analysis, I highlight a few key features of the data here. In terms of the sources to isolate events, the approach taken here is to rely on news summaries.[25] I use the financial news sources *Bloomberg* and *EuroIntelligence* both of which compile a daily news briefing for European economic news. The former is released in the afternoon and the latter in the morning. Both contain around 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 include 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.

To be classified as an event, a news story must satisfy certain criteria. These are discussed in more detail in the appendix. However, the most important are: *(1)* the story must relate to a single crisis-hit country;

---

play out in a monetary union (see for example, Aguiar M and Gopinath (2012); Cooper (2012); Corsetti and Dedola (2013); Lorenzoni and Werning (2013)).

[25] This is the approach taken in Beetsma et al. (2013). 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.

specifically, either Greece, Cyprus, Portugal, Ireland, Italy or Spain; *(2)* the event must be timeable in the sense that it is possible to isolate when it occurs so as to determine the market's reaction. There also limits on the type of news considered, generally events are correspond to policy or political announcements. News revolving around private companies or (importantly) data releases do not enter the proxy.

Events are timed to the minute when the first headline related to the story appears on the *Bloomberg* newswire. An untimeable event is one where it is impossible to identify an initial headline in an objective fashion. For events that can be timed, the market reaction is considered over a 20 minute window on either side of the initial headline. The market reaction is defined as the change in the mid-yield to maturity on the benchmark 10 year sovereign bond. 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 financial institutions.

A critical point is dealing with events that occur when markets are closed. European policymakers' 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 that there is more chance of another piece of important information being released and distorting the market's reaction. As a compromise, and 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 period where the market is open presented in section 6.

It is worth noting that some of the events may have little informational content or may be entirely anticipated by the time of the announcement. However, beyond the use of "headline" stories for overnight events, there is no need for any additional indicator of event importance, predictability, or whether the news is positive or negative. This information is contained within the market's reaction.

Another concern is simultaneous events. Given the high frequency of the dataset, another story breaking simultaneously 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 contemporaneous news. The structure of the dataset means it is straightforward to single out any foreign event that overlaps with a local event and these are not 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, pan-European and certain international data releases are obtained from the *Bloomberg* economic calendar and events that would overlap with a windows around these releases are similarly omitted.[26] Events that overlap with ECB decisions and press-conferences are not included. Lastly, any country-specific event that overlaps with the announcement of a pan-European policy intervention is omitted. Such events are isolated using the ECB's time line of the crisis[27] and are timed in an identical fashion to the country-specific events as described above.
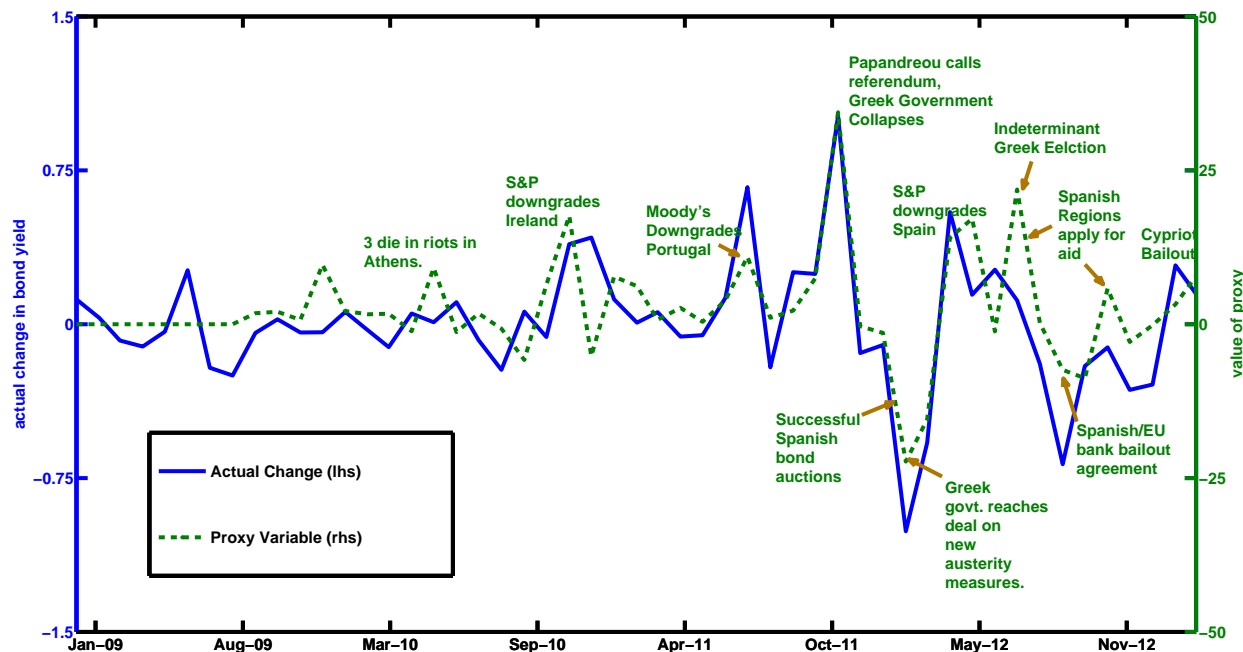
Despite these steps it is impossible to be completely certain as to what drives the market move at any point in time. There may be news from outside the Euro Area driving yields, as well as private information or rumours that cannot be picked up using the approach taken here. Furthermore, yields are 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 is 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 endogeneity but they will introduce measurement error which motivates the specification described in section 4.

---

[26]I refer readers to the online appendix for the complete list of omitted data events.

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

Figure 2: Proxy variable and actual changes in the bond yield: Italy

To give a measure of confidence in this procedure, the results presented are robust to different assumptions over the construction of the proxy, particularly over the interval window considered for market reaction and the types of events included - see section 6.

Completing the process 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. However, the depth of coverage and variety of events also means that it is difficult to describe completely the narrative in a concise manner within the main text of this paper. Therefore, readers are referred to the appendix of this document for greater detail and the online appendix for an exhaustive list of narrative events.

Greece and Cyprus are not included in the final empirical analysis due to inconsistent availability of intraday bond market data. The final proxy variable is constructed only for Italy, Spain, Ireland and Portugal. Greek and Cypriot events are included in the proxies, however. To reiterate: local events are not included in the proxy. Once one has filtered out domestic events, events that overlap with other news and events outside trading hours which are not "headline" news, the proxy variables contain between 390-450 events depending on the country. This corresponds to 9-10 per month over the sample.

Figure 2 shows the proxy variable for 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 strong at 0.75. For the other included countries the correlation is not as strong;

figure 3 presents the graphs for Spain (correlation: 0.65), Ireland[28] (correlation: 0.55) and Portugal (correlation: 0.38). These correlation coefficients are all statistically significant at the 1% level. From the point of view of the empirical strategy, what matters is not the correlation with the actual change in the bond yield but rather with the residuals in the reduced form VAR. The online appendix contains a formal analysis of the strength of the relationship between the proxy and the reduced form residuals of the VAR model. The result of this analysis is that the fit is sufficiently strong such that one needs not be concerned with weak proxy problems.

The variation in the proxy series is driven by large market reactions to a small proportion of events rather than more moderate reactions to every piece of news. Figure 4 illustrates this by ordering all events included in the proxies according to the square of the market reaction and then plotting the cumulative contribution of each ordered event against 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 important events in each proxy; still a large number but much less than the total number of identified events.

I also carry out a battery of stastistical checks on the proxies to see if the data matches the assumptions underlying the identification strategy. For the sake of brevity I again refer readers to the online appendix for exact details regarding these analyses and the associated tables and figures.

The reason for using high frequency data is that existing public information is already reflected in market prices; thus market reactions are orthogonal to preceding shocks the participants are already aware of. The natural check of the validity of this assumption is to see whether the proxy is predictable. I do this by running predictive regressions on the local market reaction to foreign events aggregated at weekly, biweekly and monthly frequencies. I find that these reactions are not explained by past events, either local or foreign, or macroeconomic data. 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.

As well as being unpredictable, aggregations of market reactions to foreign events should be uncorrelated contemparaneously with reactions to local events. In the data, there is evidence of such a correlation between market reactions to local and foreign events for the Italian bond market but not in the other countries. In small samples spurious correlations are always possible. In this case the correlation is attributable to a single outlier: in the month of November 2011 the collapse of the Greek and Italian governments occured simultaneously.[29] Abstracting from this observation removes the correlation.

A final empirical test is to gauge the extent to which the foreign events included in the proxy are a reaction to changes in local economic, monetary or fiscal conditions. One way to capture this is to see if the proxy is correlated with the market reaction to local economic and fiscal data releases or ECB announcements. This can be thought of as a test of whether foreign events are correlated to local macroeconomic shocks that are being captured by a data surprise.[30] The result of this analysis is negative: the proxy is uncorrelated with data surprises.

---

[28]Irish tick data is not available between May 2011 and October 2011. In this period the Irish proxy is constructed using the daily change in yields during foreign events that were "headline" news, i.e. only major events.

[29]There is no evidence that the collapse of Silvio Berlusconi's government was a direct consequence of George Papandreou's resignation in November 2011 and vice versa.

[30]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.

4 Methodology

This section lays out the econometric methodology and describes how the proxy SVAR identification regime in section 2 is implemented. The reduced form model follows a Bayesian panel SVAR with cross-sectional heterogeneity in the slope coefficients and covariance matrices. Implementing the proxy SVAR in a Bayesian framework adds a layer of complexity in the sense that one cannot directly apply the two-stage procedure used in the frequentist approach. The reduced form VAR is monthly, yet the proxy variable is best viewed as an aggregation of censored high frequency events that are stochastic. I show how this aggregation can be dealt with by using a leptokurtic distribution to capture the effect of the censoring process when estimating the relationship between the narrative proxy and the residuals. From here it is straightforward to derive the likelihood of the proxy conditional on the reduced form and, using Bayes rule, the joint posterior.

The Bayesian approach has a number of advantages. First, it is sufficiently flexible so as to allow for additional panel assumptions at the identification stage. Second, the information from the proxy is used when estimating the VAR coefficients, even though the proxy is not directly included in the reduced form model and, third, credible intervals can be readily constructed that include uncertainty from both the reduced form VAR estimation and the identification procedure.

4.1 The reduced form Panel VAR

The primary feature of the panel VAR model 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. 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 and scalars are lower case symbols, matrices are uppercase symbols, the indices $c = 1, ..., C$, $l = 1, ..., L$ and $t = 1, ..., T$ denote (local) 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:

$$y_{c,t} = \sum_{l=1}^{L} B_{c,l} y_{c,t-l} + \Gamma_c z_t + u_{c,t} \tag{2}$$

Where y_{ct} is a $N \times 1$ vector of endogenous country variables, $B_{c,l}$ 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 2 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 $B_c = [B_{c,1}, ..., B_{c,l}]$'. Lastly, I define the vectorised data and parameter terms as $y_c = vec(Y_c)$, $\beta_c = vec(B_c)$ and $\gamma_c = vec(\Gamma_c)$.

4.1.1 The first level of the hierarchy

The first level of the hierarchy governs the statistical form of the individual country models. The likelihood for the model corresponding to country c is given by:

$$p(y_c|\beta_c,\gamma_c,\Sigma_c) = N((I_N \otimes X_c)\beta_c + (I_N \otimes Z_c)\gamma_c, (\Sigma_{c,u} \otimes I_{T_c})) \tag{3}$$

A non-informative prior is assumed for γ_c in each country: $p(\gamma_c) \propto 1$. The country slope coefficients β_c are assumed to have an exchangeable Gaussian prior with common mean $\bar{\beta}$ and variance $\Lambda_{c,\beta}$ which is country specific:

$$\beta_c|\bar{\beta},\Lambda_{c,\beta} \sim N(\bar{\beta},\Lambda_{c,\beta})$$

The parameter vector $\bar{\beta}$ serves as the cross-country mean of the slope coefficients. I depart from Jarocinski (2010) by imposing the prior that the covariance matrix of the residuals is also drawn from a common distribution (in this case inverse-Wishart) with a common scale parameter $\bar{S}$:

$$\Sigma_{c,u}|\bar{S},\kappa \sim iW(\bar{S},\kappa) \tag{4}$$

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.[31] 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.

4.1.2 The second level of the hierarchy

The second level of the hierarchy governs the common cross-country elements, specifically the prior distributions of the hyper-parameters in the country models. I let the data determine the common means and use a diffuse prior for both $\bar{\beta}$ and $\bar{S}$: $p(\bar{\beta}) \propto 1$ and $p(\bar{S}) \propto |\bar{S}|^{-0.5(N+1)}$. The degree of shrinkage applied to slope coefficients β_c is governed by the country-specific covariance matrices $\Lambda_{c,\beta}$. I assume this covariance matrix can be decomposed into a country-specific positive definite matrix ($L_{c,\beta}$) and a common scale parameter contained in the set of positive real numbers (λ_β):

$$\Lambda_{c,\beta} = \lambda_\beta L_{c,\beta}$$

The matrix $L_{c,\beta}$ is 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). This form of specification for $L_{c,\beta}$ adheres to a 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 $L_{c,\beta}$ 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_β. This hyperparameter 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_β it is useful to consider the two extreme cases. An estimate of $\lambda_\beta = 0$ 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 homogenous panel VAR. Conversely, as $\lambda_\beta \to \infty$ the distribution about $\bar{\beta}$ is sufficiently diffuse such that there is no information contained within

---

[31]Allowing the residuals to be correlated across countries is attractive from an efficiency perspective but is computationally intensive. Re-estimating the benchmark specification allowing for such correlations does not meaningfully alter the results thus for computtional convienence I restrict the residuals to be uncorrelated.

the common mean. As a result the posterior means of β_c are almost equivalent to those as if each country has been estimated separately.[32] Any increase in λ_β is equivalent to a reduction in shrinkage and allowing the estimated country models to become increasingly different. The parameter λ_β, therefore, determines how close the model is to either of the two extremes of country-specific slopes and homogenous slopes.

While it is possible to provide an interpretation as to what changes in λ_β 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_β 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:

$$\lambda_\beta | s, v \sim IG_2 \propto \lambda_\beta^{\frac{-v+2}{2}} exp\{-\frac{1}{2}\frac{s}{\lambda_\beta}\} \tag{5}$$

with hyper-parameters s and v. The hyper-parameters are specified, as recommended Gelman (2006), such the standard deviations for the individual coefficients have a uniformly distributed prior over the positive portion of the real line, i.e. $p(\lambda_\beta) \propto \lambda_\beta^{-1/2}$. This equivalent an improper prior with $v = -1$ and $s = 0$.[33]

The parameter κ plays a similar but inverted role to λ_β for the covariance matrices. As $\kappa \to \infty$ the distribution in equation 4 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. 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.[34]

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. One can attempt to use a non-informative prior for κ and simulate the univariate non-standard conditional distribution using an adaptive rejection step in the sampling procedure. However, this parameter appears weakly identified when estimated. Therefore, I treat κ as deterministic. To be conservative I choose the value of $N + 2$ (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.

4.2 Identification

4.2.1 Aggregating high frequency events into the proxy.

Identification relies on the proxy SVAR setup described in section 2. However, the simple linear relationship between the proxy and the reduced form model described in equation 1 may be construed as an overly strong assumption in the context of the strategy used here. It is also unclear what the appropriate distributional assumption is over ther term ω_t. The construction of the proxy as described in section 3.3 presents several issues from an econometric perspective. The variable is an aggregation of high frequency bond market reactions which are themselves stochastic. Events are not continuously observed and the event inclusion

---

[32]Note that due to the prior assumption on common covariance matrices in equation 4 the estimates of β_c still depend on the cross section even as $\lambda_\beta \to \infty$.

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

[34]In practice, it is not possible to apply the no-shrinkage case and estimate a common mean covariance matrix that conforms with equation 4; 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.

criteria means certain types of news are omitted. Thus the proxy is effectively the aggregation of censored random 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 from their initial reaction. Furthermore, the market reaction may be slowed by 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. Rumours may leak in advance, attenuating the response. This means that there will be measurement error contained in the observed reaction to each event. However, it suggests there may be also scaling effects: the initial market reaction may over- or under-shoot in a regular fashion.

Despite the empirical issues with this high frequency set up, it is possible by making simplifying assumptions to return to the simple linear model in section 2. This linear framework can be used for the purposes of identification without any explicit estimation of the high frequency statistical process driving the bond yield or event occurrence.

The true underlying data-generating process is continuous. However, rather than attempt to combine a discrete time VAR and continuous events, I approximate the event-generating process by assuming it occurs over a series of sequential, non-overlapping, discrete time windows. These windows are denoted thusly: $d = \{1, 2, \ldots, M\}$ where M is the total number of windows in the month. For the purposes of exposition consider events at a daily frequency (d is one day); conceptually, this can be extended to a higher frequency by simply narrowing the time window. For notational convenience the country subscript, c, is dropped for this subsection.

Let m_{dt} be the recorded market response on day d of month t. I assume this has the following data-generating process:

$$m_{dt} = \vartheta_{dt}(\psi\varepsilon_{dt} + v_{dt}) \tag{6}$$

The variable ϑ_{dt} is an indicator for censoring, taking a value of one or zero. If an event is not observed on a particular day then no market response is recorded such that m_{dt} (and ϑ_{dt}) take a value of zero. If an event occurs then the observed market reaction is assumed to be the sum of the scalar structural shock, ε_{dt}, that occurs in that time window, scaled by parameter ψ, and some independent measurement error v_{dt} ($v_{dt} \sim NID(0, \sigma_v^2)$).

I follow the special case in Mertens and Ravn (2013b) and assume that the censoring process is random; that is to say ϑ_{dt} is an independent variable that takes a value of 1 with probability p and zero otherwise (i.e. $1 - p$ is the probability that an observation is censored).[35]

I assume that the daily series of scalar structural shocks sum perfectly to create a monthly shock of interest:

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

The daily structural shocks have the property: $\varepsilon_{dt} \sim NID(0, 1/M)$ such that the monthly shock is Gaussian with unit variance. If one interprets ε_{dt} as a structural shock to the bond yield, this aggregation assumption

---

is equivalent to the yield following a process close to a random walk at a high frequency. The data supports this.[36] One may be concerned that the market reaction to an events decay over the course of the month; in section 6 I show that restricting events only to the first week in the month leads to similar results.

The monthly proxy is the aggregate of the observed market reactions:

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

The critical identifying assumption is that both ε_{dt} and v_{dt} are uncorrelated with any other structural shocks that hit the economy in month t. These assumptions are sufficient for the relationship between the proxy and the structural shock to sastisfy the conditions in section 2; namely, $\mathbb{E}(m_t \varepsilon_t) = \phi$, $\mathbb{E}(m_t \tilde{\varepsilon}_t) = 0$ where ϕ is a scalar. The assumption of that v_{dt} is uncorrelated with other structural shocks has two justifications. First, time windows where important information about other shocks are revealed (e.g. data releases) are excluded. Second, that the time windows are sufficiently small such that any underlying correlation between yields and the state of the economy tends to zero. Furthermore, it is assumed that $\mathbb{E}(m_{dt} m_{st}) = 0$ and $\mathbb{E}(m_t m_s) = 0$ $\forall t, s$.

I prove in appendix A that, given these assumptions, the relationship between m_t and the reduced form residuals, u_t, is:

$$m_t = \Upsilon' u_t + \omega_t \tag{7}$$

where $\omega_t | u_t \sim iid(0, \sigma_\omega^2)$, $\Upsilon' = p\psi a_1$ and $\sigma_\omega^2 = pM\sigma_v^2$.

This returns us to the linear specification at a monthly frequency. There is one last problem, namely that the distribution of ω is non-Gaussian due to the censoring process. It has the same support as the normal and retains symmetry. However, there is a difference in the fourth moment as censoring takes mass from the tails to places it at the mean. This means 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 7, extended to include t-errors, has been covered extensively in the literature (see Geweke (1993) for the classic treatment). Letting $\mathcal{M}$ be the matrix of proxy observations stacked over time, one can approximate the conditional distribution of the proxy as:

$$\mathcal{M} | U, \Upsilon, \sigma_w^2; \nu \sim t(U\Upsilon, \sigma_\omega^2 I_T; \nu) \tag{8}$$

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.

Extending the model away from purely random censoring to allow for the probability of being censored, p, to depend on ε_{dt}, is technically complex. Unconditional random censoring ensures a homoskedastic, linear relationship between the proxy and the reduced form residuals, with error of known variance, as in equation 8. Conditional censoring can be expressed in a similar way to equation 7. Υ is defined differently but still has the form $\phi a_1'$, where ϕ is a scalar. The distinction is the errors are heteroskadastic conditional on the latent high frequency shocks with no tractable analytical expression for the variance of the error term conditional on u_t. However, as is well known, one can reinterpret Studentised errors as a form heteroskedasticity; draws

---

[36]To clarify, this assumption implies that the yield should enter the VAR as the monthly close rather than the monthly average used in the benchmark specification. The average yield is more commonly used for macroeconomic analysis to clean high frequency fluctuations present in the close. In this paper, the results are not sensitive to using the close or the average - or, indeed, other forms of aggregation.

from the tail of the distribution are conceptually the same to those drawn from a distribution with higher variance. Indeed, as discussed in the appendix, this is how the Studentised errors are implemented in practise.

4.2.2 Priors on the identification parameters.

As much of the data entering the model is at a monthly frequency, jointly estimating the 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. This is equivalent to the maximum likelihood estimator of p. From the narrative series described in section 3.3 this results in a $p = 0.15$ (averaged across countries). Similarly M is deterministic and is set to 30. The extent of the excess kurtosis that arises from the censoring process is a function of only 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, I simulate the ω in equation 7 for arbritrary σ_v^2 and ψ. I then calculate ν by matching the fourth moment. Plugging in the numbers from above results in $\nu = 13$ to the nearest integer. It is worth emphasising the results are not too sensitive to this choice of parameter. Figure 5 illustrates this approximation 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. Visually, for these values of p and M, the Student-t approximation appears closer to the true density of ω when compared to the normal.

The parameters in equation 7 are assumed to be country specific with a cross-sectional relationship along the same lines as the parameters in the reduced form VAR. The country slope coefficients Υ_c are assumed to have prior normal distribution with a common mean, $\bar{\Upsilon}$, and variance, $\Lambda_{c,\Upsilon}$, which is country-specific:

$$\Upsilon_c | \bar{\Upsilon}, \Lambda_{c,\Upsilon} \sim N(\bar{\Upsilon}, \Lambda_{c,\Upsilon})$$

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, $\Lambda_{c,\Upsilon} = \lambda_\Upsilon L_{c,\Upsilon}$, with $L_{c,\Upsilon}$ set along the same lines as $L_{c,\beta}$. The parameter λ_Υ has the same prior as λ_β as described in equation 5, and plays the same role. The parameters, $\bar{\Upsilon}$, and, $\sigma_{c\omega}^2$, have diffuse priors: $p(\bar{\Upsilon}) \propto 1$ and $p(\sigma_{c\omega}^2) \propto \sigma_{c\omega}^{-1}$.

4.3 Discussion of the impact of the hierarchical model

For the purposes of 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 slope coefficients in both the reduced form and identification stages leads to estimates that take 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).[37] For the model as a whole, what determines how close the estimation is to each extreme is the parameter λ_β in the case of the reduced form slope coefficients and λ_Υ in the case of the identification model. However, the extent of the pooling also varies from country to country depending on how well each country's model fits the data. 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 the precision of that country's model. Thus 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.

---

[37]In the case of the β_c terms there is an additional penalty term arising from the need for the reduced form residuals to fit the proxy.

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.

Partial pooling has implications for the identification strategy. It implies that in a country where the proxy variable 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.

4.4 Estimation

An advantage of working with the $\mathcal{A}$ matrix for the proxy SVAR is that joint likelihood of the data is heirarchical and straightforward to define. To simplify the notation define the parameter space in the model as Θ, 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 $\mathcal{M} = \{\mathcal{M}_1, \ldots, \mathcal{M}_C\}$. By Bayes rule the likelihood of the data is equal to the product of the likelihood of the proxy variables conditional on both the reduced form model and data and the likelihood of the reduced form model unconditional on the proxy: $p(\mathcal{M}, Y|\Theta) = p(\mathcal{M}|Y, \Theta)p(Y|\Theta) = \prod_c p(\mathcal{M}_c|Y_c, \Theta)p(Y_c|\Theta)$. The form of $p(Y_c|\Theta)$ is given in equation 3. The condition density of the proxy, $p(\mathcal{M}_c|Y_c, \Theta)$, is defined in equation 8. As is well known, regression with Student-t errors can be rephrased as a mixture of normals. Therefore $p(\mathcal{M}_c|Y_c, \Theta)$ is Gaussian and from above $p(Y_c|\Theta)$ is Gaussian; hence, the joint density, $p(\mathcal{M}_c, Y_c|\Theta)$, is also Gaussian.

For estimation, 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 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 full form of the sampling algorithm is laid out in appendix C.

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 January 2008 to March 2013. Regarding 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 convertibility 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.[38]

To capture the impact of elevated sovereign borrowing costs on private financial conditions in a concise manner a composite measure of the private cost of finance is used. This is calculated as the weighted average in the cost of equity, debt securities and bank credit for non-financial corporations and households in each country. The series is computed internally by the ECB by weighting yields on the various sources of finance in accordance with flows of new lending.

Lastly, as a measure of the fiscal stance, the monthly general government primary balance is included as an annualised percentage of nominal GDP. The comparable fiscal data (across the countries) is only available at a quarterly basis from the flow of funds dataset available in Eurostat (net/lending or borrowing by the general government sector plus interest payments). However, all four countries publish monthly fiscal data using a variety of definitions. To generate a monthly fiscal series that has the same definition across countries I use the regression based interpolation methodology of Mitchell et al. (2005) on the quarterly series using the country specific monthly fiscal balances as interpolands. A deficit is a negative reading. Exact details of the data sources and their construction are laid out in appendix B.4.

This sets $N = 8$. The set of deterministic variables, Z, is set to include only a constant for all countries. The trended series (the CPI and Industrial Production) enter the VAR in log year-on-year differences; other series are included in levels. The lag length L is set to 2.[39]

The posterior is simulated using 600,000 draws from the MCMC sampler in the appendix; the first 100,000 are discarded as a burn-in and the remaining chain is thinned by a factor of 50 leaving 10,000 draws for inference. Results presented are the median of the 10,000 retained draws and 95% uncertainty bands are computed using standard Bayesian Monte-Carlo methods.[40]

As with most VARs, the model can be used to produce three main results of interest. The impulse response analysis provides an assessment of the propagation of systemic shocks to the included variables. Variance decompositions give an indication of the relative importance of systemic shocks in explaining unanticipated fluctuations in the included variables. Last is a 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 systemic shocks to the borrowing costs and unemployment rates in the crisis-hit countries.

5.2 Benchmark 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 regard 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 may lie in the soothing impact of policy: monetary easing follows the shock, albeit with a lag, with a peak response of a 40bp decline in the Eurepo and 10 year

---

[38]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. For example, it is not clear how the contracts would be honoured in the event of Euro break-up - see Nordvig and Firoozye, 2012. However, the OIS has less apparent embedded risk of redomination when compared to the German bond.

[39]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 adding more lags until the median estimated residuals display no serial correlation. This lag-selection matches the Schwarz-Bayesian criterion assessed on the panel version of the VAR with homogenous slope coefficients and covariance matrix. So, $L = 2$, it is robust to alternative lag selection criteria. In the online appendix I present robustness checks to alternative lag lengths.

[40]The algorithm appears to mix well; the standard diagnostic tests are passed. I refer readers to the online appendix for convergence diagnostics.

OIS rates after 4 months. The unemployment response is statistically insignificant on impact but the shock propagates and leads to a peak of 0.9ppt 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.

To give better context to the results, it is worthwhile attempting to rephrase the impact on the monthly measures of output in terms of GDP. Recent estimates of Okun's law in Europe during the crisis period (see for example Ball et al. (2013)) suggest a coefficient of around 0.5, so the unemployment response is consistent with a 1.8ppt GDP reduction at the peak. For the countries in question industrial production is about twice as volatile as GDP over the sample period, so the industrial production response is consistent with 1.0ppt off GDP growth. These two responses are not completely consistent in scale but given the uncertainty involved in the estimates they are not too dissimilar.

Figure 7 presents the results of country-specific models. What stands out is the similarity of the responses. This suggests the data is returning a model which is close to the mean country estimates. The impulses are similar across across countries both in terms of the impact response and the dynamics that follow, thereby supporting a model which is close to a slope homogeneity assumption in both the reduced form and identification stages. This is also evidence that the countries did behave similarly in response to increases in sovereign risk during the crisis period.

In a model with improper priors it is not possible to construct standard Bayesian likelihood ratio tests on a 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 in 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 144.19, while the fully-pooled model returns a DIC of 272.90 and the country by country estimate returns 938.40. 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.[41]

Figure 8 presents forecast error 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 50% of the variation in the bond yield is explained by the shock. At longer-term horizons the importance of shock for the variance of yields seems to dissipate; this matches the impulse responses. A second finding of note is that 45% of the forecast error variance 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.

In terms of other variables, systemic shocks explain around 35% of the variation in the private cost of finance on impact, but the effect disappears 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

---

[41]The results from the fully pooled and partially pooled are available in the online appendix.

explanation could simply be that industrial production is a noisy series and thus a greater proportion of its forecast error is explained by its own volatility.

This result - that the systemic shocks were an important driver of unemployment dynamics over the crisis periods - raises the question: what are the key channels by which elevated sovereign yields feed into the economy? In terms of the fiscal channel, 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 the series 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; increasing borrowing costs do not lead to a positive response in the primary balance.

It is worth considering, therefore, if this represents 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. Therefore, one cannot rule out a fiscal channel, but the evidence is only sufficient to say that increases in sovereign borrowing costs inhibit automatic stabilisers, and not that there is absolute fiscal tightening.

A couple of further channels emerge from the literature; first, that sovereign debt crises are often associated with damage to the financial system and banking crises (Reinhart and Rogoff (2009) and De Paoli et al. (2009)).[42] Sovereign bonds are used for collateral by banks; any fall in their value constrains the supply of liquidity to the banking system. Losses on bonds also reduce bank capital and constrain the supply of credit to the private sector. Private sector bonds are also affected. The notion of the sovereign ceiling (see Durbin and Ng (2005)) suggests that no corporate bond should yield less than their sovereign (due to the risk of expropriation). Corsetti et al. (2013) embed the notion that sovereign risk impacts the private cost of finance, which they describe as a "sovereign risk channel", into a New Keynesian model with financial frictions where private financial conditions are inherently linked to changes in sovereign risk premia. Under these circumstances, the intensification of a debt crisis acts as a form of financial shock curtailing private demand.

A second mechanism that emerges from the literature on debt crisis is an "external channel": the crisis reduces access to foreign markets. Countries experiencing debt crises often lose access to international capital markets (Arteta and Hale (2008)) and experience sharp declines in bilateral trade (Rose (2005)). Mendoza and Yue (2012) argue that this can result in a mechanism that inhibits the supply side of the economy: the loss of imported inputs which can not be substituted for domestically acts as a productivity shock, similar to a working capital channel.[43]

These two channels are not mutually exclusive but is worthwhile exploring whether they are apparent in the data. I augment the VAR to explore these issues further.

5.2.1 The sovereign risk channel

In the benchmark specification the pass-through of a systemic shock that raises sovereign yields by 100bps onto private borrowing costs is less than one-for-one: yields on private finance increase by only 40bp on impact and the effect is short-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

---

[42]Indeed, as Reinhart and Rogoff (2011) describe, it is often financial stress that leads to sovereign crises.

[43]As a clarifying remark: Mendoza and Yue (2012) refer to epsiodes of default rather than just a high sovereign risk premia. However, the empirical stylised facts used to motivate their mechanism are not restricted to default episodes and can be generalised in this context.

have a high weight in the composite cost of finance. However, bank lending rates are also sticky and slow to respond to market conditions.

To explore the sovereign risk channel in more detail I augment the benchmark VAR by decomposing the composite cost of finance index into its three components. The average interest rate on loans and corporate bonds can be extracted as a component of the ECB composite cost of finance. As the equity yield is not consistently 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. To account for the potential stickiness in loan rates it is important to include a variable that captures quantities of private finance as well as its price. The supply of loans by local banks to the domestic non-financial private sector is included for this purpose (sourced from the ECB).[44]

Figure 9 presents the mean country impulse response functions for VAR augmented with these four additional variables. The corporate bond yields respond a little more than one for one as compared to the sovereign: on impact the yields increase by about 130bps. Corporate yields also follow a similar dynamic in response to the shock. This result is consistent with the idea of a sovereign ceiling (see, for example, Durbin and Ng (2005)) where a country's corporates can never borrow more cheaply than their sovereign, and where this would imply at least a one to one response to a shock to the sovereign risk premium.[45] Equity prices fall with a peak decline of 7 percentage points. The two market-based variables seem to have large reactions. Loan rates react to a smaller extent in contrast, rising by around 25bp. However, there is a simultaneous decline in the quantity of credit: at its lowest ebb credit growth falls by around 1 percentage point. The responses of the original variables in the VAR are similar to the benchmark.

These results confirm that systemic shocks pass through into the domestic private financial system. Furthermore, the variance decomposition in figure 8 provides evidence that the shocks were a significant driver in the fluctuation in private financial conditions during the crisis. The sovereign risk channel appears to play an important role.

5.2.2   The external channel

A distinctive feature of the Euro Area debt crisis when compared to other debt crises is the presence of the ECB as an external source of finance. Euro area countries suffering from the crisis were less vulnerable to sudden stops of external private capital flows due to the ability of their domestic financial institutions to use the ECB's liquidity operations to obtain central bank financing. This source of finance is limited only by the eligible collateral available to the domestic banking system. One way to gauge this substitution of private capital for official financing is the change in the country's national central bank's balance with the wider Eurosystem, otherwise known as the change in the Target 2 balance. This is available at a monthly frequency from Steinkamp and Westermann (2012). I rescale these figures into a percentage GDP term, where a negative figure is equivalent to a liability.[46]

I attempt to capture the trade effects of systemic shocks using monthly goods trade data. I take the

---

[44]Outstanding loans are calculated as a notional stock based on transactions to remove revaluations effects.

[45]Durbin and Ng (2005) do show that some corporates can escape the sovereign ceiling if they have large external operations or export earnings. However, this effect will be lost in this data as bond index used is an aggregation of corporate issuers in the country in question.

[46]To see why this is true: imagine a situation where investors withdraw deposits from a Spanish bank and deposit them in a German bank. The German bank then holds these new deposits as reserves at the Bundesbank. Correspondingly, the Spanish bank needs to replace the lost funding and therefore goes to the Bank of Spain to borrow collateralised at the its discount window. This leaves the Bundesbank with a liability (the reserves of the German bank) and the Bank of Spain with an asset (the lending to the Spanish bank). The Bank of Spain finances its asset by borrowing from the Eurosystem (a Target 2 liability) while the Bundesbank lends the deposited reserves to the rest of the Eurosystem (a Target 2 asset). This clears the Eurosystem balance sheet. In effect, what we observe is a central bank loan from the Bundesbank to the Bank of Spain of the same size as the private capital flight. Although neither institution is specifically authorising such a transaction, it is just a feature of the Eurosystem.

trade balance (as a % GDP) of imports (as a % GDP) and leave exports as the balancing item. To explore the external channel, I augment the benchmark VAR with these three additional variables and reestimate the model. Figure 10 presents the mean country impulse response functions. As above, the responses of the original endogenous variables in the VAR are similar to the benchmark.

The observed decline in the Target 2 balance is consistent with private capital flight worth 1.9% of GDP on impact and 1.8% the following month - i.e. private capital worth 3.7% of GDP leaves the country. There is no correction. The outflow appears permanent in this sample. The loss of private external capital may be a further cause for the sovereign risk channel observed above. Replacing largely uncollateralised private finance with collaterised borrowing from the Eurosystem has the potential to impose an additional source of stress on the local financial system if some financial institutions are collateral constrained.

There is also evidence that systemic shocks generate an external adjustment: the trade balance improves by about 0.6% of GDP on impact and has a relatively persistent response. Imports decline by about 0.7% of GDP. The move in the balance and imports suggests that exports either experience a small decline or are about stable. The results are consistent with the mechanisms highlighted by Mendoza and Yue (2012); specifically, capital flight and a decline in imports.

However, while the evidence supports the existence of both channels it is not clear which, if either, is dominant. Mendoza and Yue (2012)'s working capital channel is supply-driven while Corsetti et al. (2013) rely on the impact disruption to the financial system has on demand. A natural way to distinguish between the two is to look at the inflation response; however, in the benchmark this is indistinguishable from zero. If the working capital effect was the dominant channel by which the elevation in sovereign yields influence the real side of the economy then the impact would be inflationary. In some specifications, such as the one presented in figure 9. and some of the alternative specifications presented in the online appendix, inflation does increase in response to the shock. So a supply driven mechanism is potentially present in the data, but the result is not robust. It is also worth noting that the inflation response is a potentially imperfect test in this context due to the reliance on increases on indirect taxes as a fiscal measure during the crisis.

5.3 Counterfactual analyses

In order to gauge the extent of the contribution of systemic shocks to changes in the observed time series of macroeconomic variables 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. 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 credible intervals. I first focus my attention on the counterfactual sovereign bond yields and then switch attention to macroeconomic consequences by considering counterfactual unemployment rates.

Figure 11 presents the results of this analysis on the bond yields for the four countries in the sample; on the top panel is the actual versus median counterfactual for the 10-year bond; the bottom panel has the difference between the two alongside accompanying Bayesian confidence intervals. For want of a better term, I refer to this difference between the true sovereign bond yield and its counterfactual equivalent as a "systemic premium". It is important to emphasise that this systemic premium is a deviation from the sample baseline caused by the systemic shocks, rather than an absolute level of systemic sovereign risk.

Several points stand out. First, and 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 is 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 trough-to-peak move in yields across the four countries can be explained by systemic shocks. This order of magnitude is about what one would expect given the variance decomposition result; 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 in 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 but are not 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 - which can be interpreted as an improvement in economic or monetary conditions lowering sovereign borrowing costs.

In Portugal, the counterfactual suggests that the yield should be slightly higher than observed , to the extent that the systemic premia is statistically significant - and negative - near the end of the sample. This effect lasts for 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 reading 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, then 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.

Figure 12 presents counterfactuals for the unemployment rates in the four countries in the sample. As above, the top panel is the actual versus median counterfactual for the unemployment rate; the bottom panel has the difference between the two alongside accompanying Bayesian confidence intervals. As is evident from the data unemployment rises sharply at the start of the sample, stabilises then rises again. This pattern is apparent with slight variations in intensity and timing across countries. In the early part of the sample, 2008-2009, actuals and counterfactuals align tightly. This is to be expected. For example, one would not anticipate that the 10 percentage point rise in unemployment in Spain between 2008 and 2009 was explained by an innovation to the riskiness of the sovereign. A financial shock during the initial intensification of the crisis appears the more likely culprit.

However, once tensions begin to arise in the sovereign debt market observed unemployment rates begin to distance themselves from their counterfactual equivalents in a manner distinguishable from zero. This occurs in mid-2010 in Ireland and Portugal and in mid-to-late 2011 in Spain and Italy which is roughly in line with when the crisis reached those countries plus a few months to allow a pass through to the unemployment rate. By the end of the sample unemployment is 4 percentage points higher than it would have been in the absense of the systemic shocks in Portugal and Ireland and around 1-1.5 percentage points higher in Italy and Spain. If one considers the rise in unemployment since the mid-2010 around a third of the move in Spain is explained by systemic shocks with an equivalent figure of one quarter for Italy. In Portugal almost all the recent rise can be attributed to systemic shocks while the dynamics in Ireland are dramatic: the model predicts that unemployment would have fell markedly in the absence of systemic shocks. In all, the unemployment counterfactuals reaffirm that sovereign risk was an important driver of this variable over the crisis period matching the evidence from the variance decompositions.

5.4 Discussion and Policy Implications

From a policy perspective a couple of discussion points emerge from these results. Firstly, it is clear that an increase in sovereign risk premia does pass through to private financial conditions. Although the results are conditional on the policy regime in place, this evidence does support the notion that measures to insulate the private financial system from the riskiness of the sovereign are effective.

Second, the response of policymakers to the crisis has been to focus on fiscal consolidation in order to lower sovereign risk premia. The argument in favour of doing so lies in the concept of expansionary austerity: that the fall in government borrowing costs will more than outweigh the negative effects of contractionary fiscal consolidation. This paper provides some evidence to quantify the first half of this argument, the impact of a fall in yields on the economy. Detailed reviews of the evidence for the second half of this argument, i.e. the effect of fiscal policy on output, are available elsewhere in the literature. For example, Ramey (2011a) suggests a probable range for spending multipliers between 0.8 and 1.5 for the US and, crucially, over sample periods where government borrowing was not constrained such that there was little feedback from the fiscal stance to the sovereign risk premia. Therefore, these estimates represent the effect of changes in fiscal policy conditional of no change in sovereign borrowing costs. Hebous (2011) suggests a similar range for spending shocks, based on a review of evidence including those based on international estimates, and slightly larger number for tax shocks - again considering studies estimated during periods government's had fiscal space.

Consider the lower bound on the fiscal multiplier proposed by Ramey (2011a), and the upper bound of the output effect of a 100bps decrease in borrowing costs, which is 1.8ppt extra GDP (based on the unemployment response). Under these circumstances a 1% fiscal adjustment needs to promote a 45bp reduction in the sovereign risk premium in order to be expansionary. This is a conservative estimate; under adverse economic conditions, and with a constrained/external monetary authority, the fiscal multiplier is potentially much larger.[47] Further, the output response to a reduction in yields is diminished if one uses the industrial production reading. Empirically, the relationship between government borrowing costs and fiscal fundamentals is highly non-linear (see Corsetti et al. (2013)).[48] As a result, it is difficult to assess whether a 45bp reaction is a realistic market reaction to 1% of GDP less debt under the circumstances associated with the crisis. I leave this issue to future research.

A third policy-related issue is the difference between the actual and counterfactual bond yield. One possible interpretation for why sovereign borrowing costs diverge from local economic conditions is self-fulfilling expectations of default. An alternative is that it reflects convertibility premia due to concerns over the future of the single currency. However, either interpretation can be used to justify interventions and official sector financing. The size of the systemic premia provide evidence that sovereign borrowing costs had divorced themselves from a level justified by the underlying local macroecomic situation. It is also telling that the decline in the premia coincides, particularly in Italy and Spain, with a period of ECB action over the summer of 2012 culminating in the announcement of Outright Monetary Transactions in September. This evidence supports the ECB's intervention and suggests it was effective in bringing yields back towards a more neutral setting thereby soothing the crisis. This is not to say that all the increase in sovereign borrowing costs seen during the crisis period were not due to macroeconomic conditions, as is clearly evidenced by the increase in the counterfactual yield.

5.5 Comparison with the literature

Neri and Ropele (2013) represents the closest paper in terms of approach and topic to this work. Qualitatively

---

[47]See Christiano et al. (2011) and Rendahl (2012) for a theoretical discussion. Auerbach and Gorodnichenko (2012) offer some empirical evidence.

[48]Bi (2012) offers some theoretical interpretation for this relationship.

their main results agree with those described herein: increase in sovereign risk premia are bad for industrial production growth and unemployment and reduce credit to the private sector (although, they do not shed light on the fiscal response or capital flows). The exact size of the responses they report are not directly comparable to those seen here; for example, in their benchmark case they consider a 350bp increase in the Greek yield. However, one can attempt to back out the responses: for example, they report their benchmark shock raised Italian borrowing costs by 60bp and results in a 1.2ppt decline in industrial production growth which is consistent with the results above. Other point estimates from their results differ, however; for example, Spanish industrial production responds more strongly, Irish less so. Their estimates of the unemployment response is also much larger in all the countries in question. Uribe and Yue (2006) carry out an investigation of the business cycle impact of innovations to country risk premia in emerging markets. They document the negative output effects of such shocks and show an improvement in the trade balance. They also find that such innovations explain a fair portion of business cycle dynamics in emerging economies, that is, around 12% - this is smaller than here but their sample contains non-crisis periods. The results presented in this paper also chime with Levy Yeyati and Panizza (2011), where fear of default can be very costly for an economy, and not just default itself.

In terms of the counterfactual systemic premia, the results are not directly comparable to the systemic component calculated in Ang and Longstaff (2013), as the counterfactuals are obviously relative to a baseline level of systemic risk. However, the fact that 50% of the short-term variance in yields is explained by the shock is consistent with Ang and Longstaff (2013). The systemic premia can also be compared with the large number of papers (see the literature review) assessing the extent to which the movements in bond yields in the Euro Area during the crisis was explained by local macroeconomic fundamentals. The answer from this paper is around half was and half was not. The remainder of the literature contains a variety of views.

6   Sensitivity analysis

This section presents two sensitivity checks for the empirical benchmark empirical results delineated in the previous section; first, how different assumptions over the construction of the proxy can affect the results is explored; and second, a placebo is study carried out to show that the results are not as due to of simply oversampling of the bond yield at a high frequency. For the sake of brevity the results of these analyses are presented only as impulse responses for the mean country model; 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. Additional robustness checks regarding the specification are presented in the online appendix.

The sensitivity checks do not qualitatively alter the message presented above. However, quantitatively the benchmark specification is an outlier with respect to the unemployment response. While the alternative specifications concur with the pattern of unemployment in response to the shock, other specifications, as presented in figures 9 and 10 for example, have a peak response of around 0.6ppt when compared to 0.9ppt in the benchmark case. This is in response to a systemic shock that raises the bond yield by 100bp. Correspondingly, these alternative specifications suggest that around 30% of the forecast error of unemployment is explained by systemic shocks. Interestingly, an unemployment increase of 0.6ppt is actually consistent with the 2.0ppt fall of industrial production growth (which is robust to alternative specifications) when one places it in GDP terms. Therefore, the benchmark may be overstating the unemployment response somewhat.

6.1   Alternative proxies

Construction of the proxy in section 3.3 involved several assumptions that should be tested for robustness. The first alternative proxy is constructed in one hour windows rather than 20 minute windows. This alternative

gives markets longer to react to events at 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 only uses events related to Greece. Greece is a small country from a Euro area perspective; thus real and financial linkages between it and the countries in the sample should play less of a role in determining bond market reactions. The fourth alternative proxy excludes all events that happen outside of trade hours.

The fifth alternative proxy drops all events involving foreign agents intervening in a local country; e.g. the Troika agreeing to bailout a particuluar country. Including events of this nature could be considered questionable given exogeneity assumptions underpinning the identification strategy. While it is plausible to argue that the market's reaction to domestic news in Greece, for example, is not caused 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 this category and rerun the model.

The last alternative proxy just looks at events that happen in the first week of the month. This is designed as a robustness check against the assumption that market reactions are persistant over the course of the month - a proxy based on events early in the month would be ineffective if reactions decayed in a meaningful fashion. Using early events also serves to strengthen the identifying assumptions further. If the market reaction to a foreign event is partly a function of preceding local structural economic shocks, by looking just at events early in the month there is less opportunity for foreign agents to react to the local shocks that happen that month and subsequently act upon them. This is a similar line of reasoning to what one would employ with regard to 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 the reduced form model is held constant at the benchmark.

Figure 13 presents 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, nearly all the median alternative impulse responses are within confidence set of the benchmark.

6.2   Placebo study

Given that there are approximately 400 events included in the proxy, a valid concern is that the approach taken here is simply sampling the actual changes in the yield. If this were the case, the proxy would simply be a noisy measure of the overall change in the month rather than picking up any particular shock. This is 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 recreate the benchmark proxy in an identical fashion using the same event time

but one trading day previously.[49] This retains 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, is made.

Figure 14 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 is completely 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 indicates the data collection procedure is not equivalent to just generating a noisy measure of the bond yield at time t; there is real information contained within the benchmark proxy.

The placebo study also supports the assumption that the measurement error associated with the proxy, v_{dt} in equation 6, is uncorrelated with other structural shocks in the economy. The placebo is an aggregation of randomly sampled high frequency bond market reactions. Conceptually, this can be written as an aggregation of censored observations on v_{dt}. The lack of meaningful responses produced by the placebo suggests its correlation with *any* structural shock is weak.

7 Conclusion

This paper provides evidence of the macroeconomic implications of sovereign risk premia by using narrative methods to identify fluctuations in bond yields that appear separate from economic conditions. This addresses an identification problem in the literature: discriminating between changes in the riskiness of the sovereign that is itself a function of macroeconomic conditions and the macroeconomic implications of fluctuations in sovereign risk. The high frequency, narrative identification strategy relies upon market reactions to foreign events during the Euro crisis. The transmission of these events provides variation in yields that have a degree of separation from local economic conditions; I show that the market reaction to these events does not seem to be explained by other local economic or monetary shocks.

I show that changes in sovereign yields driven by foreign events were a critical driver of macroeconomic dynamics in crisis-hit Euro Area countries over the crisis period. The identified systemic shocks explain 45% of the forecast error variance in the unemployment rate and 35% of variance in a composite measure of the private cost of finance. I also find that systemic shocks explain a substantial proportion of the variation in overall borrowing costs faced by crisis-hit countries, explaining 50% of the forecast error variance and 40-60% of the trough to peak move in bond yields in crisis-hit countries.

How generalisable these results are needs to be assessed. An analysis such as this is exposed to the Lucas critique. The results would be different if policymakers behaved differently. A more important caveat is that the Euro Area presents a rather specific circumstance. It is a monetary union where local sovereign's issue

---

[49]A placebo study based on a random set of time windows leads to much the same results.

debt in a currency that they do not control. This has several implications. Euro Area sovereigns may be more vulnerable to belief-driven crises than those with monetary sovereignty. The loss of access to the printing press may also limit the inflationary consequences of a crisis due to controls over monetary financing. However, external central bank finance is still available via the ECB. Whether the results presented here would be replicated if a similar analysis was conducted in countries such as the US, assuming an appropriate proxy could be found, is difficult to tell. A better comparator may be countries that have debt in an external currency such as in many emerging markets. In that regard, an interesting extension of the work here is to apply the methodology to the Asian crisis using the transmission of events that occurred in that period.

Another interesting angle for future research is whether these results can shed light more on the appropriateness of the micro-foundations of channels by which sovereign risk feeds through into the real economy. This paper highlights several mechanisms and it would be interesting to see if a medium scale DSGE model allowing for exogenous variation in sovereign default risk can replicate the impulse responses. One challenge that needs to be overcome is recognising that the costs of high sovereign risk premia and default are not necessarily equivalent. This paper documents the experience of four countries who kept debtors whole but suffered substantial output losses nonetheless. Any new theoretical framework needs to take this fact into consideration.

The methodology also presents a couple of avenues for exploration in future research. This paper is the first to adapt the proxy SVAR approach into a Bayesian setup. In this context, the Bayesian approach was useful in that it allowed for panel assumptions. However, the sampler used here could be adapted to allow for time-varying coefficients or stochastic volatility. A second avenue for potential future research is the general use of the proxy SVAR methodology combined with financial market reactions to identify shocks of a financial nature. The identification of shocks of this type is challenging as there are no valid short run restrictions that can be imposed; market prices cannot be assumed to be contemporaneously independent. One could also adapt the strategy to the high frequency identification of monetary shocks, or to oil shocks using the high frequency reaction of oil prices.

The dataset used in this paper has further applications. Although this paper is silent regarding the explanation behind observed market reactions, the narrative dataset can also be used to explore the determinants of the transmission of sovereign risk between countries at different times. For example, one could investigate if certain countries were of systemic importance at different times, in addition to which sorts of events the markets were sensitive to. Potentially, one could also use the events to test forms of market efficiency.

References

Acharya, V. V., Drechsler, I., Schnabl, P., Dec. 2011. A pyrrhic victory? bank bailouts and sovereign credit risk. CEPR Discussion Papers 8679, C.E.P.R. Discussion Papers.

Afonso, A., Furceri, D., Gomes, P., 2012. Sovereign credit ratings and financial markets linkages: Application to european data. Journal of International Money and Finance 31 (3), 606–638.

Aguiar M, M Amador, E. F., Gopinath, G., 2012. Crisis and commitment: Inflation credibility and the vulnerability to sovereign debt crises, mimeo.

Aizenman, J., Hutchison, M., Jinjarak, Y., 2013. What is the risk of european sovereign debt defaults? fiscal space, cds spreads and market pricing of risk. Journal of International Money and Finance 34 (C), 37–59.

Amisano, G., Giannini, C., 1997. Topics in Structural VAR Econometrics. Springer.

Ang, A., Longstaff, F. A., 2013. Systemic sovereign credit risk: Lessons from the u.s. and europe. Journal of Monetary Economics 60 (5), 493–510.

Arellano, C., Bai, Y., October 2013. Linkages across sovereign debt markets. Working Paper 19548, National Bureau of Economic Research.

Arezki, R., Candelon, B., Sy, A., 2011. Sovereign rating news and financial markets spillovers: Evidence from the european debt crisis. CESifo Working Paper Series 3411, CESifo Group Munich.

Arteta, C., Hale, G., January 2008. Sovereign debt crises and credit to the private sector. Journal of International Economics 74 (1), 53–69.

Attinasi, M.-G., Checherita-Westphal, C., Nickel, C., Dec 2009. What explains the surge in euro area sovereign spreads during the financial crisis of 2007-09? Working Paper Series 1131, European Central Bank.

Auerbach, A. J., Gorodnichenko, Y., October 2012. Fiscal multipliers in recession and expansion. In: Fiscal Policy after the Financial Crisis. NBER Chapters. National Bureau of Economic Research, Inc, pp. 63–98.

Ball, L. M., Leigh, D., Loungani, P., Jan. 2013. Okun's law: Fit at fifty? NBER Working Papers 18668, National Bureau of Economic Research, Inc.

Beetsma, R., Giuliodori, M., de Jong, F., Widijanto, D., 2013. Spread the news: How the crisis affected the impact of news on the sovereign bond markets. Journal of International Money and Finance 34, 83–101.

Bi, H., 2012. Sovereign default risk premia, fiscal limits, and fiscal policy. European Economic Review 56 (3), 389–410.

Bi, H., Traum, N., 2013. Estimating fiscal limits: The case of greece. Journal of Applied Econometrics, forthcoming.

Borensztein, E., Panizza, U., November 2009. The costs of sovereign default. IMF Staff Papers 56 (4), 683–741.

Brutti, F., Saure, P. U., March 2013. Transmission of sovereign risk in the euro crisis, mimeo.

Burnside, C., Eichenbaum, M., Fisher, J. D. M., March 2004. Fiscal shocks and their consequences. Journal of Economic Theory 115 (1), 89–117.

Calvo, G. A., September 1988. Servicing the public debt: The role of expectations. American Economic Review 78 (4), 647–61.

Canova, F., Ciccarelli, M., January 2013. Panel vector autoregressive models a survey. Tech. Rep. 1507, European Central Bank Working Paper Series.

Christiano, L., Eichenbaum, M., Rebelo, S., 2011. When is the government spending multiplier large? Journal of Political Economy 119 (1), 78 − 121.

Cloyne, J., June 2013. Discretionary tax changes and the macroeconomy: New narrative evidence from the united kingdom. American Economic Review 103 (4), 1507–28.

Cochrane, J. H., Piazzesi, M., May 2002. The fed and interest rates - a high-frequency identification. American Economic Review 92 (2), 90–95.

Cole, H. L., Kehoe, T. J., January 2000. Self-fulfilling debt crises. Review of Economic Studies 67 (1), 91–116.

Cooper, R., September 2012. Fragile debt and the credible sharing of strategic uncertainty. Working Paper 18377, National Bureau of Economic Research.

Corsetti, G., Dedola, L., April 2013. The mystery of the printing press: Self-fulfilling debt crises and monetary sovereignty, mimeo.

Corsetti, G., Guimaraes, B., Roubini, N., April 2006. International lending of last resort and moral hazard: A model of imf's catalytic finance. Journal of Monetary Economics 53 (3), 441–471.

Corsetti, G., Kuester, K., Meier, A., Müller, G. J., 02 2013. Sovereign risk, fiscal policy, and macroeconomic stability. Economic Journal 0, F99–F132.

De Grauwe, P., 2011. The governance of a fragile eurozone. Tech. rep., CEPS working documents.

De Grauwe, P., Ji, Y., 2013. Self-fulfilling crises in the eurozone: An empirical test. Journal of International Money and Finance 34 (C), 15–36.

De Paoli, B., Hoggarth, G., Saporta, V., Feb. 2009. Output costs of sovereign crises: some empirical estimates. Bank of England working papers 362, Bank of England.

De Santis, R. A., Feb. 2012. 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.

Durbin, E., Ng, D., June 2005. The sovereign ceiling and emerging market corporate bond spreads. Journal of International Money and Finance 24 (4), 631–649.

Faust, J., Swanson, E. T., Wright, J. H., September 2004. Identifying vars based on high frequency futures data. Journal of Monetary Economics 51 (6), 1107–1131.

Favero, C., Giavazzi, F., May 2012. Measuring tax multipliers: The narrative method in fiscal vars. American Economic Journal: Economic Policy 4 (2), 69–94.

Furceri, D., Zdzienicka, A., 2012. How costly are debt crises? Journal of International Money and Finance 31 (4), 726–742.

Gade, T., Salines, M., Glöckler, G., Strodthoff, S., Jul. 2013. Loose lips sinking markets?: the impact of political communication on sovereign bond spreads. Working Paper Series 150, European Central Bank.

Gelman, A., 2006. Prior distributions for variance parameters in hierarchical models. Bayesian Analysis 1 (3), 515–533.

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

Giannone, D., Lenza, M., Primiceri, G. E., Nov 2012. Prior selection for vector autoregressions. Working Paper Series 1494, European Central Bank.

Giordano, R., Pericoli, M., Tommasino, P., 06 2013. Pure or wake-up-call contagion? another look at the emu sovereign debt crisis. International Finance 16 (2), 131–160.

Gorodnichenko, Y., Weber, M., Feb. 2013. Are sticky prices costly? evidence from the stock market. NBER Working Papers 18860, National Bureau of Economic Research, Inc.

Gürkaynak, R. S., Wright, J., Mar. 2013. Identification and inference using event studies. CEPR Discussion Papers 9388, C.E.P.R. Discussion Papers.

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

Hewitt, G., 2013. The Lost Continent. Hodder & Stoughton.

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

Leeper, E. M., Walker, T. B., Yang, S.-C. S., 05 2013. Fiscal foresight and information flows. Econometrica 81 (3), 1115–1145.

Levy Yeyati, E., Panizza, U., January 2011. The elusive costs of sovereign defaults. Journal of Development Economics 94 (1), 95–105.

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

Lorenzoni, G., Werning, I., 2013. Slow moving debt crises. NBER Working Papers 19228, National Bureau of Economic Research, Inc.

Manasse, P., Zavalloni, L., Jan 2013. Sovereign contagion in europe: Evidence from the cds market. Working paper series, The Rimini Centre for Economic Analysis.

Mendoza, E. G., Yue, V. Z., 2012. A general equilibrium model of sovereign default and business cycles. The Quarterly Journal of Economics 127 (2), 889–946.

Mertens, K., Ravn, M. O., 2013a. A reconciliation of svar and narrative estimates of tax multipliers. Journal of Monetary Economics, forthcoming.

Mertens, K., Ravn, M. O., June 2013b. The dynamic effects of personal and corporate income tax changes in the united states. American Economic Review 103 (4), 1212–47.

Mitchell, J., Smith, R. J., Weale, M. R., Wright, S., Salazar, E. L., 02 2005. An indicator of monthly gdp and an early estimate of quarterly gdp growth. Economic Journal 115 (501), F108–F129.

Nakamura, E., Steinsson, J., Jul. 2013. High frequency identification of monetary non-neutrality. NBER Working Papers 19260, National Bureau of Economic Research, Inc.

Neri, S., Ropele, T., 2013. The macroeconomic effects of the sovereign debt crisis in the euro area, mimeo.

Nordvig, J., Firoozye, N., 2012. Planning for an orderly break-up of the european monetary union. Submission to the wolfson economics prize 2012, Nomura.

Pescatori, A., Sy, A. N. R., June 2007. Are debt crises adequately defined? IMF Staff Papers 54 (2), 306–337.

Ramey, V. A., September 2011a. Can government purchases stimulate the economy? Journal of Economic Literature 49 (3), 673–85.

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

Ramey, V. A., Shapiro, M. D., June 1998. Costly capital reallocation and the effects of government spending. Carnegie-Rochester Conference Series on Public Policy 48 (1), 145–194.

Reinhart, C. M., Rogoff, K. S., May 2009. The aftermath of financial crises. American Economic Review 99 (2), 466–72.

Reinhart, C. M., Rogoff, K. S., August 2011. From financial crash to debt crisis. American Economic Review 101 (5), 1676–1706.

Rendahl, P., Feb 2012. Fiscal policy in an unemployment crisis. Cambridge Working Papers in Economics 1211, Faculty of Economics, University of Cambridge.

Romer, C. D., Romer, D. H., December 1989. Does monetary policy matter? a new test in the spirit of friedman and schwartz. In: NBER Macroeconomics Annual 1989, Volume 4. NBER Chapters. National Bureau of Economic Research, Inc, pp. 121–184.

Romer, C. D., Romer, D. H., September 2004. A new measure of monetary shocks: Derivation and implications. American Economic Review 94 (4), 1055–1084.

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

Rose, A. K., June 2005. One reason countries pay their debts: renegotiation and international trade. Journal of Development Economics 77 (1), 189–206.

Sims, E. R., Jun. 2012. News, non-invertibility, and structural vars. Working Papers 013, University of Notre Dame, Department of Economics.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., Van Der Linde, A., 2002. Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 64 (4), 583–639.

Steinkamp, S., Westermann, F., 2012. On creditor seniority and sovereign bond prices in europe. CESifo Working Paper Series 3944, CESifo Group Munich.

Stock, J. H., Watson, M. W., May 2012. Disentangling the channels of the 2007-2009 recession. NBER Working Papers 18094, National Bureau of Economic Research, Inc.

Tirole, J., Jun. 2012. Country solidarity, private sector involvement and the contagion of sovereign crises. TSE Working Papers 12-375, Toulouse School of Economics (TSE).

Tomz, M., Wright, M. L. J., 04-05 2007. Do countries default in "bad times" ? Journal of the European Economic Association 5 (2-3), 352–360.

Uribe, M., Yue, V. Z., 2006. Country spreads and emerging countries: Who drives whom? Journal of International Economics 69 (1), 6 − 36.

A   The conditional distribution of the proxy.

This appendix details the distribution of the proxy variable under random censoring and derives the linear relationship with the reduce form errors. Recall the assumed data generating process:

$$m_{dt} = \vartheta_{dt}(\psi \varepsilon_{dt} + v_{dt}) \tag{9}$$

Under random censoring, the density of m_{dt}, given the true structural shock, ε_{dt}, can be expressed as:

$$P(m_d|\psi, \varepsilon_d, p) = \left((2\pi\sigma_v^2)^{-\frac{1}{2}} exp\left\{-\frac{1}{2}\left(\frac{m_d - \psi\varepsilon_d}{\sigma_v}\right)^2\right\} p\right)^{1-\mathbf{1}_0(m_d)} (1-p)^{\mathbf{1}_0(m_d)} \tag{10}$$

Where $\mathbf{1}_0(m_d)$ is an indicator variable which returns 1 if $m_d = 0$ and zero otherwise and $\mathbb{E}(\vartheta_{dt}|\varepsilon_{dt}) = \mathbb{E}(\vartheta_{dt}) = p$. Dropping the time t subscripts, this density implies the moment generating function is given by:

$$M_{m_d}(g) = \int_{-\infty}^{\infty} exp\{gm_d\}\left(p(2\pi\sigma_v^2)^{-\frac{1}{2}} exp\left\{-\frac{1}{2}\left(\frac{m_d - \psi\varepsilon_d}{\sigma_v}\right)^2\right\}\right)^{1-\mathbf{1}_0(m_d)} (1-p)^{\mathbf{1}_0(m_d)} dm_d \tag{11}$$

This integral can be rewritten as:

$$M_{m_d}(g) = \int_{-\infty}^{\infty} \left(p(2\pi\sigma_v^2)^{-\frac{1}{2}} exp\left\{-\frac{1}{2}\left(\frac{m_d - \psi\varepsilon_d}{\sigma_v}\right)^2\right\} exp\{gm_d\} + \delta(m_d)(1-p)exp\{gm_d\}\right) dm_d$$

Where $\delta(m_d)$ is the Dirac delta function. Solving the integral yields:

$$M_{m_d}(g) = (1-p) + pexp\left\{g\psi\varepsilon_d + \frac{g^2}{2}\sigma_v^2\right\} \tag{12}$$

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

$$M_m(g) = \prod_{d=1}^{M}((1-p) + p.exp\left\{g\psi\varepsilon_d + \frac{g^2}{2}\sigma_v^2\right\}) \tag{13}$$

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

$$\frac{\partial M_m(g)}{\partial g} = \sum_{d=1}^{M}(p(\psi\varepsilon_d + g\sigma_v^2)x_d \prod_{s\neq d}^{M}((1-p) + px_s) \tag{14}$$

which implies the first moment is given by:

$$E(m|\varepsilon_d) = p\psi\sum_{d=1}^{M}\varepsilon_d = p\psi a_1 u = E(m|u) \tag{15}$$

To calculate the second moment, the second differential is:

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

hence:

$$\mathbb{E}(m^2|\varepsilon_d) = p\sum_{d=1}^{M}\left[(\sigma_v^2 + (\psi^2\varepsilon_d^2)) + p^2\psi^2\varepsilon_d \sum_{s\neq d}^{M}\{\varepsilon_s\}\right]$$

Since the individual daily observations are unobservable, what we care about is $\mathbb{E}(m^2|u)$, which, by the law of iterated expectations, is simply:

$$\mathbb{E}(m^2|u) = pM\sigma_v^2 + p^2\psi^2 \sum_{d=1}^{M} \left[\mathbb{E}(\varepsilon_d^2|u) + \sum_{s \neq d}^{M} \mathbb{E}(\varepsilon_s\varepsilon_d|u) \right]$$

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

B   Data and Sources

B.1   Constructing the proxy variable

The proxy is designed as an aggregation of the local bond market reactions to foreign events such as those described in section 3. Therefore, in order to construct it 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 is conducted from July 2009 to March 2013.

The Euro crisis is well-documented and comprises 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 for the information flow is, by definition, media outlets, which 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 Euro Area bond market reaction is the ultimate variable of interest, restricting attention to financial news sources is appropriate.

I use the financial news sources *Bloomberg* and *EuroIntelligence* both of which compile a daily news briefing for European economic news, with the former released in the afternoon and the latter in the morning. The *Bloomberg* news coverage for the Euro Area 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 news based series. 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, because the country specific event must be of sufficient international interest to make the briefing. As described, this is not to say the market reacts strongly to every news story within the summaries.

While there is a large overlap in terms of events between these two sources, the timing of the news summary turns out to be of importance. Twenty-four hours represents a long period of time in the context of 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 released the next morning. In combination these two briefings provide half-day snapshots of the key news stories that can be assumed to be affecting Euro Area bond markets.

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 additional sources have almost complete overlap with the two already considered.

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, either Greece, Cyprus, Portugal, Ireland, Italy or Spain. As discussed previously pan-Euro Area policy interventions are not included because the identifying assumptions are harder to justify. Experiments with political events in non-crisis countries revealed that bond markets do not react strongly to this form of news and as such these countries are omitted for the sake of parsimony.[50] Stories that relate to foreign policymakers intervening in a specific crisis country are not omitted from the benchmark specification although they are in a robustness analysis.

2. The event must be timeable in the sense that it is possible to isolate when it occurs so 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, private individuals are also not included unless that individual has an official policy-making or political capacity.

   (c) Editorial statements.

   (d) Data releases are 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.

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; 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. An untimeable story is one where it is impossible to identify an initial headline in an objective fashion. News stories for which 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 at which one can assess the market's reaction.[51] Alternatively,

---

[50] A second reason for is that Germany and France are large economies tightly integrated with the rest of the currency union. Events in these countries have implications above and beyond the bond market reaction. For example, German fiscal policy clearly affects the rest of the union via demand channels; while Greek or Portuguese fiscal policy is less relevant.

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

it could be a 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.[52]

For events that can be timed, the market reaction is considered over a 20 minute window on either side of the initial headline. This window starts slightly earlier than suggested by Gürkaynak 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.

The market reaction is defined as the change in the mid-yield to maturity on the benchmark 10 year sovereign bond. Note that this is not the country where the event occurred; thus if the local 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 so as to remove any spikes at a very high frequency; the market reaction is calculated as changes in the averaged minute by minute series. Only ticks between 07:30 and 16:30 London time are used, i.e. the time at which the London market is open. Ticks outside this interval are too infrequent to be relied upon. The market reaction to included events that occur outside of normal trading hours is calculated as the change from the previous close to 08:30 London time on the morning of the first trading day after the event. The period between 07:30am-08:30 is noisy and subject to spikes, thus for overnight reactions I record the market position at 08:30.

B.2   Properties of the proxies

Readers are referred to the online appendix, see links below, for an exhaustive list of narrative events. However, some clarity over the type of events included in the narrative series can be achieved by placing the events in loose classifications with accompanying examples:

1. *Event country political news:* 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. *External interventions:* These refer to statements by external policymakers, particularly Troika members,[53] about activities relating only to the specific event country. The various bailout agreements are obvious examples, as well as the approval of critical disbursements. Other examples included are the

---

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

[53]Only statements by international or pan-European policymakers are included.

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
	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

Event Type (%) label appears vertically along the left of the Event Type rows.

Notes: The table summarises the 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.

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 event country 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. *Event country 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. *Event country 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 announcements concerning when strikes and protests will take place. Also included are violent events that occur during a protest; for example, and with reference to 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 of events identified in each crisis country. As one would expect given the country's role as the initial focal point 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 this small country's bailout. The breakdown of events into five classifications are similar across countries with the exception of Italy, where foreign interventions are less prevalent. This is reasonable as Italy is the only country not to receive a bailout in some form over the time period.

Table 2 offers some descriptive statistics for the proxy variables in the four countries of interest. 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 events display similar statistical properties across countries which is somewhat surprising given that, 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 share is roughly in line with the relative number of events that are of Greek origin - it is not the case

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 section 3.3. 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.

that markets are reacting more strongly to Greek news on average , merely that there are more Greek events to react to.

B.3 Weblinks to the narrative 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 constraints 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 occurred, 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 found in 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.4   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 amalgamation 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 analogous to an overall interest rate on financial liabilities. The cost can be decomposed into its constituent 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 accruals 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 of 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 is 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 revealed 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 accruals 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 (in 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.

Trade Balance: The goods trade balance in millions of euros is sourced from Eurostat (*code: ext_st_27msbec*). Included traded sectors are those contained in the BEC industry classification. The data is seasonally and working day adjusted and corresponds to the balance with the rest of the world (both Euro Area and non-Euro Area trading partners). The trade balance is placed into percentage of GDP terms by dividing through by nominal GDP interpolated with a cubic spline (see primary fiscal balances).

Target 2 Balance: Data for target 2 balances are sourced from the updated dataset of Steinkamp and Westermann (2012) (available online at http://www.eurocrisismonitor.com/). The data is in millions of euros. The balance is converted into percentage of GDP terms by dividing through by nominal GDP interpolated with a cubic spline (see primary fiscal balances).

C MCMC Sampler

Define the parameter space in the model as:

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

To simplify the notation 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 rule the likelihood of the data is equal 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: $p(\mathcal{M}, Y|\Theta) = p(\mathcal{M}|Y, \Theta)p(Y|\Theta) = \prod_c p(\mathcal{M}_c|Y_c, \Theta)p(Y_c|\Theta)$.

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:

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

Where the matrix Ξ_c is a diagonal vector of unknown parameters equal to $diag\{\xi_{c1}, \ldots, \xi_{cT}\}$. With the additional prior assumption $\nu/\xi_{ci} \sim \chi^2(\nu) \, \forall i = 1, \ldots, T$, 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: $V_c = (\mathcal{M}_c - U_c \Upsilon_c)$. This gives:

$$p(\mathcal{M}_c|Y_c, \Theta) = \sigma_{c\omega}^{-T-1} \prod_t \xi_{ct}^{-1/2} exp\left[-\sum_{t=1}^T \frac{(m_{c,t} - \Upsilon_c u_{c,t})^2}{2\xi_{ct}\sigma_{c\omega}^2}\right] = \sigma_{c\omega}^{-T-1}|\Xi_c|^{-\frac{1}{2}} exp\left\{-\frac{1}{2}\left\{\sigma_{c\omega}^{-2} V_c' \Xi_c^{-1} V_c\right\}\right\}$$

The likelihood of the reduced form VAR model, $p(Y|\Theta)$, is the product of the country specific Gaussian distributions as defined in equation 3. Combining these two densities with the priors gives a joint posterior

density, $p(\mathcal{M}, Y|\Theta)p(\Theta)$, proportional to:

$$|\bar{S}|^{\frac{C\kappa-(N+1)}{2}}(\lambda_\Upsilon\lambda_\beta)^{-\frac{v+2+C}{2}}\prod_c\left(\sigma_{c\omega}^{-T-1}|\Xi_c|^{-\frac{1}{2}}|\Sigma_{c,u}|^{-\frac{T+\kappa+N+1}{2}}\right)exp\left\{\frac{s}{\lambda_\beta}+\frac{s}{\lambda_\Upsilon}\right\}\dots\quad\dots$$

$$exp\left\{-\frac{1}{2}\left(\sum_c\left\{tr\left[(U_c'U_c\Sigma_{c,u}^{-1})+\bar{S}\Sigma_{c,u}^{-1}\right]+(\beta_c-\bar{\beta})'(\lambda_\beta L_{c,\beta})^{-1}(\beta_c-\bar{\beta})+(\Upsilon_c-\bar{\Upsilon})'(\lambda_\Upsilon L_{c,\Upsilon})^{-1}(\Upsilon_c-\bar{\Upsilon})+\sigma_{c\omega}^{-2}V_c'\Xi_c^{-1}V_c-\nu tr(\Xi_c^{-1})\right\}\right)\right\}$$

This is a convienent way to express the posterior. However, it is also apparent that the VAR data and the proxy are jointly Guassian:

$$\begin{pmatrix}y_c\\\mathcal{M}_c\end{pmatrix}|\Theta\sim N\left(\begin{pmatrix}(I_N\otimes X_c)\beta_c+(I_N\otimes Z_c)\gamma_c\\0\end{pmatrix},\begin{bmatrix}(\Sigma_{c,u}\otimes I_T)&(\Sigma_{c,u}\Upsilon_c\otimes I_T)\\(\Upsilon_c'\Sigma_{c,u}\otimes I_T)&(\Upsilon_c'\Sigma_{c,u}\Upsilon_c)\otimes I_T+\sigma_{c,\omega}^2\Xi_c\end{bmatrix}\right)$$

Let:

$$\begin{bmatrix}(\Sigma_{c,u}\otimes I_T)&(\Sigma_{c,u}\Upsilon_c\otimes I_T)\\(\Upsilon_c'\Sigma_{c,u}\otimes I_T)&\sigma_{c,\omega}^2\Xi_c+(\Upsilon_c'\Sigma_{c,u}\Upsilon_c\otimes I_T)\end{bmatrix}=\Phi_c=\begin{bmatrix}\Phi_{c,11}&\Phi_{c,21}\\\Phi_{c,12}&\Phi_{c,22}\end{bmatrix}$$

For all the parameters in the model the conditional densities used in the Gibbs Sampler are in the form of classical distributions. The conditional density of the slope coefficeints is:

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

Using the joint Gaussian density of the proxy and the reduced form one can show:

$$p(\beta_c|Y,\mathcal{M},\Theta\setminus\beta_c)\propto N(D_c^{-1}d_c,D_c^{-1})\tag{16}$$

where

$$D_c=(I_N\otimes X_c)'(\Phi_{c,11}-\Phi_{c,12}\Phi_{c,22}^{-1}\Phi_{c,22})^{-1}(I_N\otimes X_c)+\lambda_\beta^{-1}L_{c,\beta}^{-1}$$

$$d_c=(I_N\otimes X_c)'(\Phi_{c,11}-\Phi_{c,12}\Phi_{c,22}^{-1}\Phi_{c,21})^{-1}(y_c-(I_N\otimes Z_c)\gamma_c-\Phi_{c,12}\Phi_{c,22}^{-1}\mathcal{M}_c)+\lambda_\beta^{-1}L_{c,\beta}^{-1}\bar{\beta}$$

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

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

This is also Gaussian:

$$p(\gamma_c|Y,\mathcal{M},\Theta\setminus\gamma_c)\propto N(F_c^{-1}f_c,F_c^{-1})$$

where

$$F_c=(I_N\otimes Z_c)'(\Phi_{c,11}-\Phi_{c,12}\Phi_{c,22}^{-1}\Phi_{c,21})^{-1}(I_N\otimes Z_c)$$

$$f_c=(I_N\otimes Z_c)'(\Phi_{c,11}-\Phi_{c,12}\Phi_{c,22}^{-1}\Phi_{c,21})^{-1}(y_c-(I_N\otimes X_c)\beta_c-\Phi_{c,12}\Phi_{c,22}^{-1}\mathcal{M}_c)$$

The conditional posterior of Σ_c is proportional to:

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

which is consistent with an inverse-Wishart distribution:

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

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

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

$$G_c = (\lambda_\beta L_{c,\beta})^{-1}$$

$$g_c = (\lambda_\beta L_{c,\beta})^{-1}\beta_c$$

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

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

which corresponds to a Wishart distribution:

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

Note that $E(\bar{S}|Y, \Theta \setminus \bar{S}) = C\kappa(\sum_c\left[\Sigma_{c,u}^{-1}\right])^{-1}$. 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, $\bar{\Sigma}$, for use in impulses etc. By setting $\bar{\Sigma} = \bar{S}/\kappa$, one obtains a matrix that is analogous to a covariance matrix and in (conditional) expectation is equivalent to the harmonic mean of the estimated country covariances.

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

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

or

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

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_β 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:

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

where:

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

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

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

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

which is consistent with an inverse-Gamma distribution:

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

The conditional posterior of ξ_{ct}, the diagonals in Ξ_c, 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:

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

$$J_c = (\lambda_\Upsilon L_{c,\Upsilon})^{-1}$$

$$j_c = (\lambda_\Upsilon L_{c,\Upsilon})^{-1}\Upsilon_c$$

Last, the posterior of λ_Υ is proportional to:

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

Where iG_2 refers to an inverted Gamma-2 distribution.

Figure 3: Proxy variable and actual changes in the bond yield, other countries.
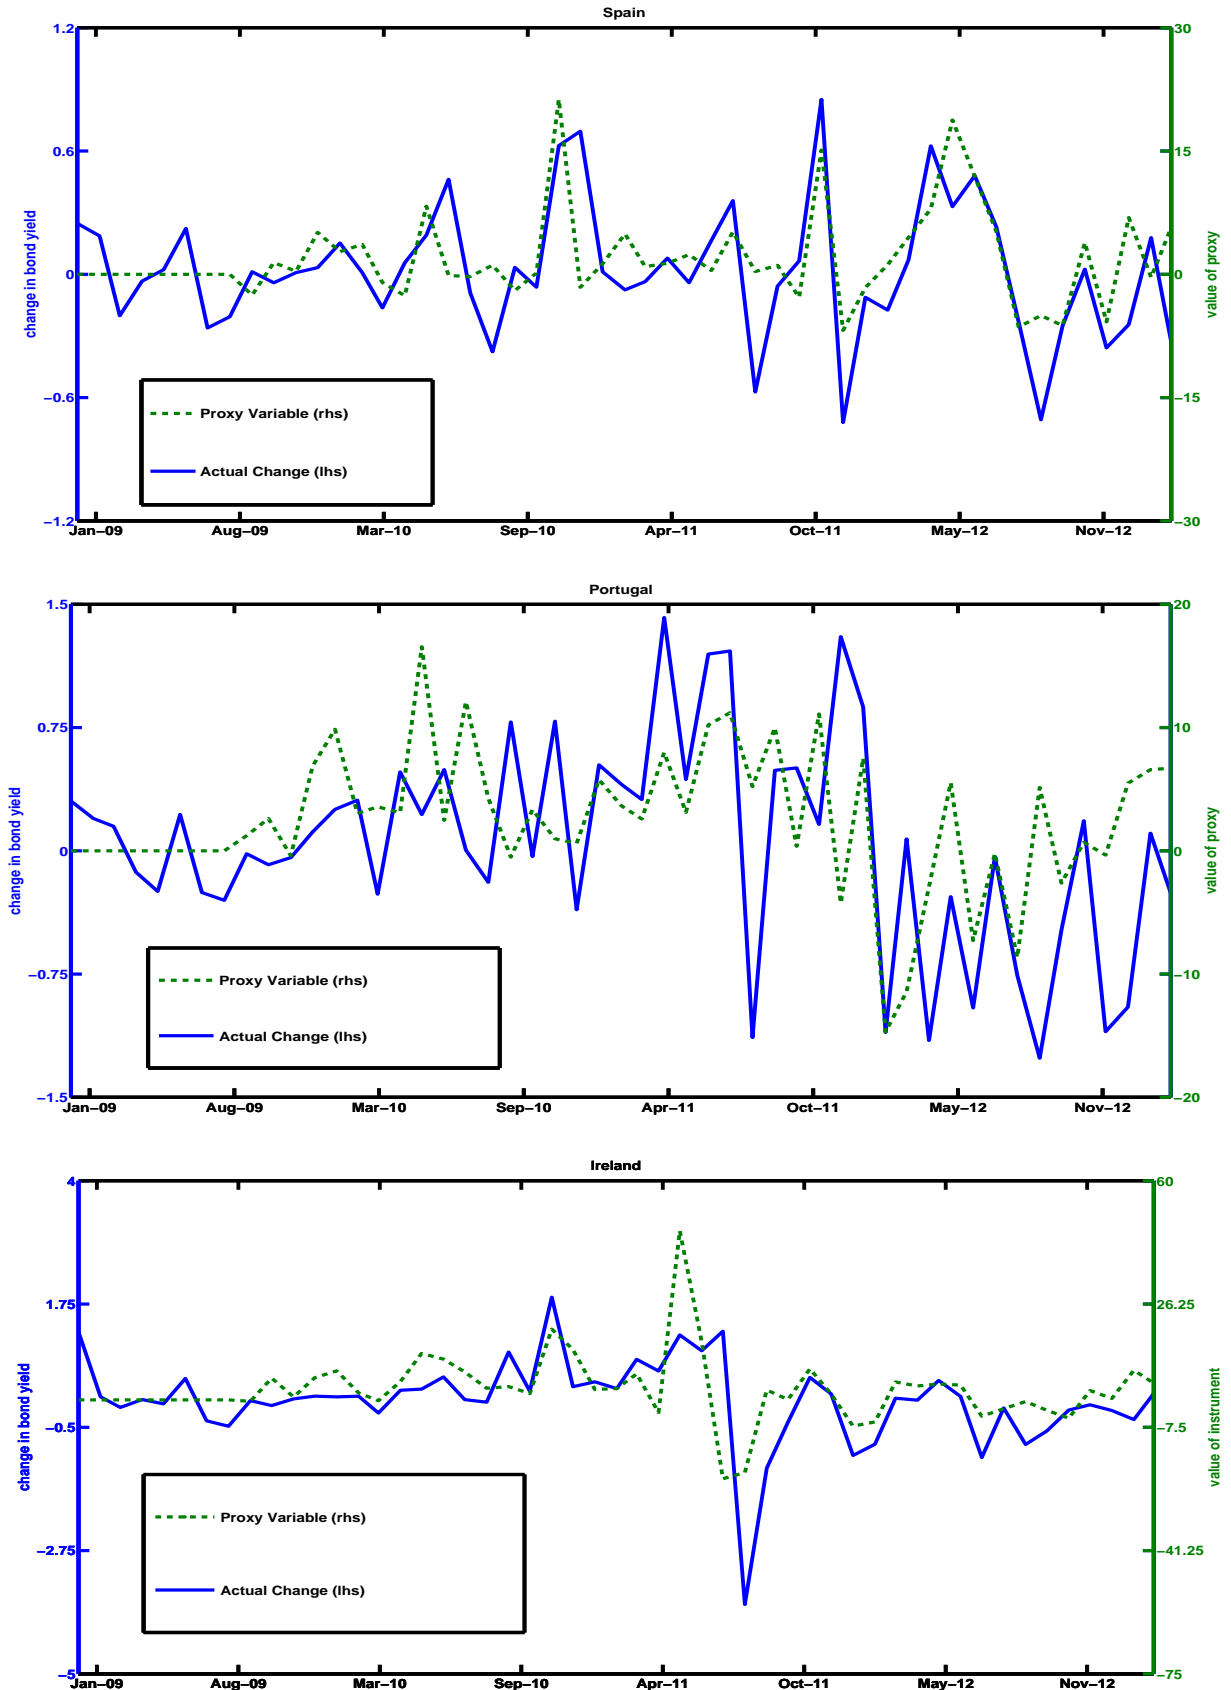
Notes: Comparison between the actual behaviour of Spanish, Portuguese and Irish bond yields and the relevant proxy variable. Y-axis denotes percentage moves in the month. Green line is the proxy variable calculated as the summed value of changes during events in that month, right hand axis, blue line is the the actual change in the average daily 10-year bond yield, left hand axis. Irish proxy uses daily changes on days of headline events from May-October 2011 due to a break in intra-day data.

Figure 4: Events ranked by their squared market reaction
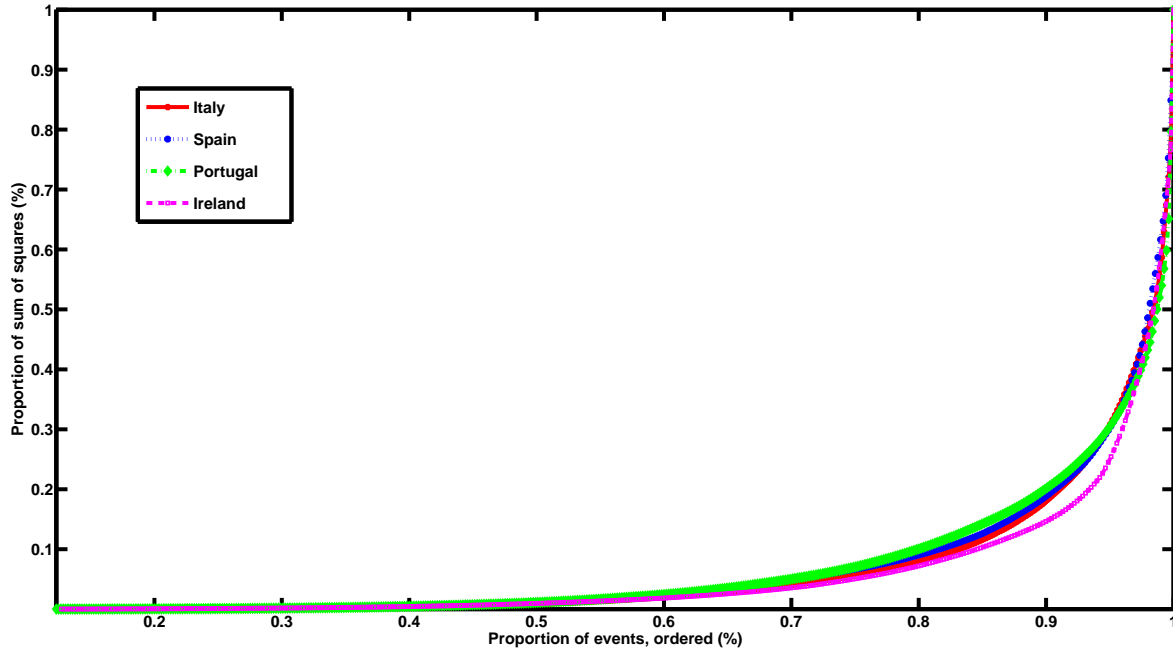
Notes: The chart illustrates the relative importance of events contained in each country's proxy in terms of market reactions. X-axis denotes the cumulative share of events in proxy ordered by size of market reaction. Y-axis denotes cumulative share total sum of squared market reactions. Sample is events from 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 5: Simulated ω compared with classical distributions
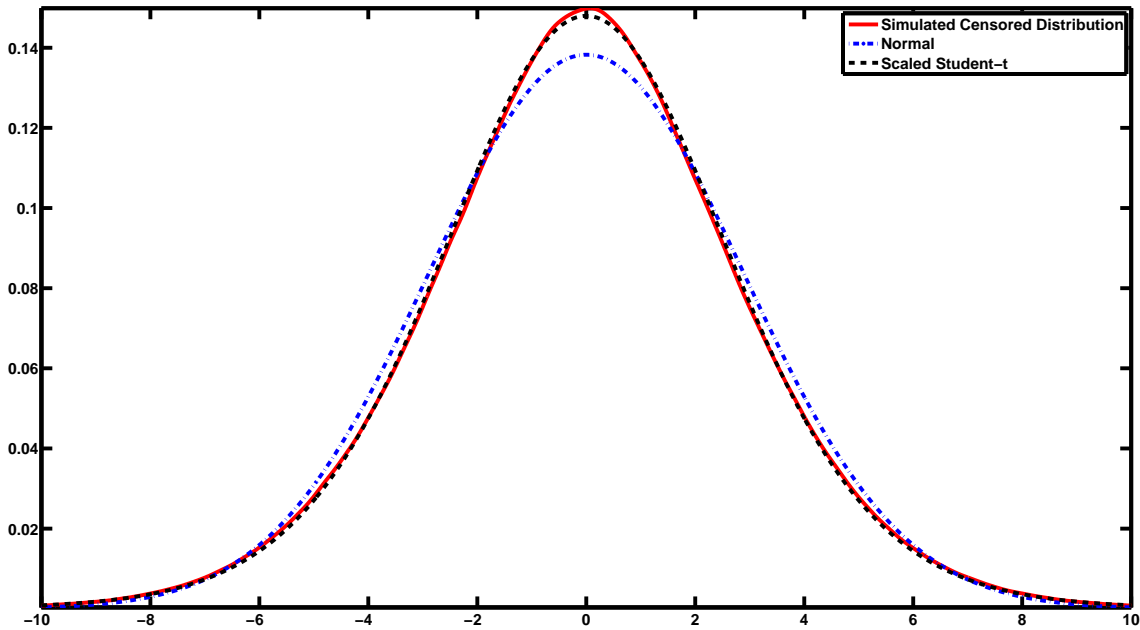
Notes: Censored distribution the kernel density of 1 million draws from the random censoring model with $p = 0.15$, $M = 30$ and $\psi = \sigma_v^2 = 1$. The normal distribution is selected to match the first two moments, the scaled t the first four moments.

51

Figure 6: Mean Country Impulse Responses to a Systemic Shock (Benchmark Specification)
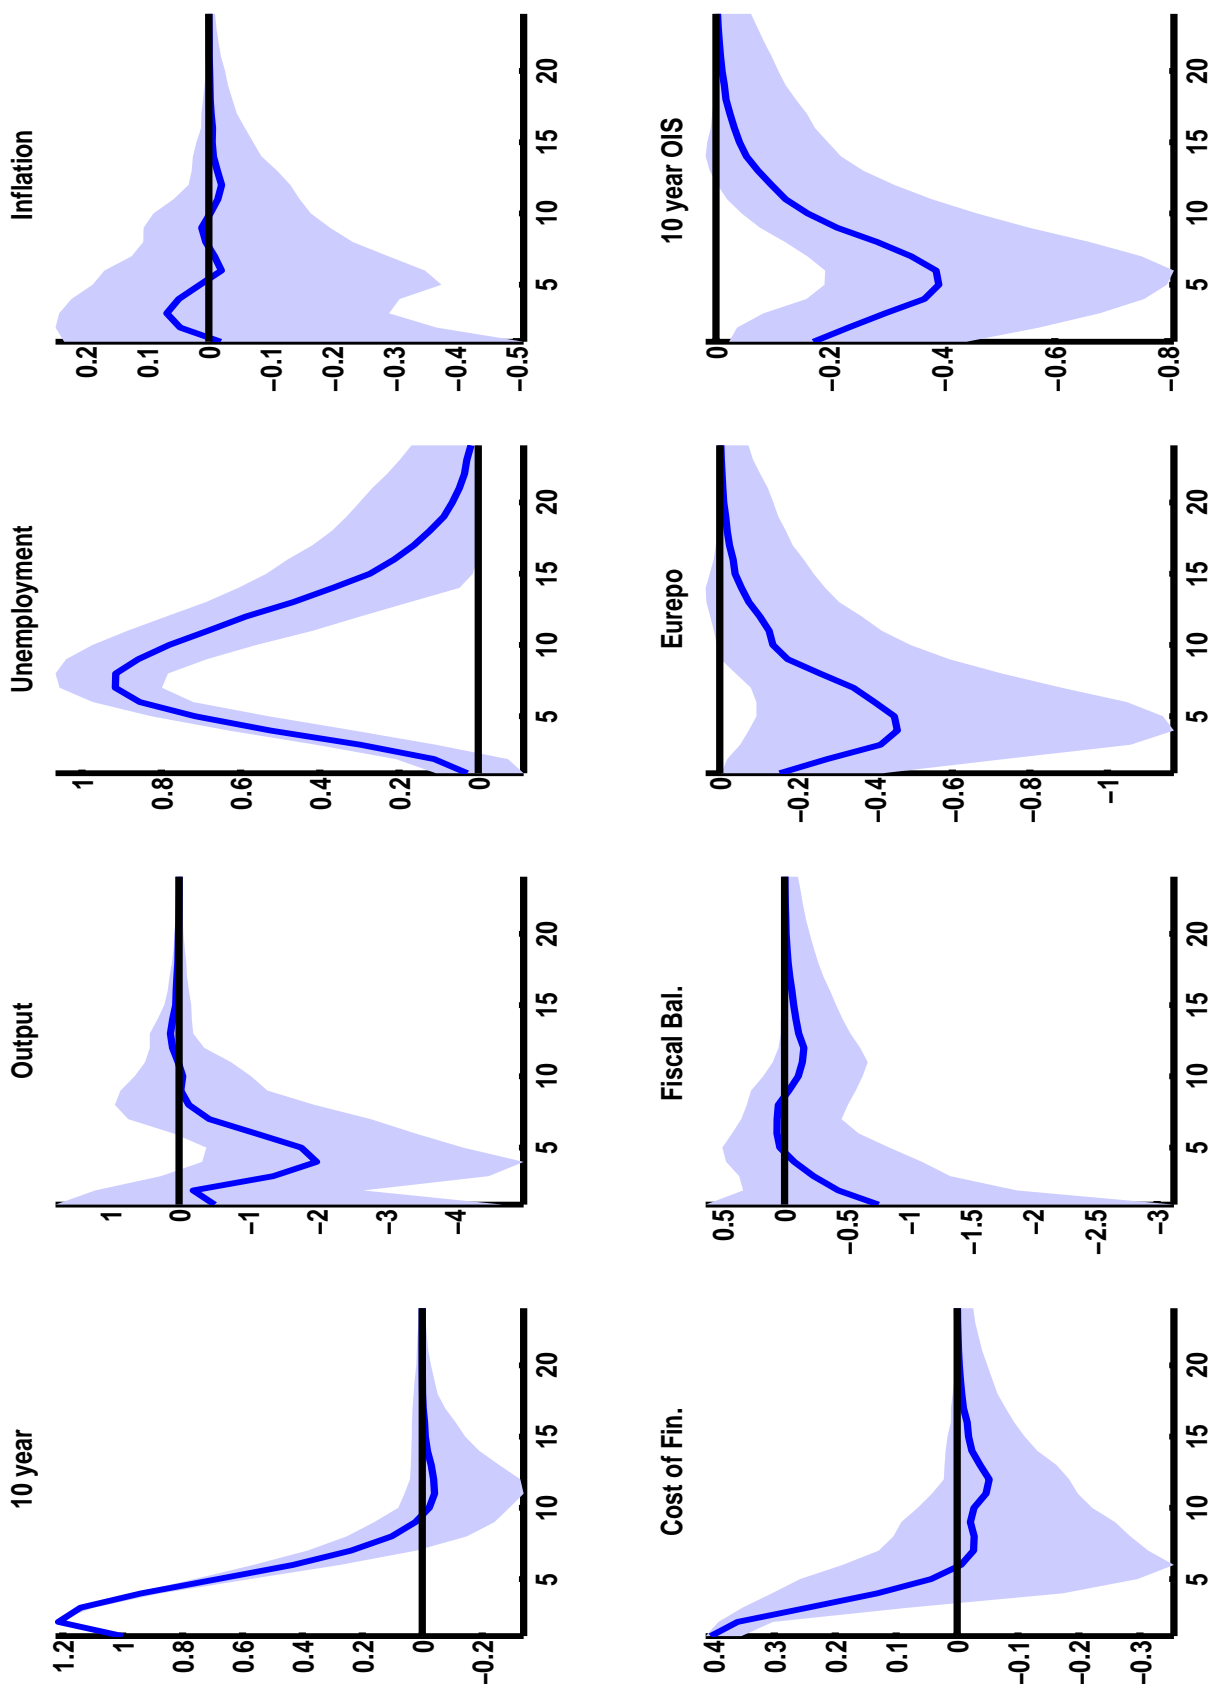

Notes: Impulse responses to a systemic shock scaled to be consistent with a 100bp increase in sovereign yield on impact and computed over 24 months. Y-axis is percentage points in all cases. Mean country model refers to impulse responses estimated using $\bar{\beta}$, $\bar{\Upsilon}$ and $\bar{\Sigma}$. 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 sovereign bond, output refers to industrial production, Cost of fin. to the private cost of finance. For exact data definitions see main text and the data appendix.

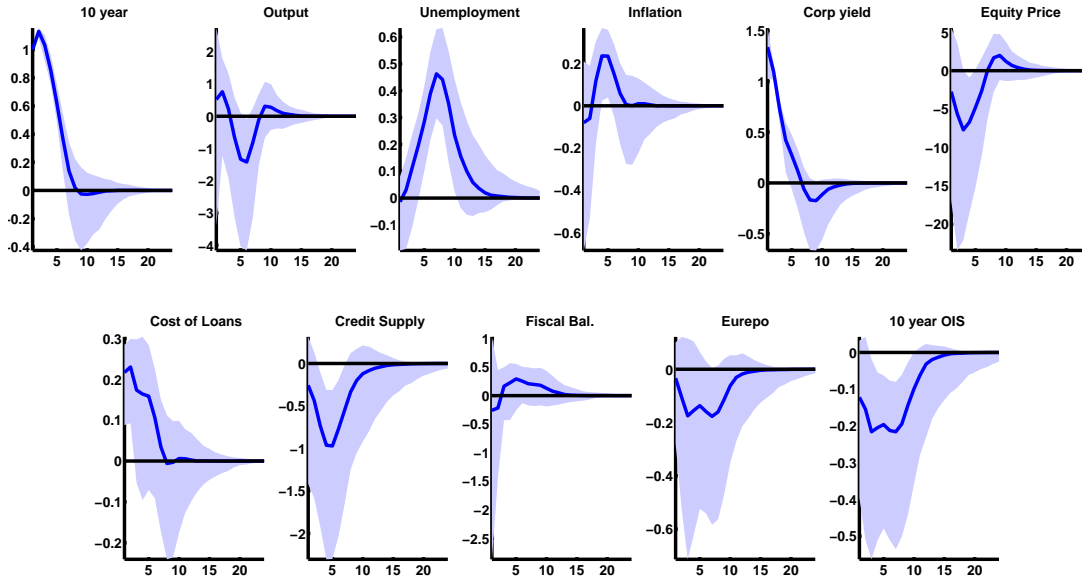Figure 7: Country Specific Impulse Responses to a Systemic Shock
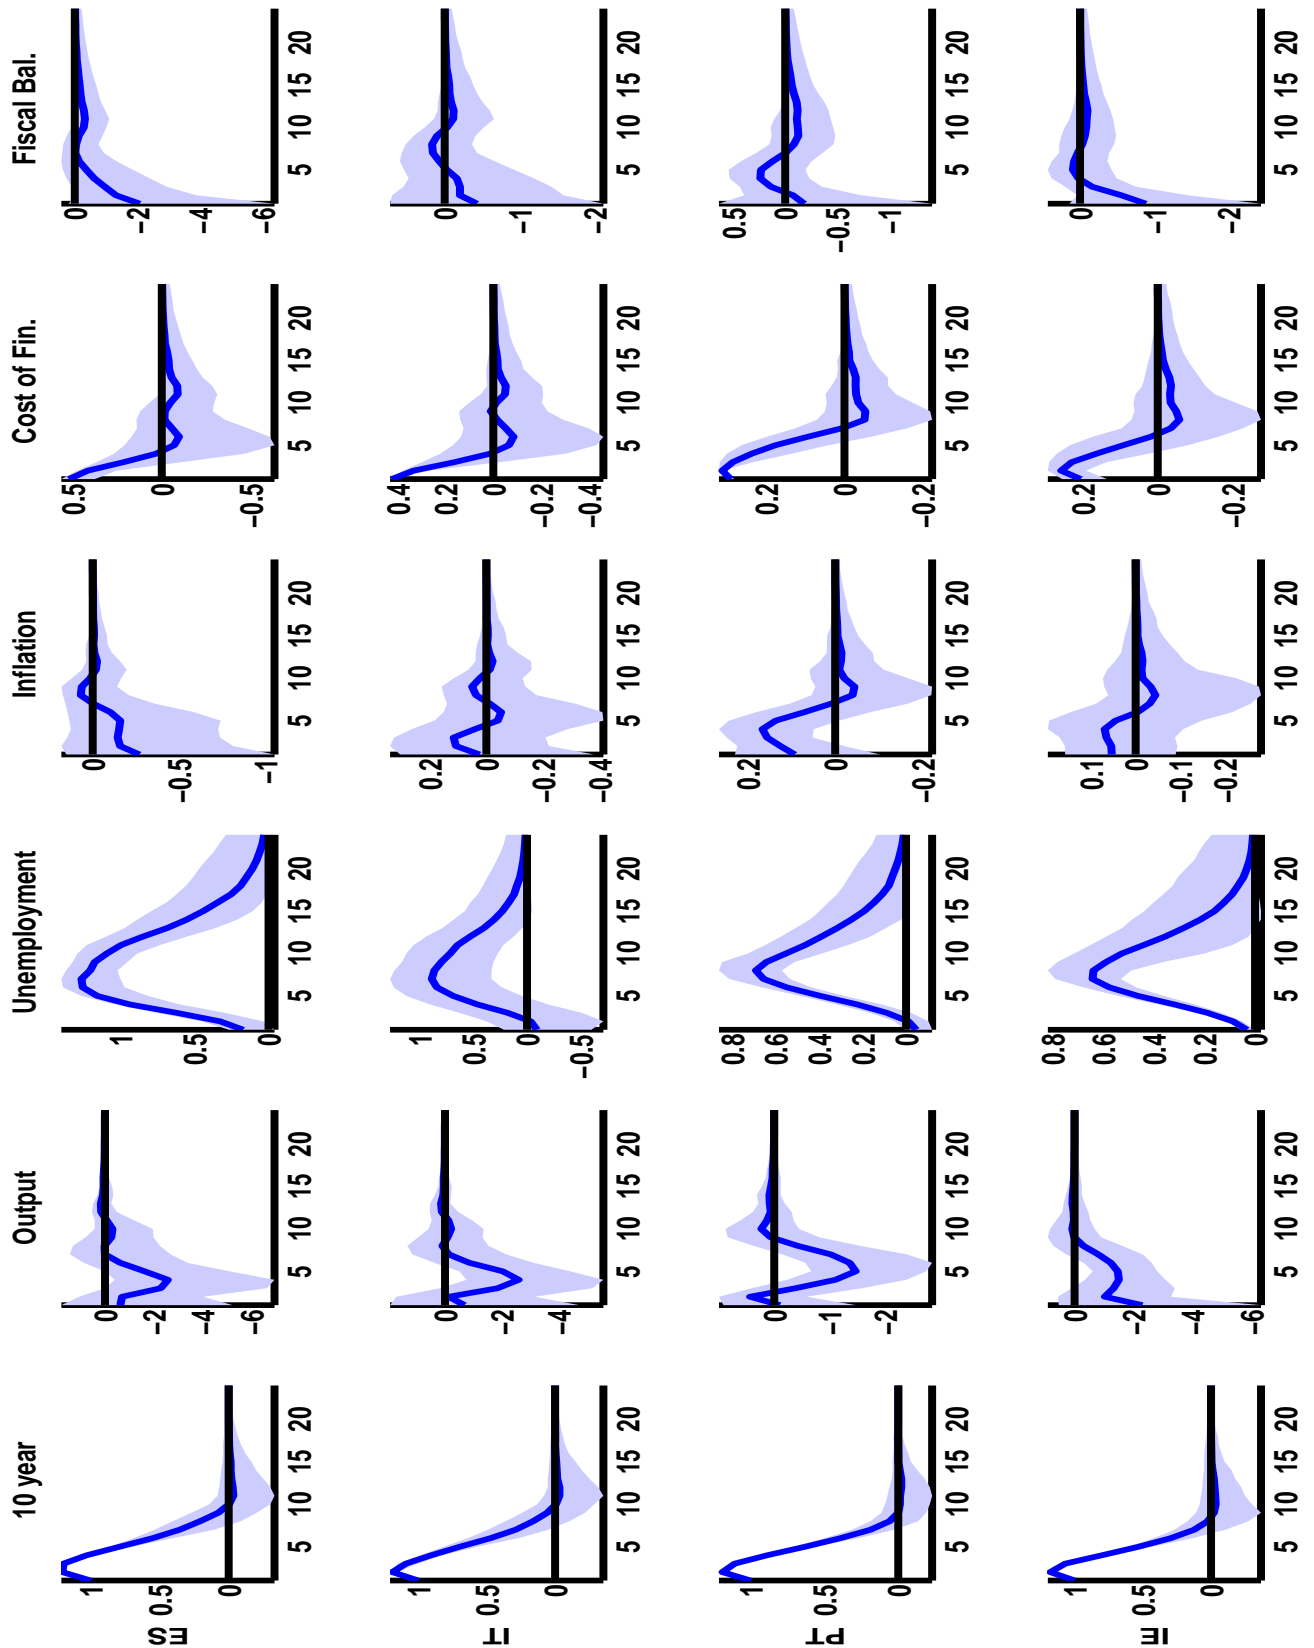
Notes: IRFs are scaled to be consistent with a 100bp increase in sovereign yield on impact and are computed over 24 months. Y-axis is percentage points in all cases. Centre line is the median of 10000 non-sequential draws from the simulated posterior. Error bands are 95% Bayesian credible intervals. Due to similarity with mean country models responses of the Eurepo and OIS rates are not presented for compactness. 10 year refers to the 10 year sovereign bond, output refers to industrial production, Cost of fin. to the private cost of finance. For exact data definitions see main text and the data appendix.

Figure 8: Forecast Error Variance Decomposition (Contribution of Systemic Shock)
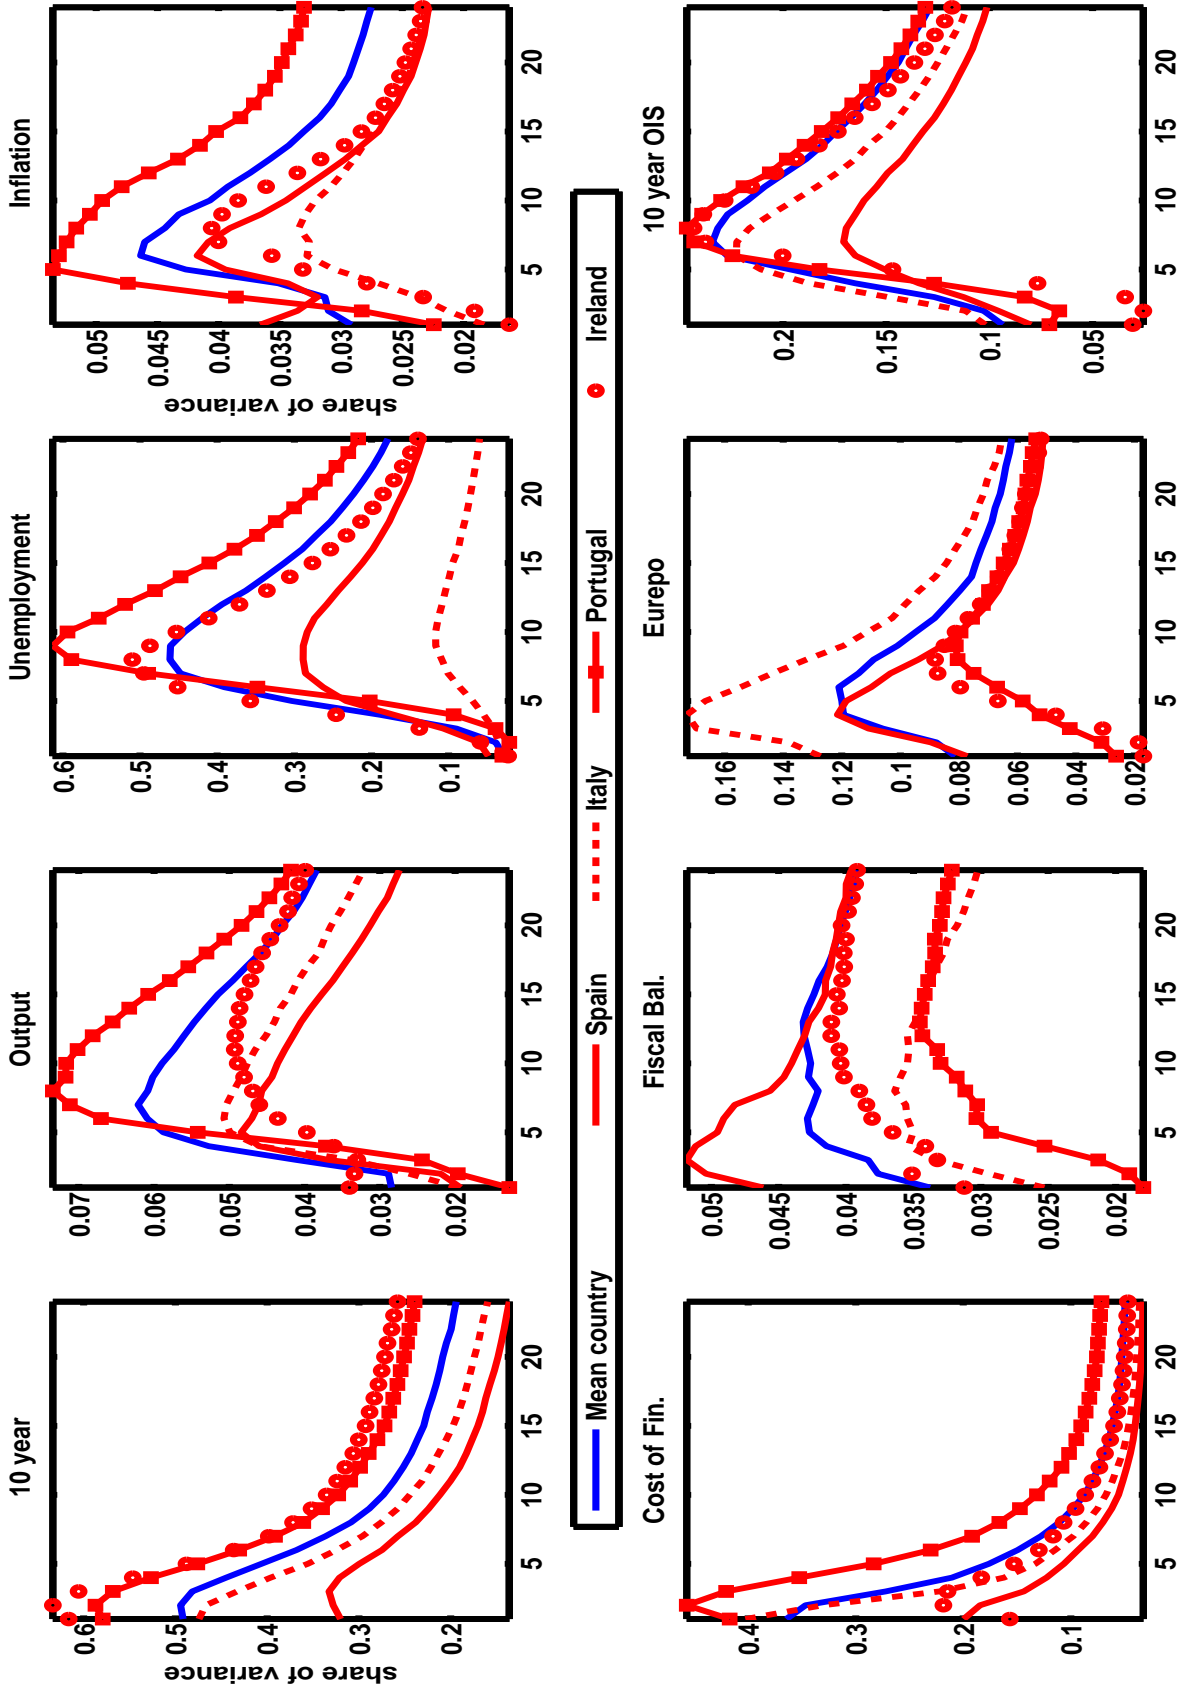
Notes: Forecast error variance decompositions to systemic shock. Computed over a 24 month horizon (x-axis). Line is the median of 10000 non-sequential draws from the simulated posterior. Mean country model (in blue) refers to decomposition estimated using $\bar{\beta}$, $\bar{\Upsilon}$ and $\bar{\Sigma}$. 10 year refers to the 10 year sovereign bond, output refers to industrial production, Cost of fin. to the private cost of finance. For exact data definitions see main text and the data appendix.

Figure 9: Mean Country Impulse Responses (augmented with additional private financing sources)


Notes: Impulse responses to a systemic shock scaled to be consistent with a 100bp increase in sovereign yield on impact and computed over 24 months. VAR is augmented to include four additional financial variables: A composite corporate bond yield (corp yield), the cost loans faced by private sector, equity prices and bank credit supply. Y-axis is percentage points in all cases. Mean country model refers to impulse responses estimated using $\bar{\beta}$, $\tilde{\Upsilon}$ 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 sovereign bond, output refers to industrial production. For exact data definitions see main text and the data appendix.

Figure 10: Mean Country Impulse Responses (augmented with additional external series)
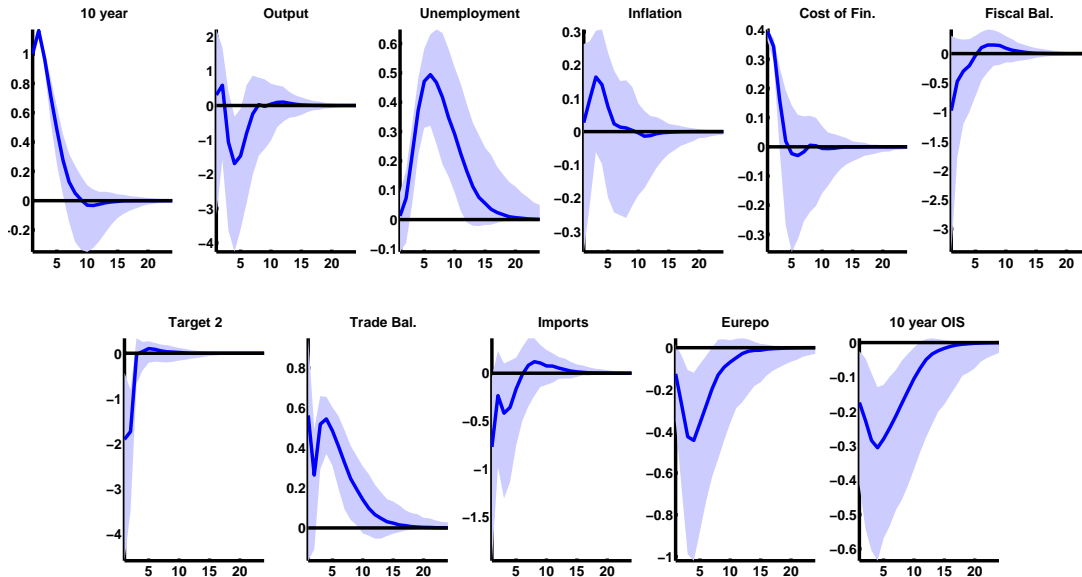
Notes: Impulse responses to a systemic shock scaled to be consistent with a 100bp increase in sovereign yield on impact and computed over 24 months. VAR is augmented to include three additional external variables: the change in the target 2 balance, the trade balance and imports. Y-axis is percentage points in all cases. Mean country model refers to impulse responses estimated using $\bar{\beta}$, $\tilde{\Upsilon}$ 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 sovereign bond, output refers to industrial production, Cost of fin. to the private cost of finance. For exact data definitions see main text and the data appendix.

Figure 11: Counterfactual Analysis: Bond Yields
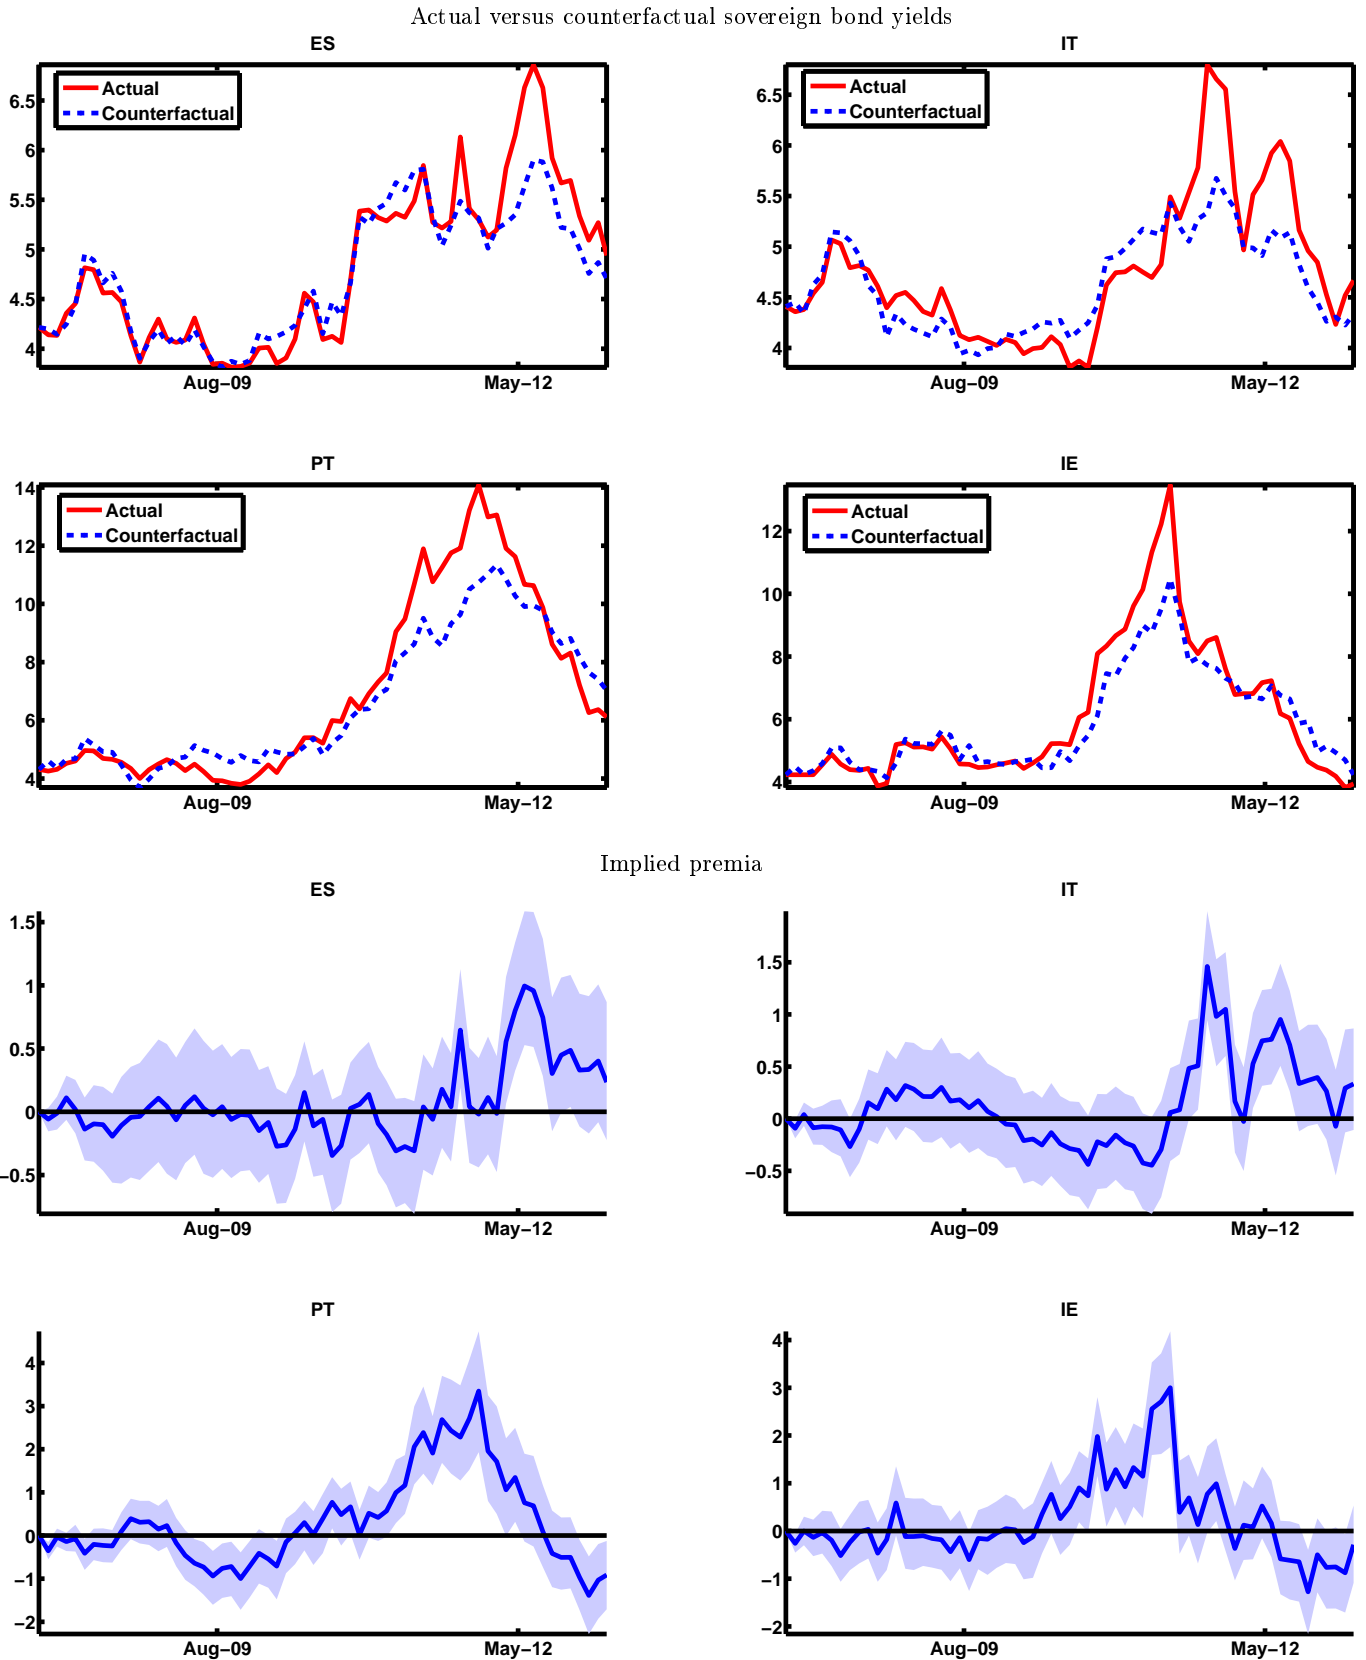
Actual versus counterfactual sovereign bond yields
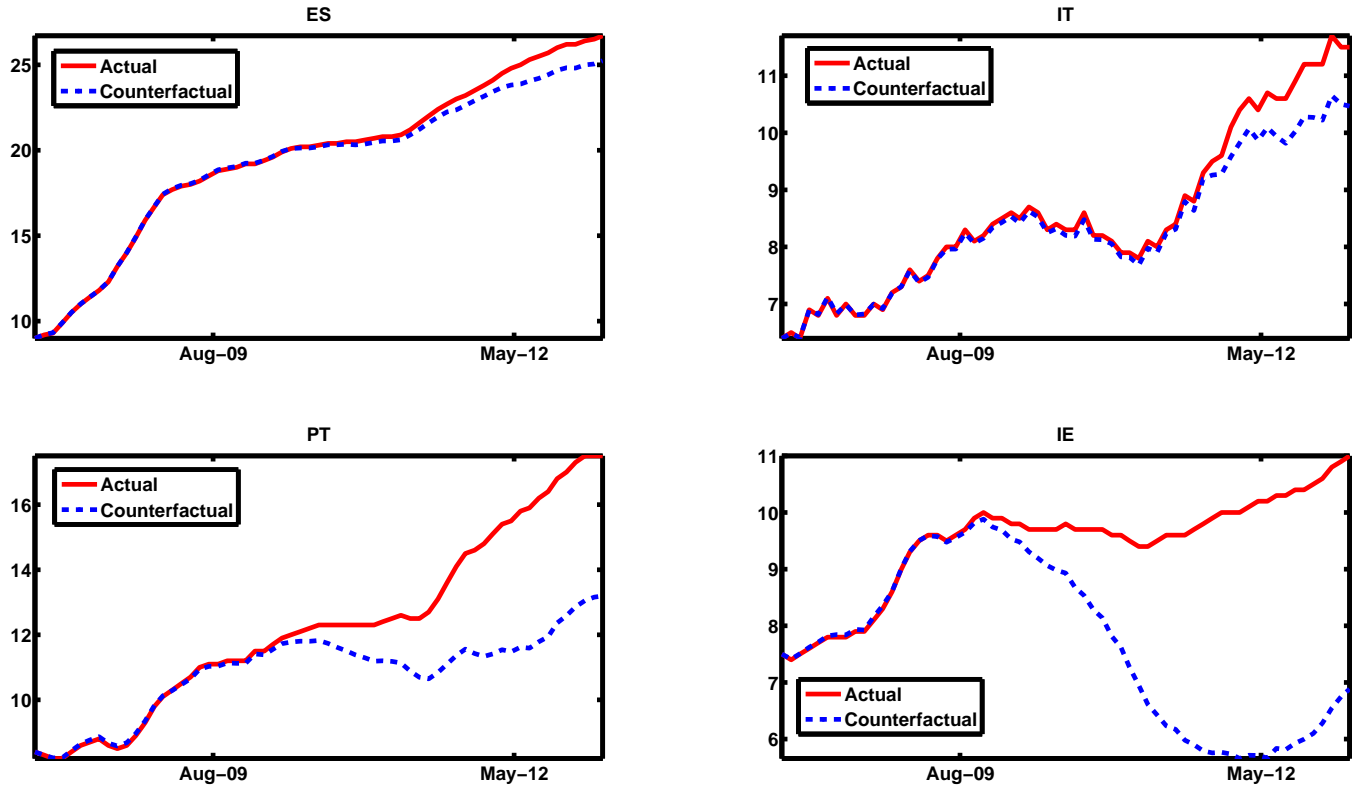

Implied premia

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. Implied premia (lower pane) is equivalent to Actual−Counterfactuals. Error bands are 95% confidence intervals.
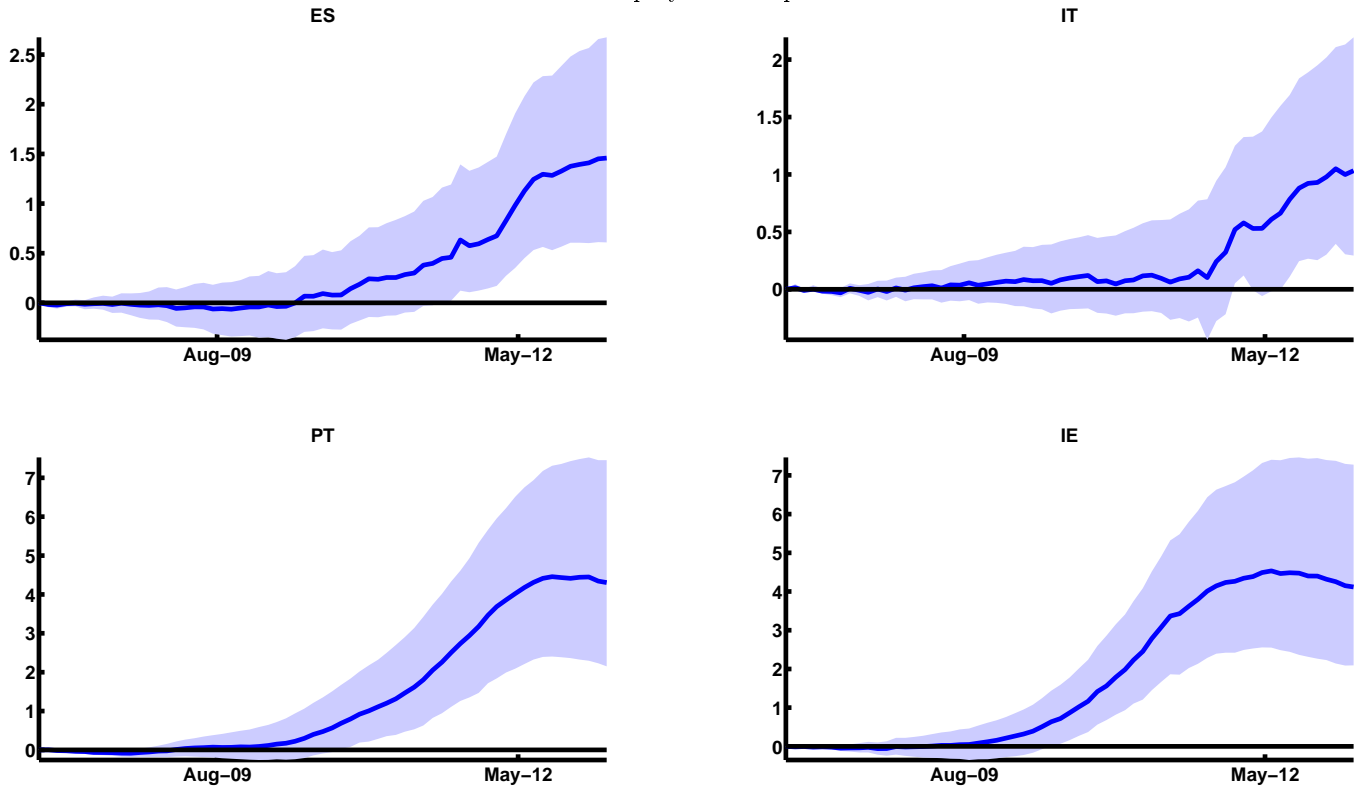
Figure 12: Counterfactual Analysis: Unemployment

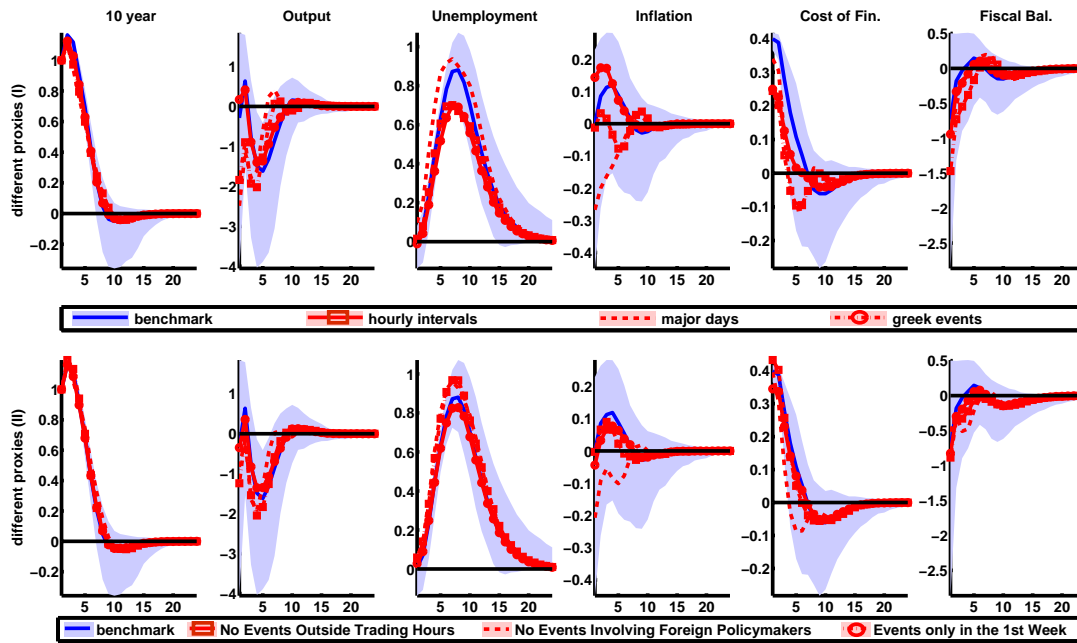Actual versus counterfactual unemployment rates


Unemployment Gap

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. Unemployment gap (lower pane) is equivalent to Actual−Counterfactuals. Error bands are 95% confidence intervals.

Figure 13: Mean country impulse response under alternative proxy definitions

benchmark — **hourly intervals** — **major days** — **greek events**

benchmark — **No Events Outside Trading Hours** — **No Events Involving Foreign Policymakers** — **Events only in the 1st Week**

Notes: Comparison of results with alternative proxy definitions. Blue line and shaded areas are the benchmark case; the red lines represent different proxy definitions. Top pane has alternative proxies constructed using hourly event windows, using daily yield changes on days of major events and using events only related to Greece. Lower pain has alternative proxies omitting certain types of events including: events overnight, events involving foreign policymakers and events in the last 3 weeks of the month. 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 14: Placebo Study
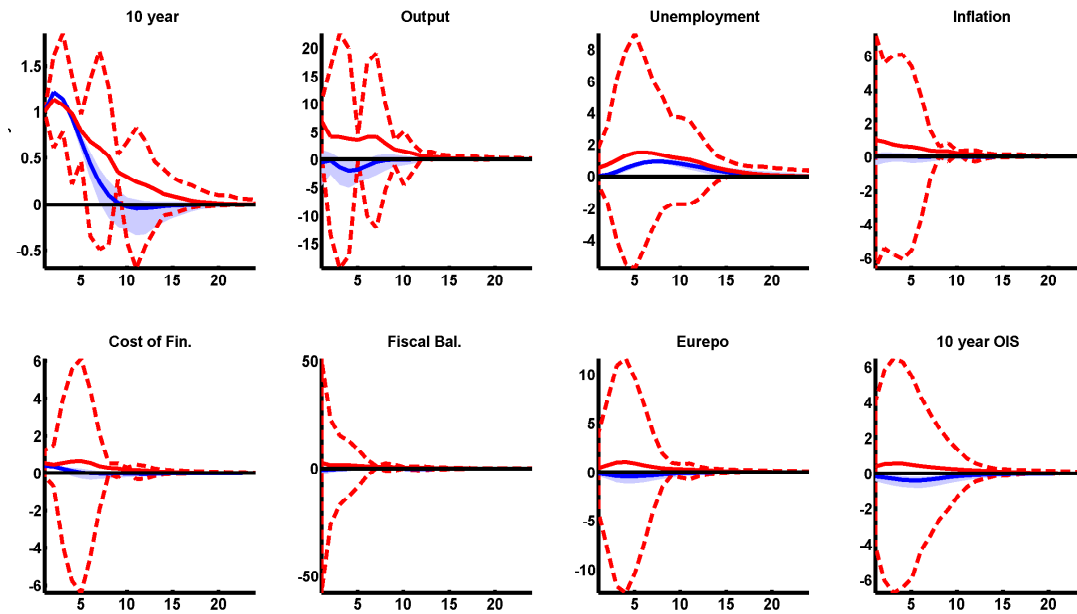
Notes: Results from the Placebo study compared with the benchmark. Placebo proxy constructed using same events timed to the previous trading day. Red line with dash error bands refers to placebo study with 95% credible intervals. Blue line and shaded areas refer to the benchmark case. Y-axis is in percentage points, X-axis is months. Impulse responses scaled to be consistent with a 100bps increase in the bond yield and computed over 24 months.