Credit risk contagion and arbitrage: Evidence from sovereign bond and credit default swap markets

Zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

an der Fakultät für Wirtschaftswissenschaften des Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

Jörg Urban

Tag der mündlichen Prüfung:	9. August 2017
Referent:	Prof. Dr. Melanie Schienle
Korreferent:	Prof. Dr. Peter Kugler
Karlsruhe	20. Juni 2017

I am deeply indebted to Dr. Kristyna Ters for all her support and encouragement during a fruitful and inspiring collaboration since 2011. My thanks also go to Dr. Peter Hördahl and Dr. Jacob Gyntelberg for our joint work at the Bank for International Settlements and for their support. I am also grateful to my supervisors Prof. Dr. Melanie Schienle and Prof. Dr. Peter Kugler whose insights have strongly contributed to the success of my research. In addition, thanks are due to Dr. Sebastian Deininger for many discussions and advice. Last but not least, I am grateful to my family for their support, motivation and patience.

To Silke, Felix, Max, mum and dad

Contents

Preface

The Global Financial Crisis started in 2007 with the collapse of the subprime mortgage market in the USA and later developed into a full scale international crisis lasting until beginning of 2009. It is considered the worst financial crisis since the Great Depression (1929-1939). The causes of the crisis are multi-dimensional and include, amongst others, subprime lending and ease of credit conditions leading to a growth of the housing bubble, weak and fraudulent underwriting practices, deregulations, financial innovations leading to complexity of financial products as well as over-leveraging.

The excessive risk-taking by banks and their interconnectedness magnified the impact globally. The collapse of the investment bank Lehman Brothers on 15 September 2008 was a key event of the financial crisis. In order to avoid a global collapse of the financial system, policy makers and politicians were forced to start large scale bailouts of financial institutions and introduce unconventional monetary and fiscal policies. However, these unprecedented measures were not sufficient to avoid a global economic downturn.

The focus of this thesis is the European sovereign debt crisis, which followed the financial crisis, starting around the end of 2009. Several euro zone member states (eg Greece, Ireland, Portugal) were unable to repay or refinance their government debt. Spain had a relatively low debt-to-GDP ratio compared to the other GIIPS countries[1] before the crisis but was forced to bailout out over-indebted banks under their national supervision. Greece, Ireland, Portugal and Spain as well as Cyprus were unable to meet their financial obligations without the assistance of third parties such as the European Commission, the European Central Bank (ECB), and the International Monetary Fund (IMF).

By measuring credit risk in terms of credit default swap (CDS) spreads[2] it is obvious from Figure 1 that the GIIPS or peripheral countries were affected by the euro area sovereign debt crisis more severely than by the preceding global financial crisis. On the other hand, the US and the UK were, as expected, clearly more affected by the Global Financial Crisis than by the euro area sovereign debt crisis (see right-hand graph of Panel B in Figure 1). The GIIPS countries' CDS spreads were on average four times higher during the euro area sovereign debt crisis compared to the financial crisis, except for Greece, where the CDS spread increased tenfold until the Greek debt underwent restructuring (starting in March 2010).

All euro area countries saw a continuing increase in credit risk concerns from around the end of 2009 until the speech of ECB president Mario Draghi on the 26 July 2012,

[1] GIIPS countries are Greece, Ireland, Italy, Portugal and Spain.

[2] A CDS is a financial instrument which offers protection against a potential loss from a credit event on debt issued by a borrower. Usually a protection buyer pays a quarterly premium in basis points (bps) of the insured amount to the protection seller. CDS contracts are over-the-counter contracts, but increasing standardisation is being enforced by the International Swaps and Derivatives Association.

where he clearly stated that the ECB was prepared to do "whatever it takes" to save the euro and to calm down credit risk markets. Most CDS spreads started to decline from mid-2012.

Figure 1: Sovereign Credit Default Swaps

The figure illustrates the evolution of CDS spreads in basis points from the beginning of the financial crisis up to the end of 2014. Panel A shows that credit risk of the GIIPS countries is priced much higher during the euro area sovereign debt crisis, relative to the Global Financial Crisis. The situation is opposite for non-euro area members, most notably the US and UK, where the financial crisis was considered by credit markets to be more severe. The vertical lines correspond to the Lehman Brothers default (15.9.2008), the announcement by the Greek government that official statistics had been fabricated (20.10.2009), the start of the sovereign debt crisis (1.4.2010) as defined by van Rixtel and Gasperini (2013) and the speech by ECB president Mario Draghi (26.7.2012). The missing data for Greece in 2012 and 2013 is due to the Greek debt restructuring. Source: Markit.

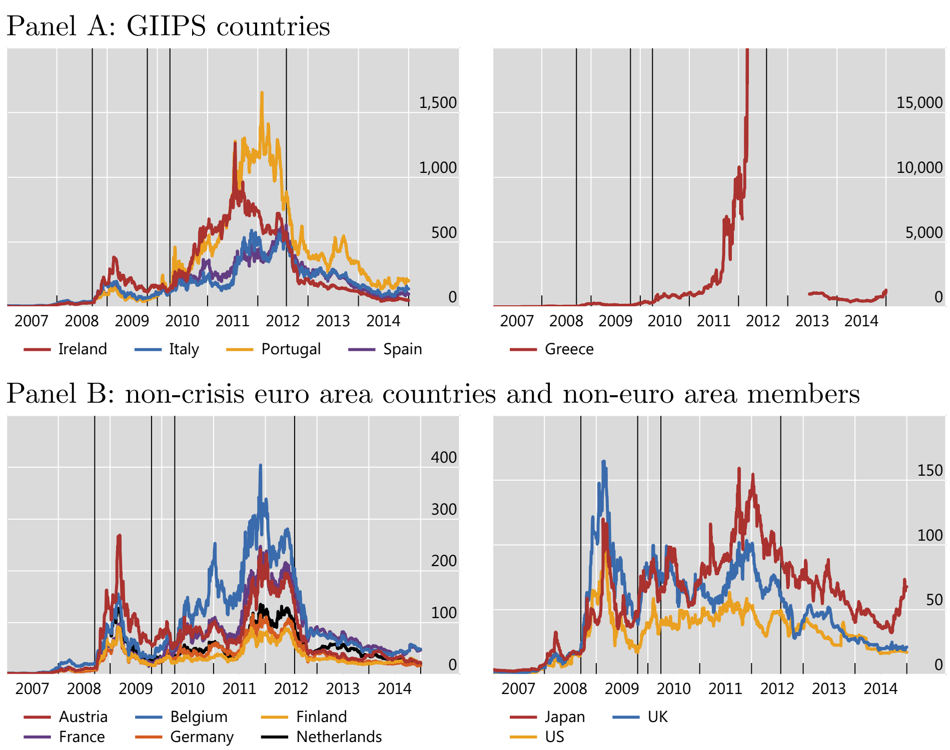

In Part I of this thesis we examine the role of the CDS and bond market as transmission channels for credit risk contagion between sovereign entities in the pre-crisis and crisis period. This first part of the thesis is fully based on three publications: Komarek et al. (2016), Ters and Urban (2017b) and Ters and Urban (2017a). The contributions are manifold, such as we use intraday data for bonds and CDS before and during the euro area sovereign debt crisis to model credit risk contagion amongst the GIIPS countries and achieve sound statistical inferences. We find that the CDS market is the most important venue of credit risk trading during the crisis period, while before the crisis the bond market was of similar importance. During the crisis we see significant flight-to-safety effects in the German bond market. The crisis period was coined by contagion amongst the GIIPS countries. Therefore, we have also performed an event study to quantify the effectiveness of the economic adjustment programmes of the European Commission, the European Central Bank and the International Monetary Fund. We found that the programmes had a stabilising effect for the country under bailout. The second bailout for Greece resulted in a significant stabilisation of the entire network of the GIIPS countries. We further contributed to the literature by extending the analysis to include central European countries in the analysis, such as the Czech Republic, Hungary, Poland and Slovakia. These countries have strong trade linkages to the euro zone. Despite these strong trade linkages, the negative effects of the euro area crisis were only marginally visible in these countries in a sense that we have found comovement of credit risk markets and no contagion effects.

On average, the CDS and the bond market should price credit risk of a given reference entity equally. However, the pricing of sovereign risk and the dynamics in the CDS and the bond markets have diverged during the euro area sovereign debt crisis. Therefore, it is necessary to analyse both credit markets in order to understand the relative importance of the CDS and bond markets as transmission channels for credit risk contagion.

Sovereign entities are mainly issuing fixed coupon (plain vanilla) bonds. These bonds carry in addition to the credit risk also an interest rate risk component. Through an interest rate swap, the exchange of the fixed coupon for a floating rate, eg Euribor, it is possible to maintain the original credit risk exposure, but transfer the interest rate risk to another party. The resulting asset swap spread A in the right-hand side of Equation (1) is then technically identical to a floating rate note with identical credit exposure and hence ensures proper comparability with a CDS. The par asset swap spread (ASW) is the most commonly traded asset swap spread and is defined as:[3]

$$\underbrace{100-P}_{\substack{\text{Upfront payment for bond}\\ \text{asset in return for par}}} \quad + \overbrace{\underbrace{C\sum_{i=1}^{N_{\text{fixed}}} d(t_i)}_{\text{Fixed payments}} = \underbrace{\sum_{i=1}^{N_{\text{float}}} \big(L_i + A\big) d(t_i)}_{\text{Floating payments}}}^{\text{Interest rate swap}}, \tag{1}$$

[3] See O'Kane (2000) and Gale (2006) for detailed discussions of the mechanics and pricing of asset swaps.

where P is the full (dirty) price of the bond, C is the bond coupon, L_i is the floating reference rate (e.g. Euribor) at time t_i and $d(t_i)$ is the discount factor applicable to the corresponding cash flow at time t_i.

The use of ASW in this project is in line with the practice used in commercial banks when trading the CDS-bond basis and also a new contribution to the academic literature. No-arbitrage arguments can be employed to show that CDS and ASW for the same reference entity and maturity should be priced identically (Duffie; 1999). However, for this identity to hold, the markets must be perfect and frictionless, there should be no tax effects and it should be possible to short bonds without restrictions and costs. These ideal conditions do not hold in periods of normal trading and we see strong distortions during the crisis period (see the CDS-ASW basis in Figure 2). For example, the basis for Ireland and Portugal went up to 400bps during the crisis period. The deviation of credit risk pricing in the two markets have reduced since mid-2012 reaching pre-crisis levels in 2014.

Figure 2: CDS-ASW basis

The figure shows the CDS-ASW basis for four countries which were strongly affected by the euro area sovereign debt crisis. The basis was computed using intraday data on 30 minute sampling frequency. Apart from a period around 2013, the basis is positive. We find strong deviations in the pricing of credit risk in the CDS and bond market during the euro area sovereign debt crisis. Source: CMA, EuroMTS, author's calculations.

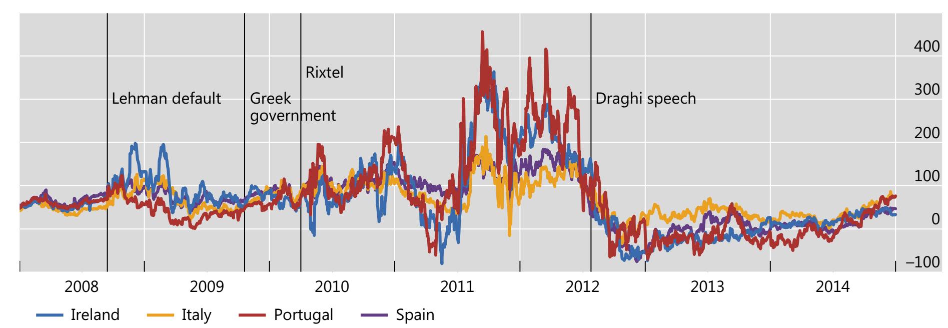

We have used intraday datasets from the MTS trading platform for government bonds and from CMA Datavision for sovereign CDS. We have aggregated our time series data into 30 minute equidistant time intervals over the trading day from 8:30 until 17:30. The use of intraday data allows us to achieve much higher statistical inferences as compared to daily data. Further, we are able to capture intraday effects, which is clearly essential in fast moving financial markets.

Our baseline model is the following bi-variate panel VAR model:

$$\begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{it} = \begin{pmatrix} A_{01} \\ A_{02} \end{pmatrix}_i + \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}_{ji} (\mathrm{L}) \begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{j\,t-1} + \begin{pmatrix} \tilde{F}_1 \\ \tilde{F}_2 \end{pmatrix}_i (\mathrm{L}) X_t + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}_{it}, \tag{2}$$

where A_{mn} and $\tilde{F}_m$ are scalars, L is the lag operator, X_t are exogenous variables and u_i are i.i.d. shocks. The indices i and j run from 1 to N, representing the different sovereign entities. $t = 1, \dots, T$, where T is the length of our time series.

Our analysis reveals a couple of interesting and intuitive results. Prior to the crisis (Jan. 2008 - Oct. 2009), the CDS and bond markets were similarly important in the transmission of sovereign risk, but the importance of the bond market waned during the crisis (Oct. 2009- Dec. 2011). We find flight-to-safety effects during the crisis in the German bond market that are not present in the pre-crisis sample. The estimated sovereign risk contagion was greater during the crisis, with an average timeline of one to two hours in the GIIPS countries. By using an exogenous macroeconomic news shock, we can show that, during the crisis period, increased credit risk was not related to economic fundamentals.

Figure 1 shows that markets become increasingly worried about the sustainability of sovereign debt especially in Greece, Ireland and Portugal in mid-2010 and so intensified concerns on the side of policy makers about the fiscal outlook for the entire euro zone. These concerns led to the implementation of a series of financial support measures such as the European Financial Stability Facility (EFSF) and the European Stability Mechanism (ESM). The ECB and other central banks also contributed to resolve the crisis by lowering interest rates and providing emergency liquidity to European banks. The time lines of the bailouts/economic adjustment programmes (EAP) for Greece, Ireland and Portugal by the Troika, a decision group formed by the European Commission, the ECB and the IMF, are shown in Figure 3. We perform an event study to quantify the impact of the individual bailouts in order to measure their success in reducing the potential risk of contagion. Our findings imply that the introduction of the EAP had a positive effect on financial market's risk perception for the individual country under bailout.

We also investigated whether contagion led to higher sovereign risk in the Visegrad group (Czech Republic, Hungary, Poland and Slovakia) and furthermore, whether the EAP by the Troika had been able to stabilise and reduce sovereign risk (see Figure 3 for the time line of the EAP) for the Visegrad group. We find comovement[4] effects in the Visegrad group member countries as they were only marginally affected by the

[4] When the impulse response is gradual we refer to it as comovement rather than contagion. Our contagion identification strategy builds upon the criteria described by the ECB where contagion is defined as (i) the transmission is in excess of what can be explained by economic fundamentals; (ii)

turmoil in the peripheral countries during the sovereign debt crisis. In contrast, we find strong contagion effects amongst the GIIPS countries in our sample. The EAP have been essential for the GIIPS countries in terms of reducing contagion and sovereign risk across the euro area while the Visegrad group reacted only with a moderate reduction of credit risk spillovers. In conclusion, we find that the Visegrad group countries were not affected by sovereign credit risk contagion, irrespective of their debt level and currency.[5]

Figure 3: Economic adjustment programmes

The figure shows 30-minute CDS spreads for Greece, Ireland and Portugal. The vertical lines correspond to the four bailout events during our sample period. We define the bailout events as the announcement dates of the Memorandum of Understanding for the economic adjustment programmes (EAP) in Ireland, Portugal and the first bailout in Greece. The event date of the second bailout for Greece is defined as the announcement of the preliminary draft of the second EAP.

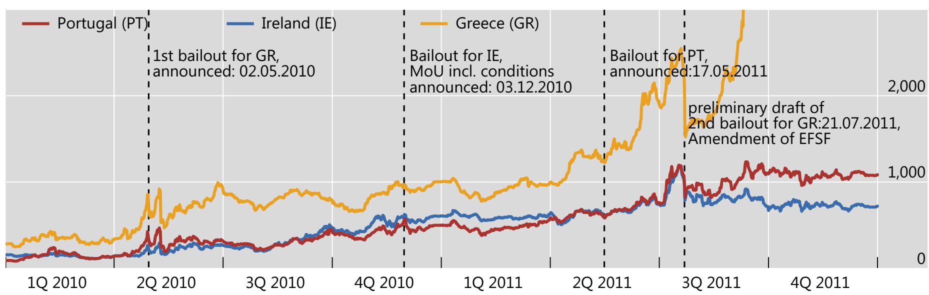

Part II of this thesis is fully based on the following two papers: Urban (2017) and Gyntelberg et al. (2017). The key contributions to the academic literature are that we develop a methodology that endogenously estimates unknown arbitrage costs for a positive and a negative basis trade in markets with a persistent non-zero basis between two similar financial market instruments traded in the spot and the derivative market. We perform a thorough statistical analysis of this computationally complex threshold vector error correction model (TVECM) and apply it as an illustration to the following commodity markets: gold and platinum as well as for the stock market indices: DAX and S&P 500. As a further contribution we perform an in depth analysis, using a TVECM, of the credit risk markets during the euro area sovereign debt crisis. By using intraday data we find strong evidence for threshold effects, which likely reflect costs that arbitrageurs face when implementing their trading strategies. The transaction costs are significantly higher during the crisis as compared to the pre-crisis.

the transmission is different from regular adjustments observed in tranquil times; (iii) the events constituting contagion are negative extremes; and (iv) the transmission is sequential.

[5] Slovakia is an euro area member, while the other three Visegrad group countries maintain their own legal tender. Hungary has the highest debt-to-GDP ratio (around 80% during the crisis) amongst the Visegrad group countries.

We are using a regime switching error-correction model which enables us to estimate the area where arbitrageurs have no incentives for trading (see Figure 4). Only when the basis[6], exceeds a critical threshold, where the potential gain from the basis trade is above the overall transaction costs, can we expect arbitrageurs to step in and carry out the respective trade. This leads to non-linear adjustment dynamics and different regimes. Our methodology allows us to estimate the overall transaction costs for an arbitrage trade in markets where transaction costs are opaque or unknown such as markets for credit risk or index trading.

Figure 4: Vector error correction models

The vector error correction model (VECM) in the left-hand graph represents markets or periods, where the basis does not deviate too strongly from zero. The markets are perfect and frictionless for arbitrageurs to step in immediately to correct any pricing differential between the spot and derivative market. The middle and right-hand graph represents classical multi-regime models, where arbitrageurs engage in basis trades only outside the neutral regime. The middle figure would be mirrored in case of a persistent negative basis, resulting in a negative threshold. In that case there would be arbitrage on basis strengthening. The horizontal axis is the time-axis and the vertical axis show the price of the basis (eg in credit risk markets the basis is measured in basis points).

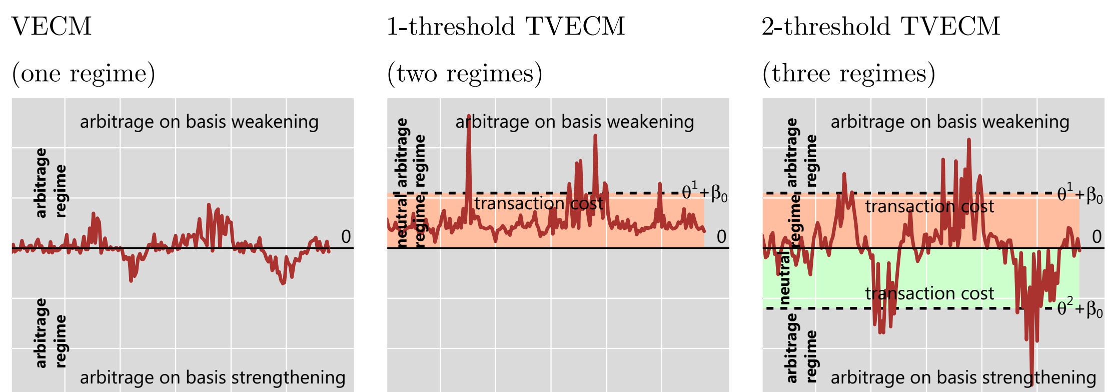

We have used the following threshold-VECM with l regimes and $l-1$ thresholds in part II of this thesis:

$$\Delta y_t = \sum_{j=1}^{l} \left[\underbrace{\lambda^j (S - \beta_1 D - \beta_0)_{t-1}}_{\text{error correction term}} + \underbrace{\Gamma^j(L)\Delta y_t}_{\text{VAR}} \right] d_t(\beta_0,\, \beta_1,\, \theta^{j-1},\, \theta^j) + \varepsilon_t, \tag{3}$$

6 We define the basis as the spot price minus the futures or derivative price. In credit risk markets the basis is defined as derivative minus spot price, ie the basis is CDS-bond.

where we focus on the bivariate case $y_t = (S_t \ D_t)^\mathsf{T}$, with S denoting the time series of spot prices and D denoting the time series of derivative prices for the same reference entity. All thresholds θ^j are ordered and

$$d_t(\beta_0, \beta_1, \theta^{j-1}, \theta^j) = I(\theta^{j-1} \leq (S - \beta_1 D - \beta_0)_{t-1} < \theta^j)\,, \tag{4}$$

where the indicator function $I(\cdot)$ is 1 if the error correction term $(S - \beta_1 D - \beta_0)_{t-1}$ is in the interval $[\theta^{j-1}, \theta^j)$ and otherwise 0. $\varepsilon_t = (\varepsilon_t^S \ \varepsilon_t^D)^\mathsf{T}$ is a vector of i.i.d. shocks and $j \in \{1, 2, ..., l\}$ are the regimes. The thresholds θ^0 and θ^l are by definition $-\infty$ and ∞, respectively.

Equation (3) constitutes a vector autoregressive model in first-order difference with $\Gamma^j(L) = \sum_{k=1}^{m} \alpha^{j,k} L^k$ and L as lag operator, m as number of VAR lags and an additional error correction term $\lambda^j(S - \beta_1 D - \beta_0)_{t-1}$. The speed of adjustment parameters $\lambda^j = (\lambda_1^j \quad \lambda_2^j)^\mathsf{T}$ and the lagged VAR terms are regime-dependent.

The error correction term represents the long-term equilibrium and the VAR-term represents the short-run dynamics coming from market imperfections (Baillie et al.; 2002). The overall transaction costs for the basis trade are the estimated threshold plus the intercept β_0 in the error correction term.

In line with part I of the thesis, we conduct a thorough analysis using the proposed TVECM method to the euro area sovereign CDS and bond markets. We find evidence that in the market for euro area sovereign credit risk, arbitrageurs engage in basis trades between CDS and bond markets only when the CDS-bond basis exceeds a certain threshold. This threshold effect is likely to reflect costs that arbitrageurs face when implementing trading strategies, including transaction costs and costs associated with committing balance sheet space for such trades. During the euro area sovereign credit crisis, we find very high transaction costs, compared to around 100bps before the crisis, which is shown for the case of Italy (transaction costs during the crisis are around 150bps) and Portugal (transaction costs during the crisis are around 300bps) in Figure 5.

Figure 5: Transaction costs

The figure shows the CDS-ASW basis for Italy and Portugal (blue curve) and the estimated transaction costs of a basis trade before and during the crisis (red line). As expected, the transaction costs are higher during the crisis. Source: CMA data vision, MTS, author's calculations.

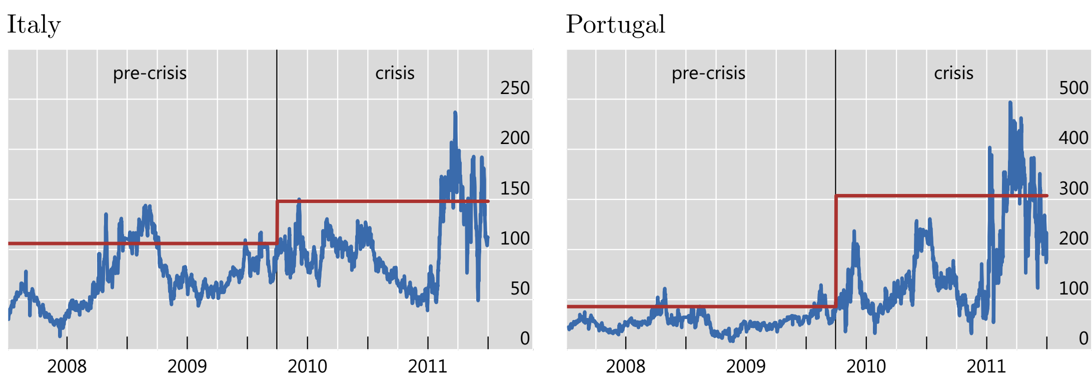

Our results further show, that even when markets are liquid in times of stress, the basis can widen as high market volatility makes arbitrage trades riskier, leading arbitrageurs to demand a higher compensation for increased risk. Our findings help explain the persistent non-zero CDS-bond basis in euro area sovereign debt markets and its increase during the last sovereign debt crisis.

Part I

The transmission of euro area sovereign credit risk

Part Ia: Intraday dynamics of euro area sovereign credit risk contagion[1]

1 Introduction

The 2008-09 financial crisis caused investors to look more critically at the fiscal outlook in a number of countries, including several in the euro area. This resulted in a sharp rise in sovereign credit spreads for a number of euro area countries. At their peak, yield spreads on sovereign bonds relative to German bonds reached several hundred basis points, while before the global financial crisis these spreads had averaged only a few basis points. At the same time, market interest in trading credit risk protection on euro sovereign borrowers via credit default swaps (CDS) grew substantially and spreads on such instruments also surged. Due to these developments, policy makers and regulators paid increased attention to the market for sovereign CDS. Of particular interest was, first, the interplay between the pricing of sovereign risk in CDS and bond markets, secondly the possibility that one market could be systematically leading the other, and thirdly the potential for sovereign credit risk contagion effects. In this paper we focus on the latter.

We analyse euro area sovereign credit risk contagion effects in GIIPS[2] countries plus France and Germany from January 2008 to end-December 2011, which we split into a pre-crisis and crisis period. Further, we investigate if and how sovereign credit risk contagion was transmitted from the GIIPS countries to central European countries (Austria, the Czech Republic, Hungary, Poland and Slovakia) during the euro area sovereign debt crisis. Austria is included as low-risk reference country for the Czech Republic. The use of intraday CDS and bond data lets us estimate credit risk contagion effects with substantially more accuracy than existing studies on sovereign credit markets have done. In addition, little is yet known about the transmission channels of credit risk contagion through the CDS and the bond market, and their relative importance in the euro area sovereign debt crisis. As we have data for both the CDS market and the bond market, we are able to assess the contagion impacts conditioned on the credit channel.

The use of intraday data allows us to capture the intraday patterns of credit risk contagion. Indeed, shocks that may seem to affect several countries simultaneously when viewed at a daily or lower data frequency are revealed, through the lens of intraday data, to have possible origins in one particular country with clear contagion effects on other countries. Also, Gyntelberg et al. (2013) discuss the advantages of using intraday data

[1] This chapter is joint work with Lubos Komarek (Czech National Bank) and Kristyna Ters (University of Basel) and has been published as: Komarek, L. Ters, K. and Urban, J. (2016) Intraday dynamics of euro area sovereign credit risk contagion. *BIS Working Papers No 573*, Bank for International Settlements and *CNB WP 4/2016*, Czech National Bank.

[2] Greece, Ireland, Italy, Portugal, Spain

due to the higher accuracy of the results as compared with lower-frequency data. They find that when using daily data, due to the dramatically smaller number of observations, the confidence bands of the estimated coefficients are extremely wide and therefore in most cases not significant. Existing research has differentiated between cross-country and intra-country analysis. Using a panel VAR methodology we can control for both country-specific risk and contagion effects across countries. Panel VARs are built on the same logic as standard VARs but, by adding a cross-sectional dimension, they become a much more powerful tool for addressing policy questions of interest related, for example, to the transmission of shocks across borders (Canova and Ciccarelli; 2013). By using the method of Canova and Ciccarelli (2013), we are able to shock the credit risk of an individual country and derive the individual response for each country in the panel.

A large body of literature concerns itself with the potential reasons and transmission channels for contagion as well as with the theoretical modelling of contagion. A whole strand of this literature focuses on empirical tests for the existence of contagion in a given stress period, that is, it asks if there are stronger cross-market linkages in times of crisis. This paper belongs to the latter type, as we focus on testing for the existence of contagion during the euro area sovereign debt crisis. We extent existing research by analysing the relative importance of the CDS and the bond market as transmission channels for sovereign risk contagion.

An important motivation for providing financial support to Greece, despite the no-bailout clause in the Maastricht Treaty, was the fear on the part of policy makers that a Greek default would spill contagiously over to other highly indebted countries in the euro area (Constancio; 2011). As pointed out by Corsetti et al. (2011), there is much disagreement among economists about the exact definition of contagion and how it should be tested. For Constancio (2011) and Forbes (2012) contagion occurs when financial or macroeconomic imbalances (shocks) create a systemic risk beyond that explained by economic fundamentals. Contagion differs from macroeconomic interdependence (comovement) among countries in that the transmission of risk to other countries is different under normal economic conditions. Forbes (2012) defines contagion as spillovers resulting from extreme negative effects. If comovements of markets are similarly high during non-crisis periods and crisis periods, then there is only evidence of strong economic linkages between these economies (Missio and Watzka; 2011). Kaminsky et al. (2003) describe contagion as an episode in which there are significant immediate effects in a number of countries following an event, such as when the consequences are fast and furious and evolve over a matter of hours or days. When the effect is gradual, Kaminsky et al. (2003) refer to it as spillovers rather than contagion. We rely on the contagion and comovement definitions according to Kaminsky et al. (2003), Constancio (2011) and Forbes (2012).

As there is a vast literature on contagion, we limit our discussion to papers that measure contagion among sovereign credit markets. Kamin and von Kleist (1999), Eichengreen and Mody (2000), Mauro et al. (2002), Pan and Singleton (2005), Longstaff et al. (2011), and Ang and Longstaff (2011) concentrate on the relationship between sovereign credit spreads and common global and financial market factors. These papers empirically identify factors that are significant variables for CDS credit spreads, such as the U.S. stock and bond market returns as well as the embedded volatility risk premium.

The issues of financial shock contagion and cross-country spillovers among countries in the euro area during the recent sovereign debt crisis have figured prominently in recent empirical research. Caporin et al. (2012) analyse risk contagion using the CDS spreads of the major euro area countries using different econometric approaches such as Bayesian modelling. They find that the diffusion of shocks in euro area CDS has been remarkably constant, while the risk spillover among countries is not affected by the size of the shock. Other examples are Bai et al. (2015), Neri and Ropele (2013), De Santis (2012), and Arghyrou and Kontonikas (2012). They all employ time series modelling approaches for contagion and include sovereign bond spreads (yield to maturity) to reflect pure credit risk considerations and macroeconomic variables. The results are mostly discriminated in terms of core (e.g. Germany and France) and peripheral countries (GIIPS). In general, they find that the bond spreads of lower-rated countries increase along with their Greek counterparts. However, their results in terms of magnitude, responses to shocks and contagion effects on core countries are somewhat mixed. Similarly to these studies, Koop and Korobilis (2016) employ an enhanced panel approach for empirical modelling of financial contagion across countries (Canova and Ciccarelli; 2013). Their findings are at odds with the discrimination between core and peripheral countries, as they also find contagion from GIIPS to core countries, albeit smaller in magnitude. The different results reported in these studies could be due to sample differences or to how bond spreads are calculated. Most empirical research uses the "constant maturity" approach to calculate bond yield differences (relative to Germany). Further, daily or weekly data are used for the empirical analysis, which may lead to inaccurate shock and contagion estimations, especially in periods when activity in sovereign risk markets is high during times of stress.

One of the key contributions of our paper to the existing literature on sovereign risk contagion during the recent euro area sovereign debt crisis is that, in contrast to all the above-mentioned studies, we do not use simple yield differences as our measure of cash spreads. Rather, we use carefully constructed asset swap spreads (ASW) based on estimated zero-coupon government bond prices. This ensures that we are comparing like with like in our empirical analysis for sovereign credit risk, by exactly matching the maturities and the cash flow structures of the CDS and the cash components. The use of ASW is also in line with the practice used in commercial banks when trading the CDS-bond basis. In addition, by calculating ASW we are able to estimate contagion impacts on Germany such as flight-to-safety effects. Germany is not included in most contagion studies since German Bund yields are used as the risk-free interest rate in the "constant maturity" approach mentioned above. Moreover, our analysis relies on intraday price data for both CDS and bonds, allowing us to estimate the contagion dynamics and the transmission channel of contagion (CDS or bond market) substantially more accurately than existing studies.[3] Further, by extending our model to include the economic surprise index, we are able to estimate how much of sovereign risk contagion can be attributed to macroeconomic news or overreaction/lack of belief by market participants.

Finally, our findings will improve the understanding of the dynamics in the market for sovereign credit risk. Further, the segregation of the credit risk transmission channels will enable policy makers and regulators to better assess the relative importance of, and the risks arising from, the derivative and cash market.

The rest of the paper is structured as follows: Section 2 discusses our data and the relationship between CDS and bonds. Section 3 explains the set-up and estimation of the panel VAR (PVAR) model and its extension. Section 4 presents the empirical results and Section 5 concludes.

2 Data

The core data we use in our empirical analysis consists of USD-denominated five-year maturity intraday quotes on CDS contracts and government bonds for France, Germany, Greece, Ireland, Italy, Portugal and Spain. We choose this group of countries as it includes the countries most affected by the euro sovereign debt crisis, as well as Germany, which serves as the near-risk-free reference country, and France, which we consider as a low-risk control country. Further, we include Austria, the Czech Republic, Hungary, Poland and Slovakia in our sample, where Austria serves as a reference country for the Czech Republic.

According to Gyntelberg et al. (2013) when one considers the number of quotes of CDS contracts at the peak of the sovereign debt crisis in 2010, the five-year segment is the most liquid. The use of intraday data in our empirical analysis enables us to obtain much sharper estimates and clearer results with respect to market mechanisms as also shown in Gyntelberg et al. (2013). Further, Gyntelberg et al. (2013) show that sovereign credit risk dynamics follow an intraday pattern.

Our sovereign bond price data is provided by MTS (Mercato Telematico dei Titoli di Stato). The MTS data comprise both actual transaction prices and binding bid-offer quotes. The number of transactions of sovereign bonds on the MTS platform is, however,

[3] As presented in Gyntelberg et al. (2013). They discuss the advantages of using intraday data due to the higher accuracy of the results as compared with lower-frequency data.

insufficient to allow us to undertake any meaningful intraday analysis. Therefore, we use the trading book from the respective domestic MTS markets.[4] The MTS market is open from 8:15 to 17:30 local Milan time, preceded by a pre-market phase (7.30 to 8.00) and an offer-market phase (8:00 to 8:15). We use data from 8:30 to 17:30.[5]

The CDS data consist of price quotes provided by CMA (Credit Market Analysis Ltd.) Datavision. CMA continuously gathers information on executable and indicative CDS prices directly from the largest and most active credit investors. After cleaning and checking the individual quotes, CMA applies a time- and liquidity-weighted aggregation so that each reported bid and offer price is based on the most recent and liquid quotes. The CDS market, which is an OTC market, is open 24 hours a day. However, most of the activity in the CMA database is concentrated between around 7:00 and 17:00 London time. As we want to match the CDS data with the bond market data, we restrict our attention to the period from 8:30 to 17:30 CET (CEST during summer).

We construct our intraday data on a 30-minute sampling frequency on our data set, which spans from January 2008 to end-December 2011. The available number of indicative quotes for CDS does not allow a data frequency higher than 30 minutes. The euro area sovereign CDS markets were very thin prior to 2008, which makes any type of intraday analysis before 2008 impossible. Microstructural noise effects may come into play when high frequency data is used (Fulop and Lescourret; 2007). However, this does not apply to our data based on a 30-minute sampling frequency because we average the reported quotes over each 30-minute interval (for tests, robustness checks and for a more detailed discussion please refer to Gyntelberg et al. (2013)).

When implementing our analysis we split the data into two subsamples. The first covers the period January 2008 to 19 October 2009 and, as such, represents the period prior to the euro area sovereign debt crisis. While this period includes the most severe phase of the financial crisis, including the default of Lehman Brothers, it is relatively unaffected by market distortions stemming from concerns about the sustainability of public finances in view of rising government deficits and therefore represents the pre-sovereign debt crisis period. The second subsample covers the euro area sovereign debt crisis period and runs from 20 October 2009 to end-December 2011. As the beginning of the crisis period, we designate 20 October 2009, when the new Greek government announced that official statistics on Greek debt had previously been fabricated. Instead of a public deficit

---

[4] We ignore quotes from the centralised European platform (market code: EBM), as quotes for government bonds on the centralised platform are duplicates of quotes on the domestic platforms.

[5] Due to our equidistant sampling frequency we have discard either the first or the last 15 minutes of each trading day. We analysed the data quality in both intervals and found, on balance, that the first 15 minutes have a lower data quality than the interval at the end of the trading day. Consequently, we discarded the first 15 minutes of each trading day.

estimated at 6% of GDP for 2009, the government now expected a figure at least twice as high.

We employ CDS and bond data in our analysis in order to be able to differentiate between the transmission of sovereign risk contagion according to the credit risk channel from one country to another. Duffie (1999) argues that, since the CDS and the bond yield spread both price the same default of a given reference entity, their price should be equal if markets are perfect and frictionless. Thus, in a perfect market, due to arbitrage, the CDS spread equals the bond yield over the risk-free rate. However, for this parity to hold, a number of specific conditions must be met, including that markets are perfect and frictionless, that bonds can be shorted without restrictions or cost and that there are no tax effects, etc. According to Duffie (1999) par floating rate notes (FRN) should be used for the bond yield computation, because FRN, unlike plain vanilla bonds, carry only credit rate risk and no interest rate risk. However, they are relatively uncommon, in particular for sovereign entities. A further complication linked to the use of fixed-rate or plain vanilla bonds as substitutes is that it is unlikely that the maturity of these instruments exactly matches that of standard CDS contracts.

To ensure proper comparability with CDS, Gyntelberg et al. (2013) employ synthetic par asset swap spreads (ASW) for the bond leg of the basis. The use of ASW is in line with the practice used in commercial banks when trading the CDS-bond basis. By calculating ASW for our empirical analysis, we ensure an accurate cash flow matching, as opposed to studies that use simple "constant maturity" yield differences for credit risk.

An asset swap is a financial instrument that exchanges the cash flows from a given security - e.g. a particular government bond - for a floating market rate.[6] This floating rate is typically a reference rate such as Euribor for a given maturity plus a fixed spread, the ASW. This spread is determined such that the net value of the transaction is zero at inception. The ASW allows the investor to maintain the original credit exposure to the fixed rate bond without being exposed to interest rate risk. Hence, an asset swap on a credit risky bond is similar to a floating rate note with identical credit exposure, and the ASW is similar to the floating-rate spread that theoretically should be equivalent to a corresponding CDS spread on the same reference entity. Specifically, the ASW is the

[6] See O'Kane (2000) and Gale (2006) for detailed discussions of the mechanics and pricing of asset swaps.

fixed value A required for the following equation to hold[7] (O'Kane; 2000):

$$\underbrace{100 - P}_{\substack{\text{Upfront payment for bond} \\ \text{asset in return for par}}} + \overbrace{C \underbrace{\sum_{i=1}^{N_{\text{fixed}}} d(t_i)}_{\text{Fixed payments}} = \underbrace{\sum_{i=1}^{N_{\text{float}}} \big(L_i + A\big) d(t_i)}_{\text{Floating payments}}}^{\text{Interest rate swap}}, \tag{1}$$

where P is the full (dirty) price of the bond, C is the bond coupon, L_i is the floating reference rate (e.g. Euribor) at time t_i and $d(t_i)$ is the discount factor applicable to the corresponding cash flow at time t_i.

In order to compute the ASW A, several observations and simplifications have to be made. First, in practice it is almost impossible to find bonds outstanding with maturities that exactly match those of the CDS contracts and second, the cash-flows of the bonds and the CDS will not coincide. To overcome these issues, in what follows we use synthetic asset swap spreads based on estimated intraday zero-coupon sovereign bond prices. Specifically, for each interval and each country, we estimate a zero-coupon curve based on all available bond price quotes during that time interval using the Nelson and Siegel (1987) method. With this procedure, we are able to price synthetic bonds with maturities that exactly match those of the CDS contracts, and we can use these bond prices to back out the corresponding ASW. As this results in zero coupon bond prices, we can set C in Equation (1) to zero.

A CDS contract with a maturity of m years for country j in time interval k of day t, denoted as $S_j(t_k, m)$, has a corresponding ASW $A_j(t_k, m)$:

$$100 - P_j(t_k, m) = \sum_{i=1}^{N_m} \big(L_i(t_k) + A_j(t_k, m)\big) \cdot d(t_k, t_i), \tag{2}$$

with $P_j(t_k, m)$ as our synthetic zero coupon bond price.

For the reference rate L_i in Equation (2), we use the 3-month Euribor forward curve to match as accurately as possible the quarterly cash flows of sovereign CDS contracts. We construct the forward curve using forward rate agreements (FRAs) and euro interest rate swaps. We collect the FRA and swap data from Bloomberg, which provides daily (end-of-day) data. 3-month FRAs are available with quarterly settlement dates up to 21 months ahead, i.e. up to 21×24. From two years onwards, we bootstrap zero-coupon swap rates from swap interest rates available on Bloomberg and back out the corresponding implied forward rates. Because the swaps have annual maturities, we use a cubic spline to generate

[7] This assumes that there is no accrued coupon payment due at the time of the trade; otherwise, an adjustment factor would need to be added to the floating payment component.

the full implied forward curve, thereby enabling us to obtain the quarterly forward rates needed in Equation (2).

Given our interest in intraday dynamics, we follow Gyntelberg et al. (2013) and generate estimated intraday Euribor forward rates by assuming that the intraday movements of the Euribor forward curve are proportional to the intraday movements of the German government forward curve.[8] To be precise, for each day, we calculate the difference between our Euribor forward curve and the forward curve implied by the end-of-day Nelson-Siegel curve for Germany.[9] We then keep this difference across the entire curve fixed throughout that same day and add it to the estimated intraday forward curves for Germany earlier on that day to generate the approximate intraday Euribor forward curves. This approach makes the, in our view, reasonable assumption that the intraday variability in Euribor forward rates will largely mirror movements in corresponding German forward rates.

Finally, we need to specify the discount rates $d(t_k, t_i)$ in Equation (2). The market has increasingly moved to essentially risk-free discounting using the overnight index swap (OIS) curve. We therefore take $d(t_k, t_i)$ to be the euro OIS discount curve, which is constructed in a way similar to the Euribor forward curve. For OIS contracts with maturities longer than one year, we bootstrap out zero-coupon OIS rates from interest rates on long-term OIS contracts. Thereafter, we construct the entire OIS curve using a cubic spline. We use the same technique as described above to generate approximate intraday OIS discount curves based on the intraday movements of the German government curve.

To gauge the potential impact of this assumption on our empirical results, we reestimate our model using an alternative assumption that the Euribor and OIS curves are fixed throughout the day at their observed end-of-day values. Under this alternative assumption, we obviously fail to capture any movements in money market rates within the day when we price our synthetic asset swaps. Our results remain robust.

Please refer to Gyntelberg et al. (2013) for an in-depth discussion of the construction of our intraday ASW and to O'Kane (2000) for a general discussion.

According to different panel unit root tests (see Appendix C) our CDS and ASW price data (displayed in Figure 1) is I(1). Therefore, we estimate our subsequent models in first differences.

[8] Euribor rates are daily fixing rates, so we are actually approximating the intraday movements of the interbank interest rates for which Euribor serves as a daily benchmark.

[9] Here we use the second to last 30-minute interval, because the last trading interval is occasionally overly volatile.

Figure 1: CDS and ASW spreads in basis points

The figures are based on data with a 30-minute sampling frequency. Our split into the pre- and the crisis period is indicated by the vertical line in each figure. Due to the Greek debt restructuring the data for Greece ends in September 2011.

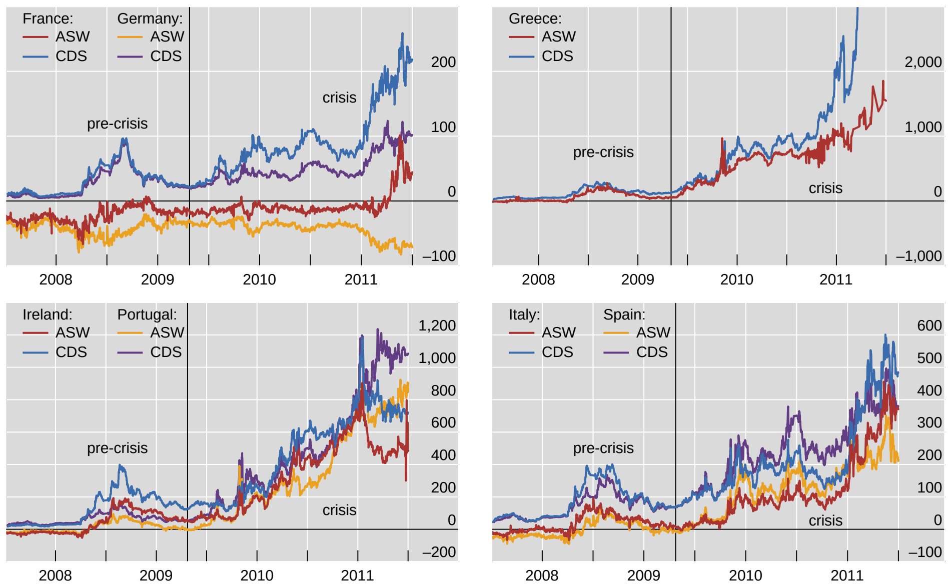

In our PVARX model extension (Section 3.2), we make use of the Citigroup economic surprise index for the euro zone[10] as an exogenous variable (see Figure 2). This index is widely recognised in academia and by practitioners for measuring unexpected economic news (such as in Goldberg and Grisse (2013), Scotti (2013), and Paulsen (2014)).

The Citigroup economic surprise index measures how economic news/data is developing relative to the anticipated consensus forecasts of market economists. According to Citigroup, the index captures objective and quantitative measures of economic news, defined as weighted historical standard deviations of data surprises (actual releases versus the Bloomberg survey median). A positive reading of the economic surprise index for the euro zone suggests that economic releases have on balance beaten the consensus, while a negative reading indicates the opposite. The index captures economic news on macroeconomic and fiscal variables such as employment change, the housing market, retail sales, debt-to-GDP, the budget deficit and consumer confidence in the euro zone. Thus, the Citigroup economic surprise index does not include news on monetary policy decisions.

[10] Bloomberg ticker: CESIEUR Index

Figure 2: Citigroup Economic Surprise Index

The figure shows the daily values of the Citigroup Economic Surprise Index (Bloomberg ticker: CESIEUR Index) from the beginning of 2008 until the end of 2011. The vertical line corresponds to 19 October 2009, the end of our pre-crisis period. Source: Bloomberg.

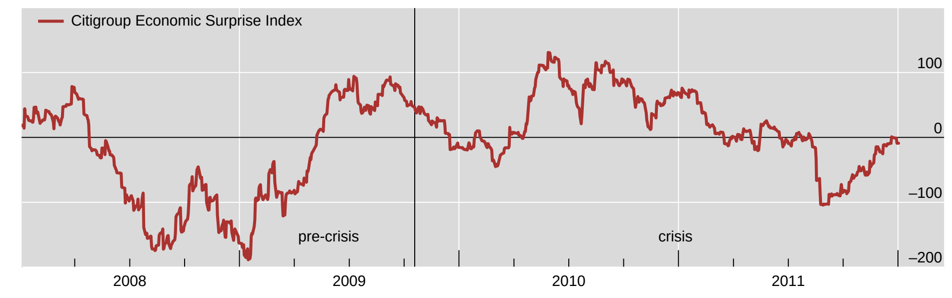

The economic surprise index has a different daily frequency from the intraday data that we are analysing in this paper. However, as market participants are exposed to economic news throughout the whole day, we disperse the actual Citigroup economic surprise index data given at the end of the trading day over that entire day. We conducted several simulations with different distributions to generate a pseudo-intraday economic surprise index. Our results remain extremely robust to this experiment. The experiment design is justified because we are interested not in the exact time line of the absorption of unexpected macroeconomic news, but rather in a qualitative picture of whether markets react to fundamental macroeconomic news in the pricing of sovereign credit risk.

3 Modelling sovereign credit spread contagion

To empirically measure the impact of euro area sovereign credit risk contagion effects according to the credit risk channel (CDS and bond market), we employ a panel vector autoregressive (PVAR) model. PVARs have the same structure as VAR models, in the sense that all variables are assumed to be endogenous but with the difference that a cross-sectional dimension is added to the representation. We define our PVAR model following Binder et al. (2005) with fixed effects when N is finite and T is large, as $i = 1, ..., N$ is the cross-sectional dimension and $t = 1, ..., T$ is the time-series dimension in our model. According to Koop and Korobilis (2016) and Canova and Ciccarelli (2013), in this setup the PVAR is the ideal tool for examining the international transmission of macroeconomic or financial shocks from one country to another.

The PVAR has several advantages over individual country VARs in a time series framework. By analysing a panel of countries, we can more accurately model contagion from one country to another since the panel approach captures country-level heterogeneity. We control for cross-sectional heterogeneity by including fixed effects in the regression. By using CDS and ASW as endogenous variables for each country in our cross-section,[11] we can differentiate the credit risk channel of contagion, which improves the understanding of sovereign risk contagion dynamics. With an extension of the PVAR using a purely exogenous variable, we can assess the effect of unexpected economic news[12] on credit risk contagion for the countries in our sample.

3.1 Panel VAR

In vector autoregressive models (VAR) all variables are treated as endogenous and interdependent in both a dynamic and static sense. The VAR model is formally defined as:

$$Y_t = A_0 + A(\mathrm{L})Y_{t-1} + u_t, \tag{3}$$

where Y_t is a G x 1 vector of endogenous variables and $A(\mathrm{L})$ is a polynomial in the lag operator, A_0 is a G x 1 vector and u_t is a G x 1 vector of i.i.d. shocks.

Panel VARs (PVAR) have the same structure as VAR models in Equation (3), as all variables are assumed to be endogenous and independent. However, a cross-sectional dimension i, in our case across countries, is added to the representation. Thus, Y_t is the stacked version of y_{it}, the vector of G variables for each country $i = 1, ..., N$, i.e. $Y_t = (y_{1t}^{'}, y_{2t}^{'}, ..., y_{Nt}^{'})^{'}$ and $t = 1, ..., T$. The major difference between a VAR and the PVAR is that the covariance σ_{ij} of the residuals is zero by definition for country i different from country j in a VAR model. The PVAR is defined as follows:

$$y_{it} = A_{0i} + A_i(\mathrm{L})Y_{t-1} + u_{it}. \tag{4}$$

A_{0i} are G x 1 vectors and A_i are G x GN matrices. We allow for country-specific heterogeneity by including a country-specific intercept. Further, lags of all endogenous variables of all entities enter the equation of country i. Canova and Ciccarelli (2013) call this feature "dynamic interdependencies". The residual u_{it} is a G x 1 vector and $u_t = (u_{1t}, u_{2t}, ..., u_{Nt})$. u_{it} is generally correlated across the cross-sectional dimension i. Canova and Ciccarelli (2013) call this feature "static interdependencies". Thus the variance-covariance matrix for a PVAR has the following property $E(u_{it}u_{jt}') = \sigma_{ij} \neq 0$ for $i \neq j$, i.e. static interdependencies occur when the correlations between the errors in two countries' VARs are non-zero. On the other hand, dynamic interdependencies occur

[11] bi-variate estimation per country

[12] By using the economic surprise index as a predetermined purely exogenous variable in the PVARX model.

when one country's lagged variables affect another country's variables. Hence, the PVAR is more flexible compared to a VAR ($\sigma_{ij} = 0$ for $i \neq j$).[13]

In our bivariate case, i.e. $G = 2$, we can rewrite the PVAR in Equation (4) as:

$$\begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{it} = \begin{pmatrix} A_{01} \\ A_{02} \end{pmatrix}_{i} + \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}_{ji} (\mathrm{L}) \begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{j\,t-1} + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}_{it}, \tag{5}$$

where A_{mn} are scalars and $i, j = 1, ..., N$ is the number of countries in the cross-sectional dimension.

For the estimation, we follow the approach proposed by Canova and Ciccarelli (2009) of an unrestricted PVAR which allows for the selection of restrictions involving dynamic interdependencies, static interdependencies and cross-section heterogeneities.[14] According to an empirical model comparison by Koop and Korobilis (2016), the proposed methodology by Canova and Ciccarelli (2009) shows the best properties compared to other PVAR approaches. Canova and Ciccarelli (2009) suggest adopting a flexible structure through a factorisation of the coefficients in Equation (4). Through the flexible coefficient factorisation, the PVAR can be rewritten as a reparametrised multicountry VAR and estimated using SUR (Canova and Ciccarelli; 2009). The advantage of this flexible factorisation is that the overparametrisation of the original PVAR is dramatically reduced while, in the resulting SUR model, estimation and specification searches are constrained only by the dimensionality of the estimated coefficient matrix (for a more in-depth discussion please refer to Canova and Ciccarelli (2009) and Koop and Korobilis (2016)).

3.2 Panel VARX

As an extension to the previous analysis, we consider the response of credit risk in CDS and bond markets in our GIIPS and low-risk country sample to unexpected macroeconomic news. We follow Canova and Ciccarelli (2013) by extending the PVAR model in Equation (4) with a predetermined purely exogenous variable X_t which results in a PVARX model which takes the following form:

$$y_{it} = A_{0i} + A_i(\mathrm{L})Y_{t-1} + F_i(\mathrm{L})X_t + u_{it}, \tag{6}$$

[13] According to Canova and Ciccarelli (2013), these features distinguish a panel VAR typically used for macroeconomics and finance from a panel VAR used in microeconomics.

[14] We use demeaned and standardised first differences.

with X_t as a M x 1 vector (M is equal to the number of exogenous variables) common to all entities i. The PVARX can also be rewritten as:

$$\begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{it} = \begin{pmatrix} A_{01} \\ A_{02} \end{pmatrix}_{i} + \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}_{ji} (\mathrm{L}) \begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{j\,t-1} + \begin{pmatrix} \tilde{F}_1 \\ \tilde{F}_2 \end{pmatrix}_{i} (\mathrm{L}) X_t + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}_{it} . \tag{7}$$

We employ the economic surprise index as a predetermined exogenous variable X_t for unexpected macroeconomic news in the euro zone, i.e. $M = 1$.

The extension to the PVARX model allows us to analyse whether credit risk responses can be attributed to macroeconomic fundamental news or if exaggerations in terms of lack of belief of economic agents also contributed to credit risk responses.

4 Results

We carry out an impulse response analysis to investigate contagion of sovereign credit risk across the euro area countries that were most affected in the sovereign debt crisis. Further, we present results on shock contagion to central European countries. We focus on individual country shocks propagating from GIIPS countries and analyse the impact of an unexpected one-unit shock to credit risk in both the CDS and ASW markets from country i to j. Finally, we present results of exogenous economic news shocks and the effect on sovereign risk in GIIPS countries.

In standard VAR models (see Equation (3)), shock identification is performed by imposing a Choleski decomposition in all countries. To reduce the number of identification restrictions in a VAR model, it is assumed that $E(u_{it}u'_{jt})$ is block diagonal, with blocks corresponding to each country. Canova and Ciccarelli (2009) state that block diagonality implies differences in the responses within and across countries. Within a country, variables are allowed to move instantaneously. But across entities, variables can only react with one lag.

The identification of shocks for PVAR models as defined in Equation (4) is more complicated, given that the PVAR model allows for static interdependencies, as u_{it} is correlated across entities i. Thus, cross-entity symmetry in shock identification cannot be assumed. We compute the impulse responses following Canova and Ciccarelli (2009) as the difference between two conditional forecasts: one where a particular variable is shocked and one where the disturbance is set to zero. For a more in-depth discussion of shock identification using conditional forecasts in PVAR models allowing for static interdependencies please refer to Canova and Ciccarelli (2009).

4.1 Results for GIIPS and low-risk countries

As a general result, we find that, pre-crisis, the bond and CDS markets are of similar importance, i.e. the response function of country i to a one-unit shock to the ASW and CDS markets of country j is of a comparable size in the two markets (see Figure 3). These results are as expected, as both markets should price the countries' credit risk equally (Duffie; 1999). During the crisis period, the CDS market becomes more relevant on balance (see Figure 4). Interestingly, the inter-market shock transmission, i.e. from CDS to ASW and vice versa, is not important during the pre-crisis period. This weak connection between the two markets during the pre-sovereign debt crisis period can be explained by different market participants and their distinct investment horizons. Insurance firms active in the bond market have a longer investment horizon than, for example, hedge funds in the CDS markets. During the crisis period, shock transmission between markets becomes relatively more important, suggesting a stronger inter-market connectivity. Market participants get more vigilant to potential bad news, which may spill over from other markets.

Further, we find that the decay of a shock is faster on average in the pre-crisis period than in the crisis period (see Figures 3 and 4). The timelines of our estimated shock contagion and absorption are dramatically shorter than in existing empirical studies, such as Koop and Korobilis (2016), who find that shock contagion spreads on average within one to two months in the case of shocks that do not decay over a timeline of 10 months. We find for both sample periods that contagion propagates immediately within the first 30-minute time interval. Therefore, responses to shock contagion are typically not lagged as found, for example, in Koop and Korobilis (2016). Further, the average response for shock absorption is around one hour in the pre-crisis period and slightly longer at one to two hours on average during the crisis period. This result is clearly in line with the generally accepted notion that financial markets react very fast to new information (Gyntelberg et al.; 2013). The slower speed of shock absorption during the crisis seems to contradict our statement above that market participants are more reactive to news during crisis periods. This could be explained by the fact that the estimated timeline of shock absorption during the crisis period is strongly affected by turmoil in financial markets, while the pre-crisis period represents a relative normal market environment for European sovereign states without fast and furious shock contagion but rather with comovements across markets as defined by Constancio (2011) and Forbes (2012).

In the pre-crisis period, a credit risk shock spreading from the ASW to the CDS market and vice versa had more or less the same impact in terms of magnitude and shock absorption. Thus, the derivatives market and the spot market were about equally significant in terms of shock contagion prior to the euro area sovereign debt crisis. However, during the crisis period we find that shock transmission from the ASW to the CDS market had a dramatically lower impact than vice versa. This leads to the assumption that the importance of the spot market as a channel of financial shock contagion decreased during the euro area sovereign debt crisis. Thus, the contagion of shocks to credit risk has been transmitted predominantly through the derivatives market.

During the pre-crisis period, a one-unit shock to either the ASW or CDS of country i results in a spread widening for all countries. However, during the crisis, we find evidence of a flight to safety to German bonds, as Germany is considered a safe haven for investors. This effect is visible in the inter-market connection, i.e. a positive shock to a GIIPS country's CDS or ASW leads to spread tightening in German ASW, while we cannot report a similar effect for German CDS. Similar behaviour is not visible for France, despite it being considered a low-risk control country.

During the pre-crisis period, we find that the magnitude of the impulse responses is similar across all countries, while during the crisis period, GIIPS countries exhibit much larger impulse responses than the rest of our sample countries do.

The forecasting precision is much more accurate during the crisis period, as the confidence bands are much tighter than in the pre-crisis period.

In contrast to the other empirical studies using this methodology, Koop and Korobilis (2016) find confidence bands for their impulse responses that all lie between positive and negative reactions to a one-unit shock to Greek bond yields relative to Germany. The advantage of our approach, using ASW and intraday data, dramatically increases the precision of the results during the crisis period.

In addition to the impulse response functions for a shock to Greek ΔASW and ΔCDS in Figures 3 and 4, we present impulse response functions for a shock to Spanish and Portuguese ASW and CDS in Appendix A[15].

[15] Impulse response functions for a shock to Irish and Italian ASW and CDS show similar results and can be provided on request.

Figure 3: Impulse responses in pre-crisis period - shock in Greece

This figure illustrates the impulse response for ΔCDS and ΔASW to a one-unit shock (increase) for the period from January 2008 to 19 October 2009. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS or ΔASW. The number of 30-minute time intervals is described by the x-axis. For each impulse response we plot the upper and lower 95% confidence bands.

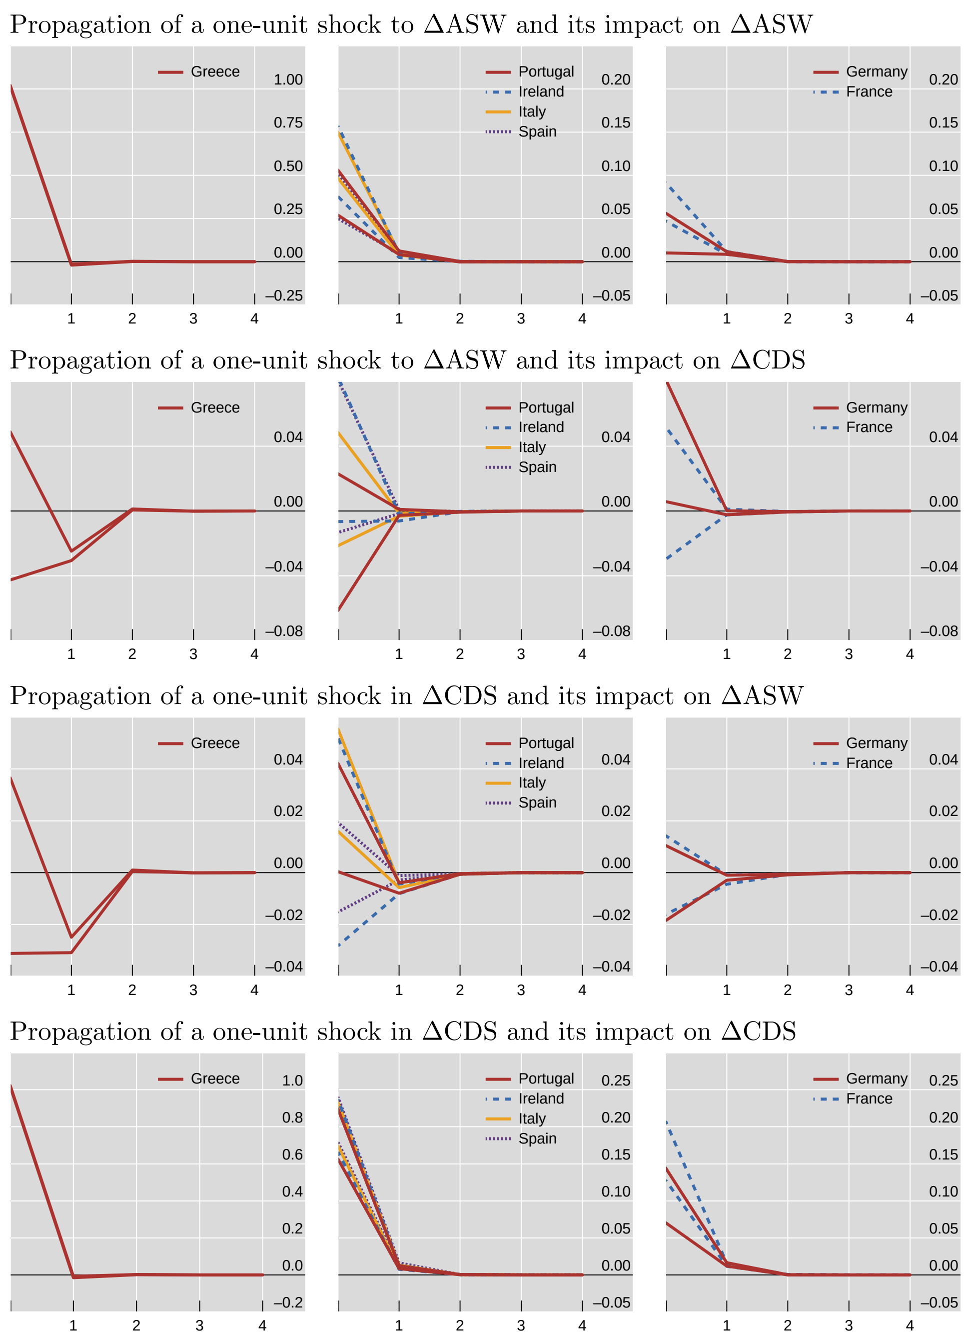

Figure 4: Impulse responses in crisis period - shock in Greece

This figure illustrates the impulse response for ΔCDS and ΔASW to a one-unit shock (increase) for the period from 20 October 2009 to end-December 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS or ΔASW. The number of 30-minute time intervals is described by the x-axis. For each impulse response we plot the upper and lower 95% confidence bands.

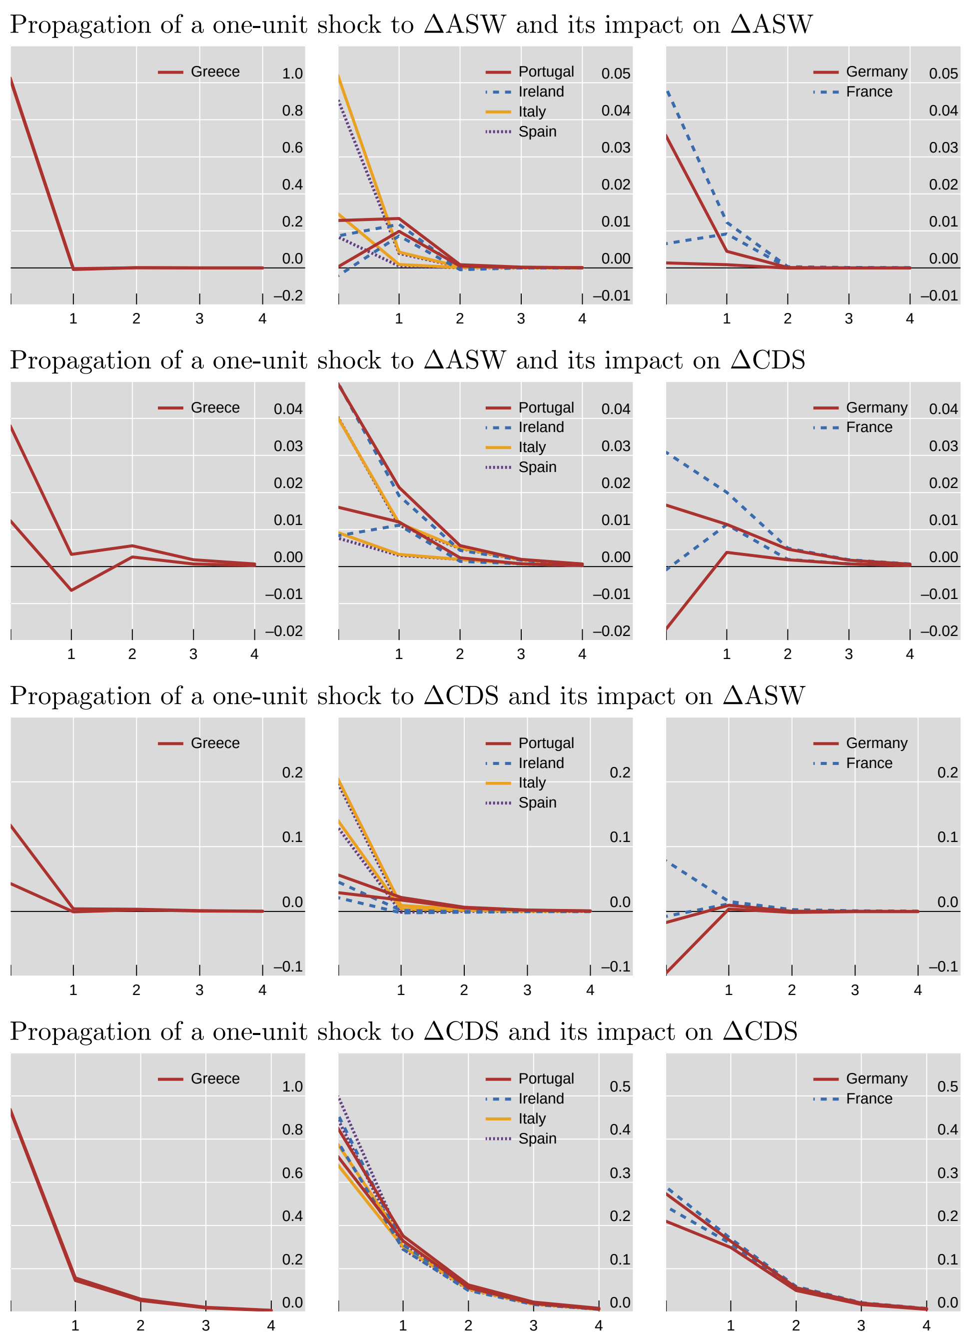

Even though policy makers may not be interested in short-lived intraday movements in sovereign credit risk, our results show that the level impacts from the short-term dynamics are persistent (see Appendix D). Hence, our results are important with regard to financial stability.

4.2 Results for central European countries during the crisis period

This section presents the results of an unexpected one-unit shock to CDS credit risk propagating from GIIPS countries and the shock response in central European countries. Due to the illiquidity of the bond markets in the central European countries in our sample, we were only able to conduct an intraday analysis based on CDS data during the crisis period. This, however, does not limit the validity of our analysis, because the results for GIIPS and low-risk countries in Section 4.1 strongly indicate that bond markets were not the main venue of sovereign credit shock contagion during the crisis. Thus, the PVAR model in Equation (4) applied in this section is estimated with $G = 1$.

Figure 5: CDS spreads in basis points

The figures are based on data with a 30-minute sampling frequency. Our split into the pre- and crisis period is indicated by the vertical line in each figure.

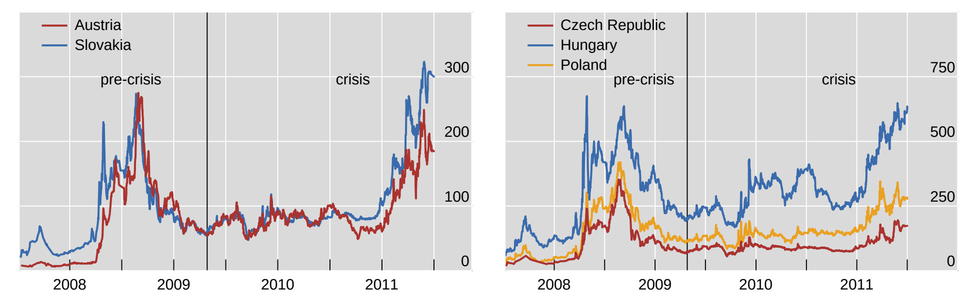

We find that the central European countries in our sample were much less affected by shocks compared to the GIIPS countries during the euro area sovereign debt crisis. We do not find differences in the impulse responses for central European euro area member countries (Austria and Slovakia) and non-euro member countries (see Figure 6, lower panel). Interestingly, we also do not find a difference in the response functions according to the debt-to-GDP levels of central European countries. The level of response for the central European countries (see Figure 6, lower panel) is almost identical. We would have expected a stronger response to shocks in central European countries with higher debt levels, such as Hungary (see Table 1).

Table 1: Debt-to-GDP levels in percent, market adjusted (Source: National data, authors' calculations)

year	Austria	Czech Republic	Hungary	Poland	Slovakia
2008	67.4	27.2	65.5	44.0	29.2
2009	77.7	32.2	77.2	49.1	36.6
2010	88.3	39.1	82.6	53.1	41.6
2011	87.1	40.8	80.7	54.9	44.9
2012	91.8	47.6	79.4	54.8	54.4
2013	91.4	49.8	82.2	56.4	57.3
2014	93.3	49.2	86.7	49.2	56.6

Figure 6: Impulse responses in central European countries in crisis period - shock in Greece

This figure illustrates the impulse response of central European countries for ΔCDS to a one-unit shock (increase) in Greece for the period from 20 October 2009 to end-December 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS. The number of 30-minute time intervals is described by the x-axis. For each impulse response we plot the upper and lower 95% confidence bands.

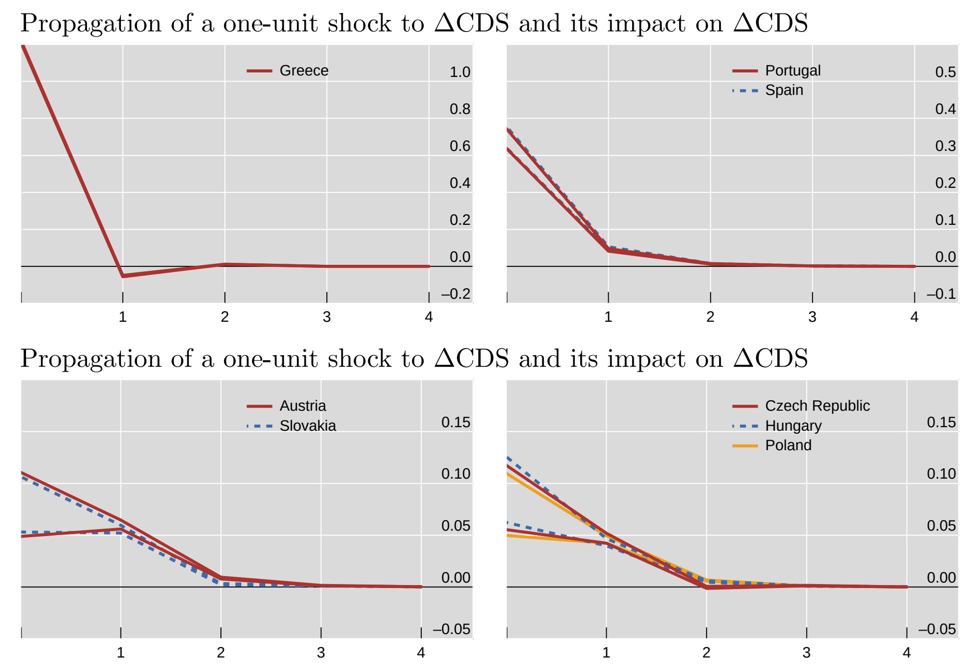

Claeys and Vasicek (2014) are in line with our results, as they also find substantial contagion effects only between countries that were most affected by the euro area sovereign debt crisis.

This leads to the conclusion that countries that lie geographically outside the crisis region are dramatically less sensitive to shocks propagating from the euro zone crisis regions. The speed of shock absorption is similar to that found for the GIIPS countries in our bivariate PVAR model discussed in Section 4.1.

Further impulse responses to shocks to Portuguese and Spanish CDS and their impact on central European countries can be found in Appendix B.

4.3 The impact of unexpected macroeconomic news on sovereign credit risk: Results from a PVARX experiment

In this section, we conduct an experiment with the aim of analysing whether responses to shocks and shock contagion can be attributed to economic fundamentals or if overreactions in credit risk during the crisis period might also be due to self-fulfilling prophecies. For Constancio (2011) and Forbes (2012), contagion occurs when financial or macroeconomic imbalances (shocks) create a systemic risk beyond that explained by economic fundamentals. Contagion differs from macroeconomic interdependence among countries in that transmission of risk to other countries is different under normal economic conditions. Gibson et al. (2012) explain the effect of self-fulfilling prophecies by interest rate spreads that were lower than justified by fundamentals prior to the crisis, owing to the role played by Greece's euro area membership in biasing investor expectations. During the crisis period, Gibson et al. (2012) define this self-fulfilling prophecy effect that interest rate spreads were higher than those predicted by fundamentals in terms of the market's disbelief that sustainable financial consolidation measures and structural reforms would be implemented.

Our experiment is designed in a similar way to that of Canova and Ciccarelli (2009), as follows: we distribute the data of the economic surprise index over each trading day (18 time intervals). The distribution is chosen such that the maximum is reached at noon, and the sum of the 18 different intraday values is equal to the value reported by the Citigroup economic surprise index. We experimented with different distributions and, despite the arbitrary distribution assumption, we found robust results. Next, the last seven values are removed from all time series in order to be close to the last maximum (in the case of a positive reading of the surprise index) or close to the last minimum (in the case of a negative reading of the surprise index). We then fit the PVARX model from Equation (6) and produce an out-of-sample forecast for eight intervals beyond the last data point,[16] which is in the case of the pre-crisis period 15 October 2009 and in the case of the crisis period 26 May 2011[17]. We call this forecast the "real forecast". Further, we repeat this same procedure, but now set the data of the surprise index of the last day to zero, i.e. we artificially remove the last positive or negative "shock" given by the data. We again produce an out-of-sample forecast, which we call the "counterfactual forecast". The difference between the real and the counterfactual forecast captures the impact of the positive or negative values of the Citigroup economic surprise index on the last day. In other words, the experiment mimics what would have happened if the last positive or negative economic news had not occurred and thus helps answer the question of whether macroeconomic fundamental news can explain changes in sovereign credit risk.

During the pre-crisis period, we find for all countries in the sample that a positive (negative) shock from the economic surprise index on the last day (15 October 2009) leads to an expected decrease (increase) in credit risk (see Figure 7). Prior to the crisis, the magnitude of the effect following an unexpected macroeconomic news shock is similar in the bond and CDS markets. Our pre-crisis results indicate that markets reacted to macroeconomic news in pricing sovereign credit risk.

During the euro area sovereign debt crisis period, a negative reading of the economic surprise index on the last available day (26 May 2011) leads surprisingly to a decrease in credit spreads in most countries (see Figure 8). In rational markets, a negative economic news shock should lead to an increase in sovereign credit risk and thus to an increase in spreads. Our results are counterintuitive, unlike those for the pre-crisis period. For the crisis period, they show that credit markets were driven not by macroeconomic news, but most likely by monetary policy, political decisions and speculations. Figure 9 displays the individual components (the real and the counterfactual forecasts) of our unexpected economic news shock experiment. Subtracting the counterfactual forecast in row 2 from the real forecast in row 1 of Figure 9 produces the forecast in row 1 of Figure 8. The same applies to the remaining rows in Figures 8 and 9. Surprisingly, in most cases a negative economic shock leads to a tightening of credit spreads (row 1 and 3 of Figure 9).

The shapes of the curves in Figures 7, 8 and 9 are due to our particular choice of decomposing the daily Citigroup economic surprise index into intraday intervals. However, other choices leading to different shapes of our curves do not change the results. This gives support to the self-fulfilling crisis theory, that changes in sovereign credit risk during the

---

16 We have chosen the forecast length of eight intervals in order to be slightly longer than the number of removed values (seven). We experimented with different values and found that the qualitative results remained robust. Again, the choice to remove the last seven values is motivated to be close to the last maximum/minimum of the surprise index, reached at midday by construction.

17 We chose end-May as the last time stamp in our experiment for the crisis period because the liquidity in the Greek bond market deteriorates. The lack of pricing data from May onwards does not allow to generate a sensible intraday forecast for our experiment.

euro area sovereign debt crisis were only partially driven by economic fundamentals, as markets did not react to economic news in contrast to the pre-crisis period.

Figure 7: Positive shock to the Economic Surprise Index during the pre-crisis period

This figure illustrates a scenario of a real positive shock to the economic surprise index minus a counterfactual scenario where we assumed that the shock did not happen. The period under consideration is from January 2008 until 15 October 2009. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

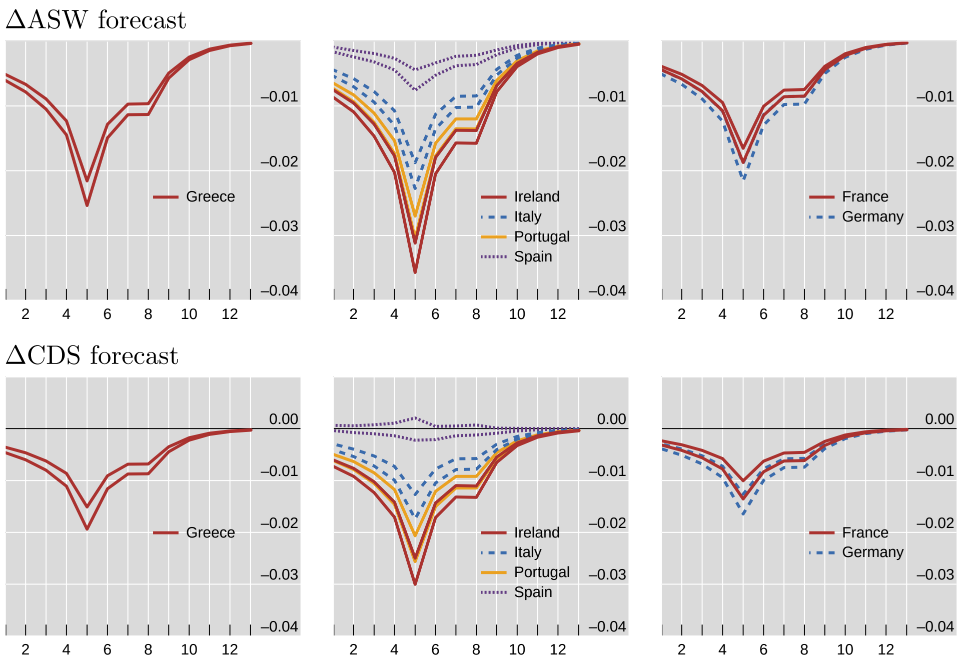

Figure 8: Negative shock to the Economic Surprise Index during the crisis period

This figure illustrates a scenario of a real negative shock to the economic surprise index minus a counterfactual scenario where we assumed that the shock did not happen. The period under consideration is from 20 October 2009 until end-May 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

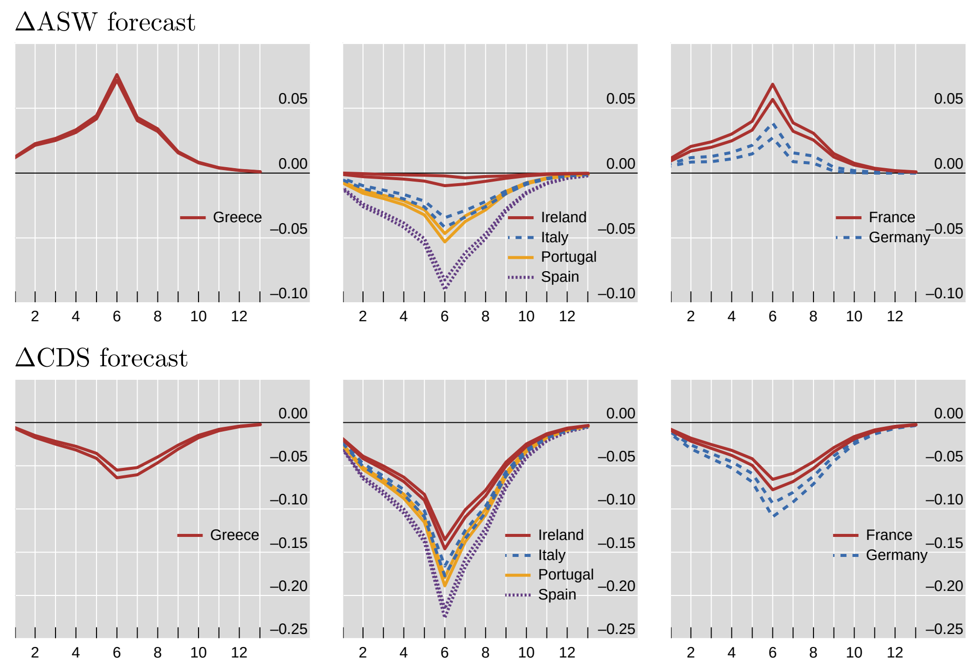

Figure 9: Real and counterfactual forecast decomposition

This figure presents the individual components, i.e. the real and the counterfactual forecasts, of our experiment for the crisis period 20 October 2009 until end-May 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

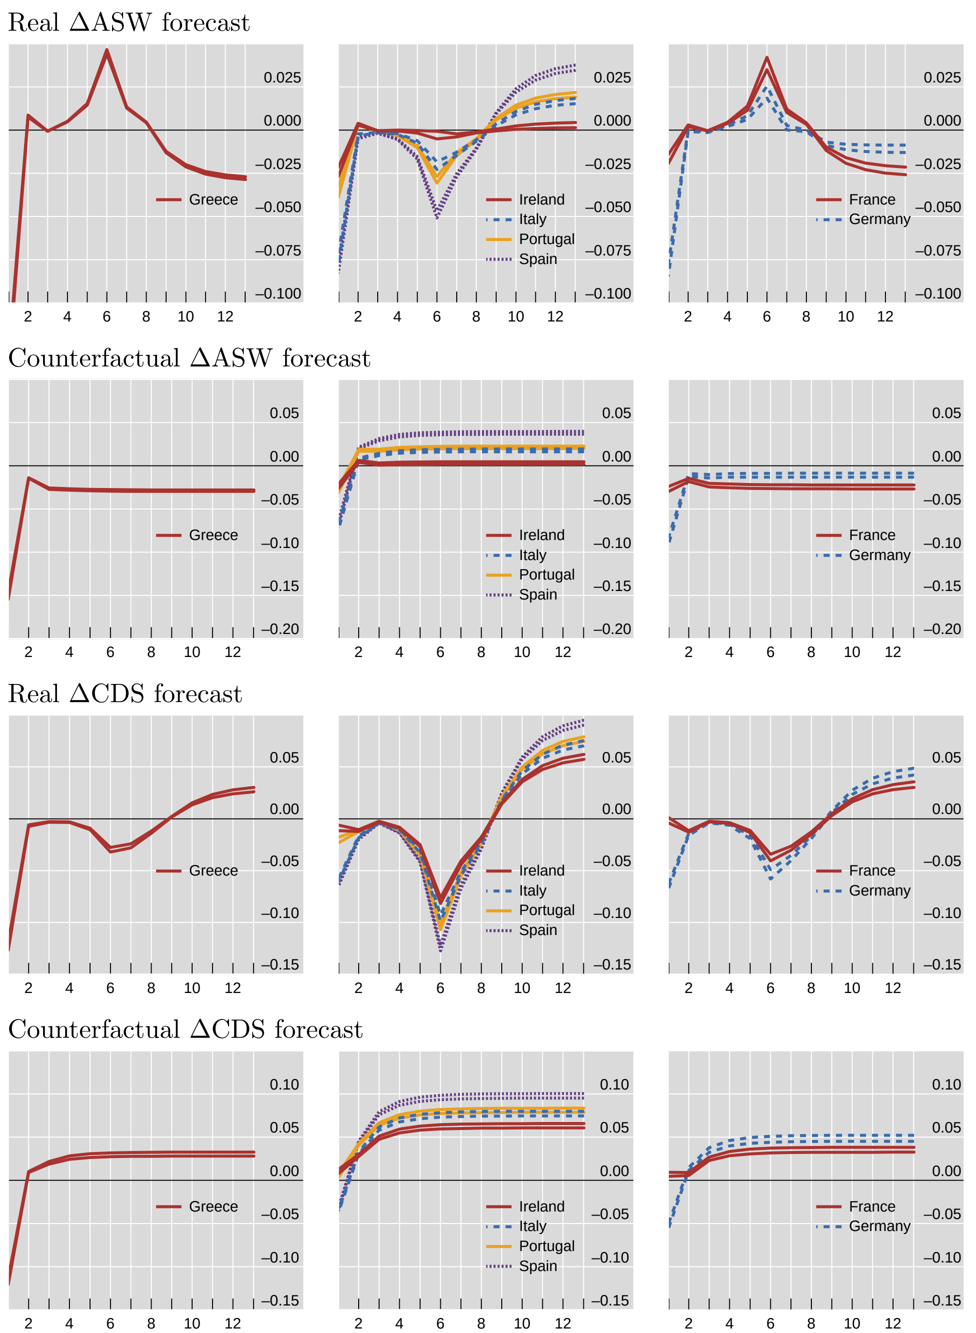

5 Conclusion

The CDS market was the main venue for the transmission of sovereign credit risk contagion during the euro area sovereign debt crisis. In contrast, we find that, prior to the crisis, the two markets (CDS and bond) were similarly important in the transmission of financial contagion, while the importance of the bond market decreased relative to the CDS market during the crisis period. We find evidence for sovereign credit risk contagion during the euro area sovereign debt crisis period, as our results show more drastic reactions to shocks in terms of magnitude and absorption compared to the pre-crisis period. Thus, our results on the responses to sovereign credit risk shocks during the crisis period confirm the contagion across euro area countries, as they result from extreme negative, systemic effects and are much larger in magnitude compared to the pre-crisis period, a fact which cannot be explained by macroeconomic fundamentals.[18] We find comovement effects rather than contagion during the pre-crisis period, as markets react rationally to economic fundamentals, while the responses to sovereign credit risk shocks remain moderate in magnitude. The use of intraday data substantially increases the precision of the results, as we find average timelines of financial shock contagion of one to two hours during the crisis period and 30 minutes to one hour prior to the crisis.

We find a flight to safety during the crisis period in the German bond market. This is not present prior to the crisis and, interestingly, is also not visible in the French bond market. The flight-to-safety effect can be explained by market participants' lack of belief in the future path of public finances (a self-fulfilling crisis), which cannot be explained by macroeconomic news.

Our results using an unexpected exogenous macroeconomic news shock suggest that, during the pre-crisis period, markets for sovereign credit risk were driven by macroeconomic news. Positive news led to a decrease in credit spreads and negative news to an increase. Using the same experiment for the euro area sovereign debt crisis period, our results show that movements in sovereign credit spreads did not respond to macroeconomic news but were rather driven by either monetary policy or exaggerations in financial markets due to lack of belief (a self-fulfilling crisis).

We find that central European countries were practically unaffected by sovereign risk contagion during the crisis. Our model further indicates no difference in the responses to shocks according to debt levels or whether the country belongs to the monetary union or not. This implies that, in general, countries that lie geographically outside of the crisis region were much less affected by sovereign risk contagion.

[18] See the contagion definitions according to Constancio (2011), Forbes (2012), and Kaminsky et al. (2003) in the Introduction.

As stated by Gyntelberg et al. (2013), the fact that CDS premia are more responsive to new information may reflect the fact that the market participants in these markets on average are more highly leveraged, are more aggressive in taking positions and hence respond more quickly to new information. Thus it is crucial for policy makers and regulators to understand the dynamics in the market for sovereign credit risk, especially in the derivative market, where contagion effects are more severe during our analysed crisis sample.

A Impulse response functions for Spain and Portugal

Figure A.1: Impulse responses in pre-crisis period - shock in Spain

For details see Figure 3.

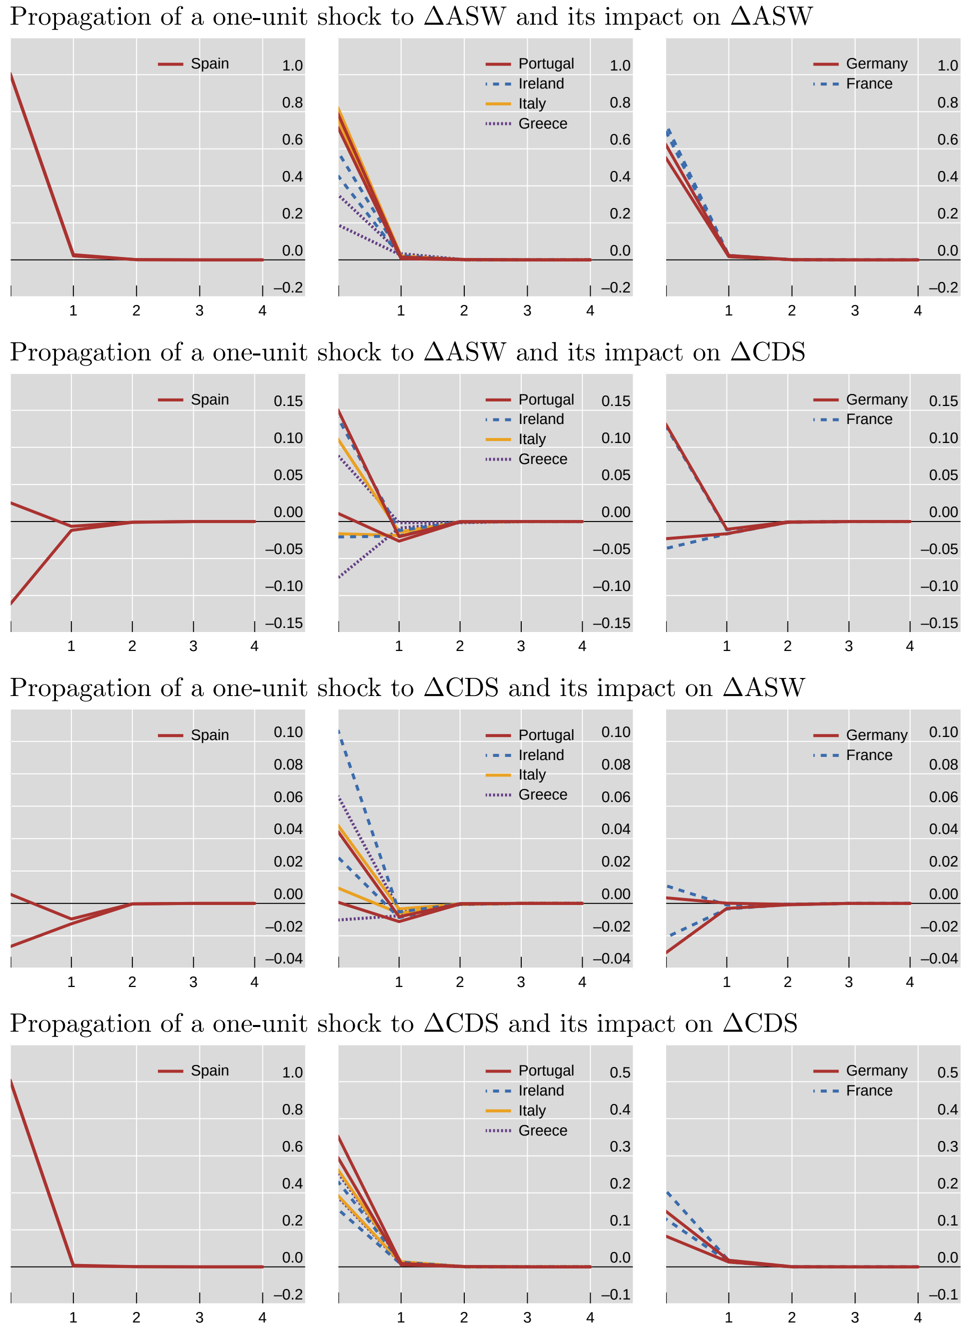

Figure A.2: Impulse responses in crisis period - shock in Spain

For details see Figure 4.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

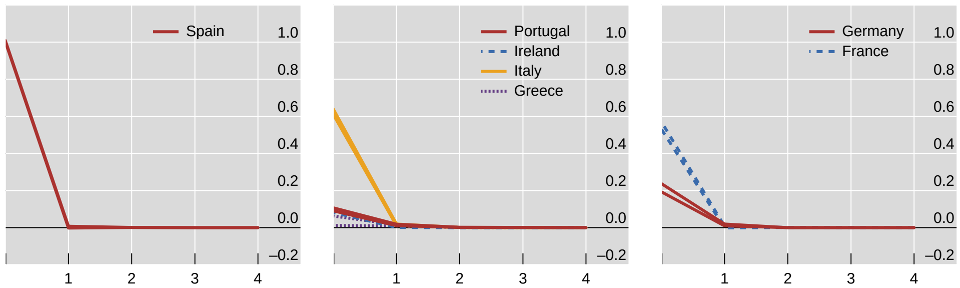

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

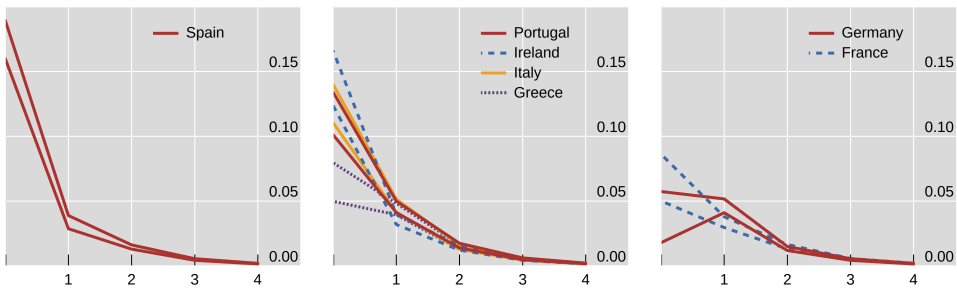

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

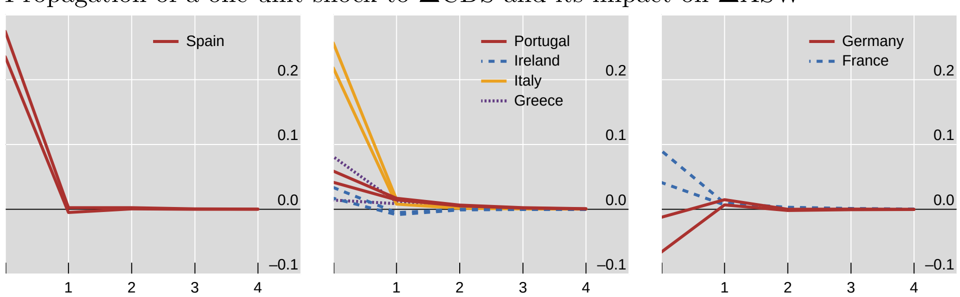

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

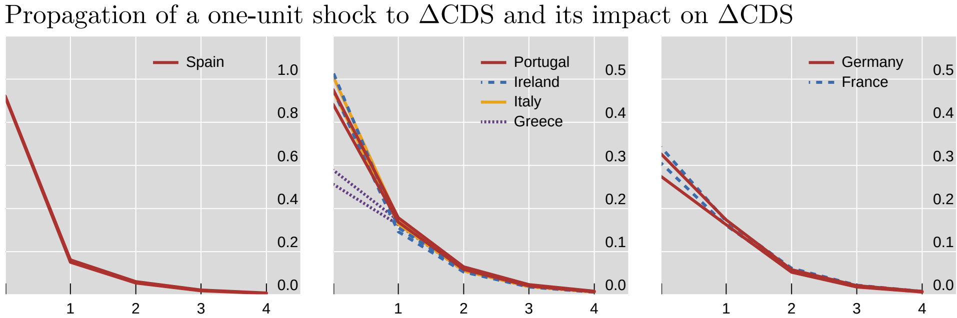

Figure A.3: Impulse responses in pre-crisis period - shock in Portugal

For details see Figure 3.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

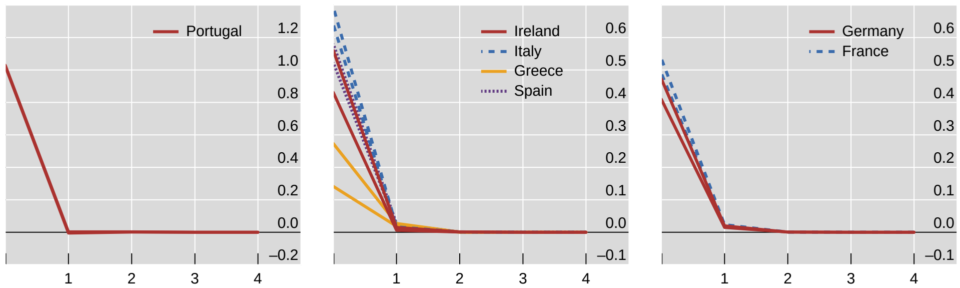

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

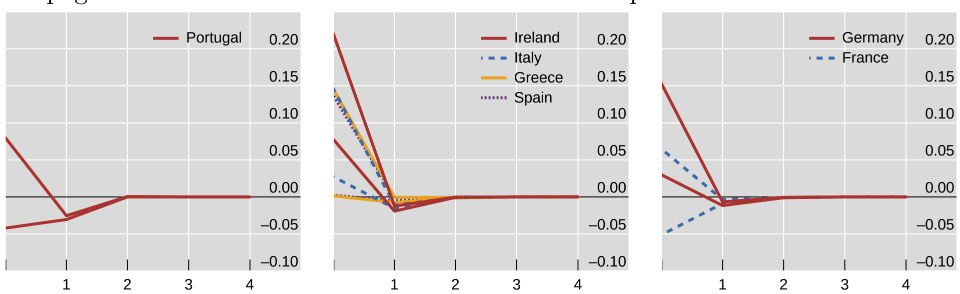

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

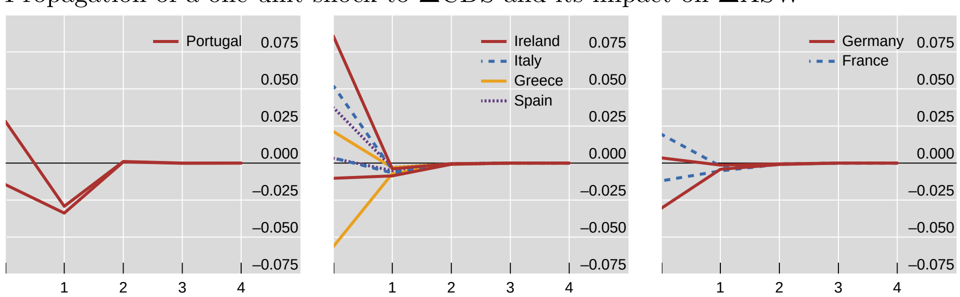

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

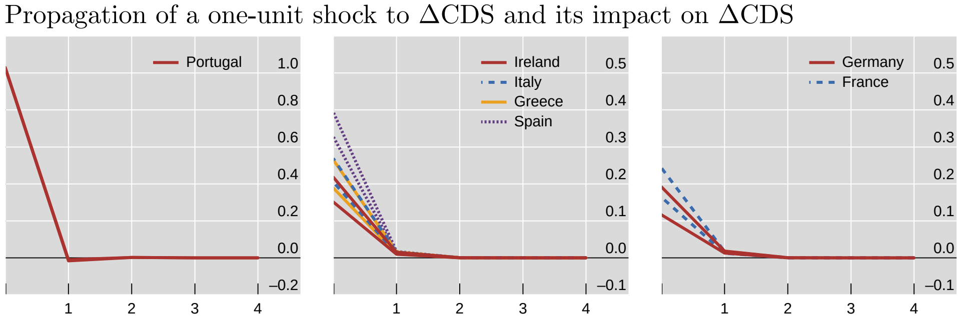

Figure A.4: Impulse responses in crisis period - shock in Portugal

For details see Figure 4.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

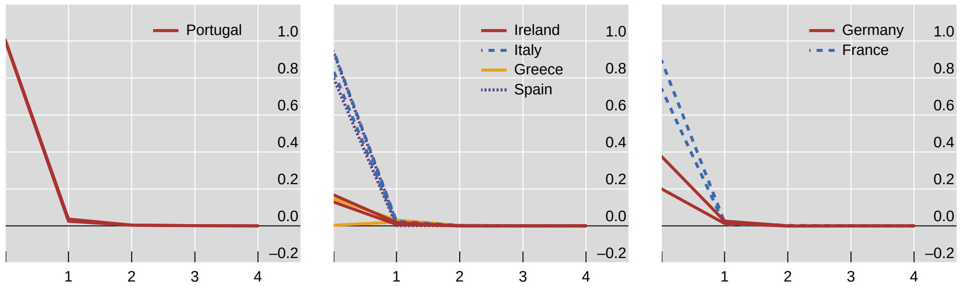

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

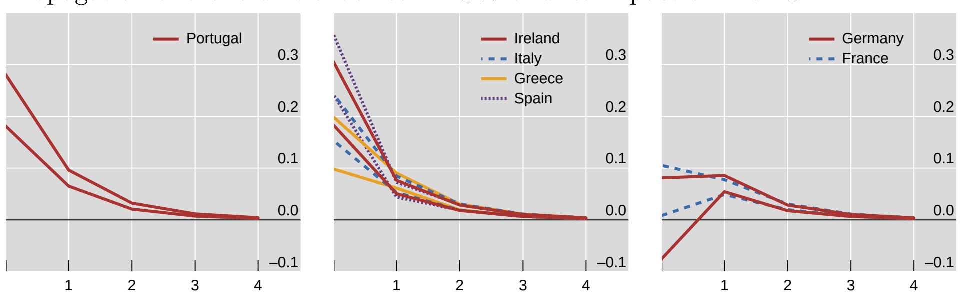

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

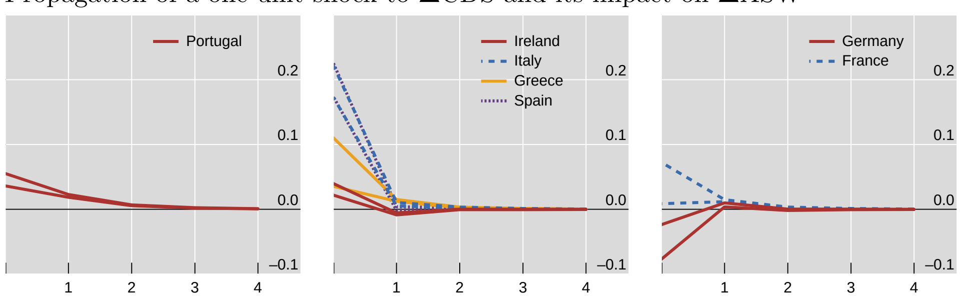

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

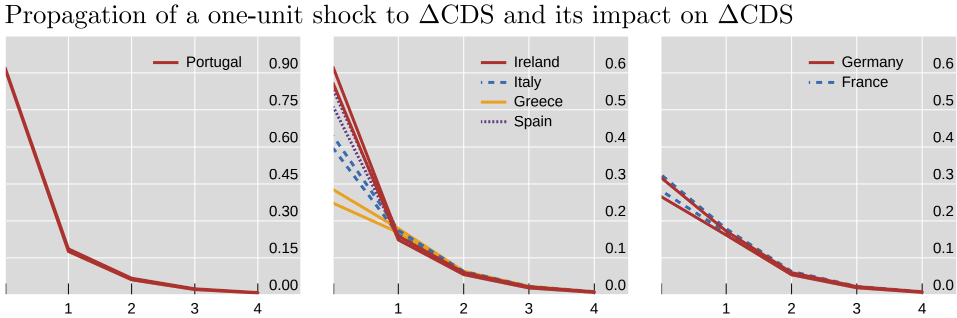

B Impulse response functions of central European countries

Figure B.1: Impulse responses in crisis period - responses of central European countries

We present the responses of central European countries to a one-unit shock to ΔCDS in Spain and Portugal. For details see Figure 6.

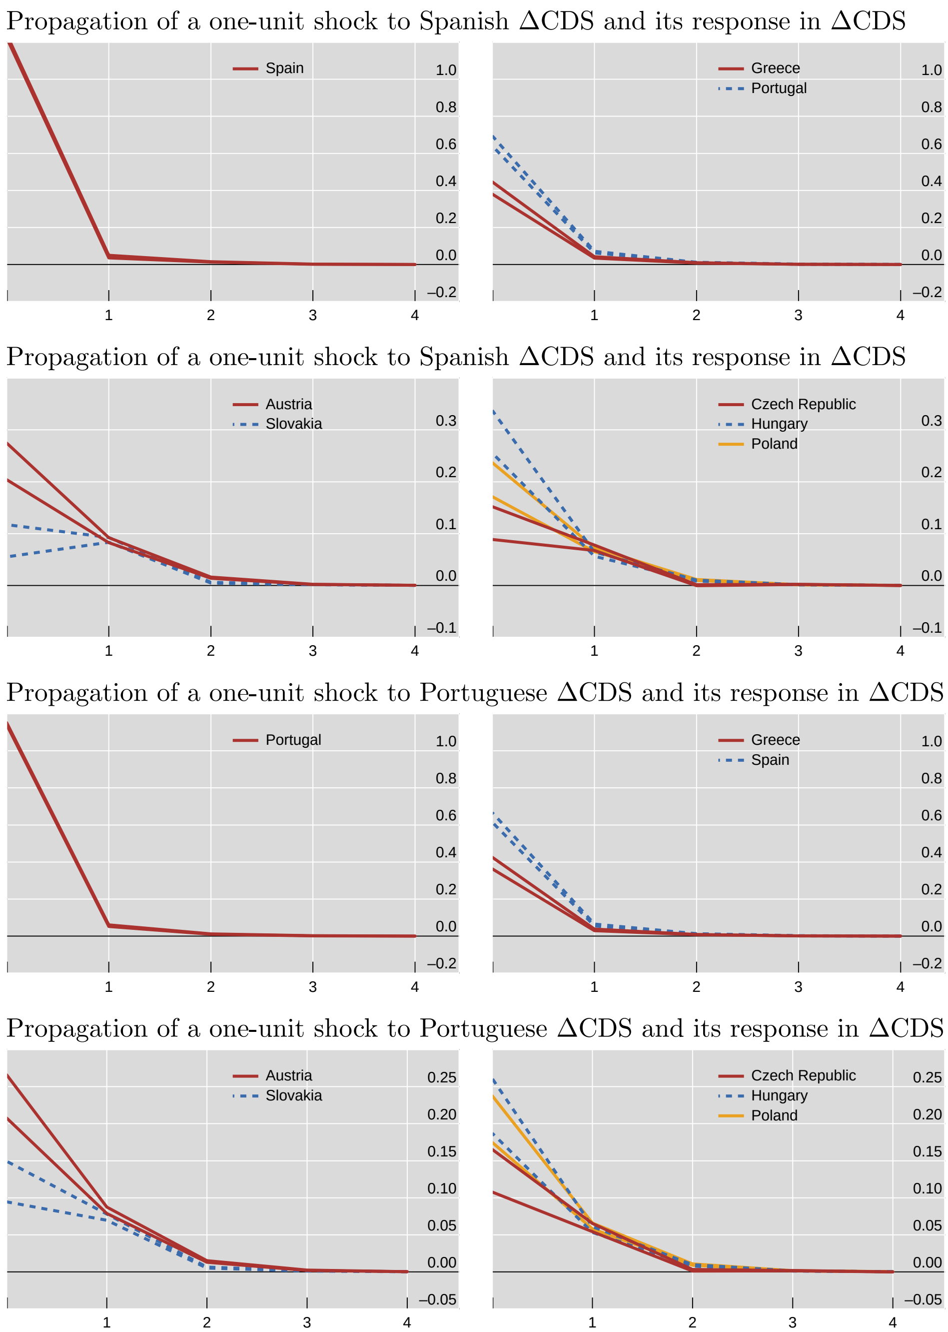

C Panel unit root

Before analysing contagion effects within a panel framework, we perform unit root and stationarity tests on our CDS and ASW price data. Canova and Ciccarelli (2013) suggest that panel-based unit root or stationarity tests have a higher power than univariate tests. For our ASW and CDS data, we cannot reject the $\mathscr{H}_0$ of a common unit root according to the Levin, Lin-, and Chu test. Further, we also can not reject the $\mathscr{H}_0$ of individual unit root processes according to the Im, Pesaran and Shin panel unit root test for our data (see Table C.1). Since all of our country series are considered simultaneously and our data for CDS and ASW are non-stationary (I(1)), we use first differences for our model estimations.

Our panel unit root test takes the following form:

$$\Delta y_{it} = \alpha_i + \rho_i y_{i,t-1} + u_{it} \qquad \text{with} \qquad \mathscr{H}_0 : \rho_1 = ... = \rho_N = 0,$$

where $i = 1, ..., N$ is the cross-sectional dimension and $t = 1, ..., T$ is the time-series dimension. Hence, all series are independent random walks under the $\mathscr{H}_0$ and non-stationary.

We perform the Levin, Lin, and Chu test, which assumes a common unit root process where the homogenous alternative takes the following form:

$$\mathscr{H}_{1a} : \rho_1 = ... = \rho_N = \rho < 0,$$

where all series are stationary under the $\mathscr{H}_1$.

Further, we perform individual panel unit root tests based on Im, Pesaran and Shin where the heterogeneous alternative takes the following form:

$$\mathscr{H}_{1b} : \rho_1 < 0, ..., \rho_{N_0} < 0, \qquad \text{where} \qquad N_0 \leq N.$$

Hence, $N_0 \leq N$ series are stationary, with potentially different AR parameters.

Table C.1: Panel unit root tests - p-values

This table reports the p-values of the panel unit root test with individual intercepts for the period from January 2008 to end-December 2011, 30-minute sampling frequency and the seven countries in our sample.

	Levin, Lin, Chu	Im, Pesaran, Shin
ASW	1.00	0.23
CDS	0.59	0.15

D Accumulated impulse response functions for Greece

Figure D.1: Accumulated impulse responses in pre-crisis period - shock in Greece

This figure illustrates the accumulated impulse response for CDS and ASW to a one-unit shock (increase) for the period from January 2008 to 19 October 2009. For further details see Figure 3.

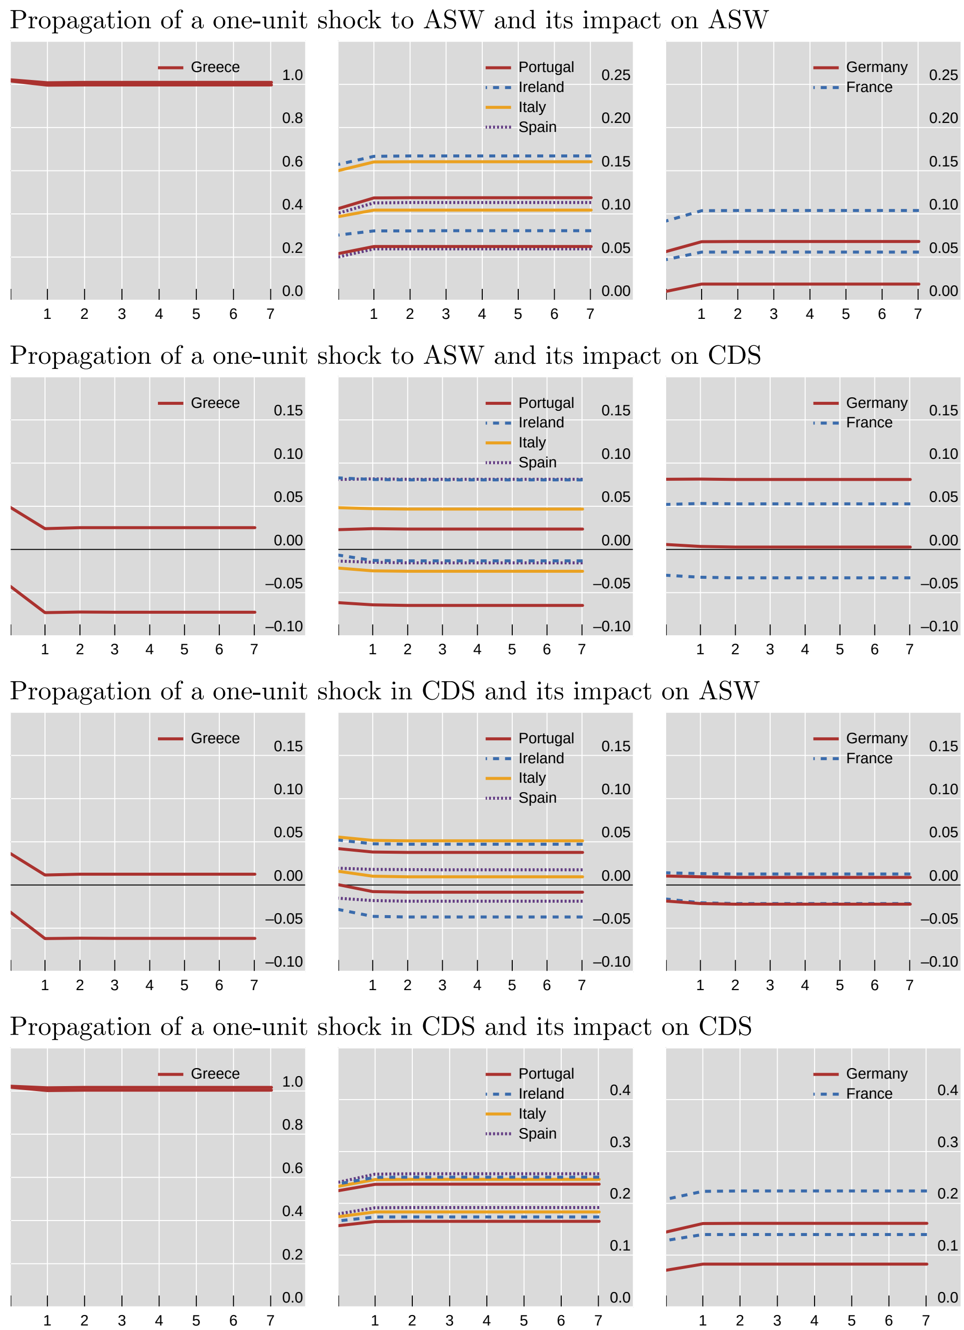

Figure D.2: Accumulated impulse responses in crisis period - shock in Greece

This figure illustrates the accumulated impulse response for CDS and ASW to a one-unit shock (increase) for the period from 20 October 2009 to end-December 2011. For further details see Figure 4.

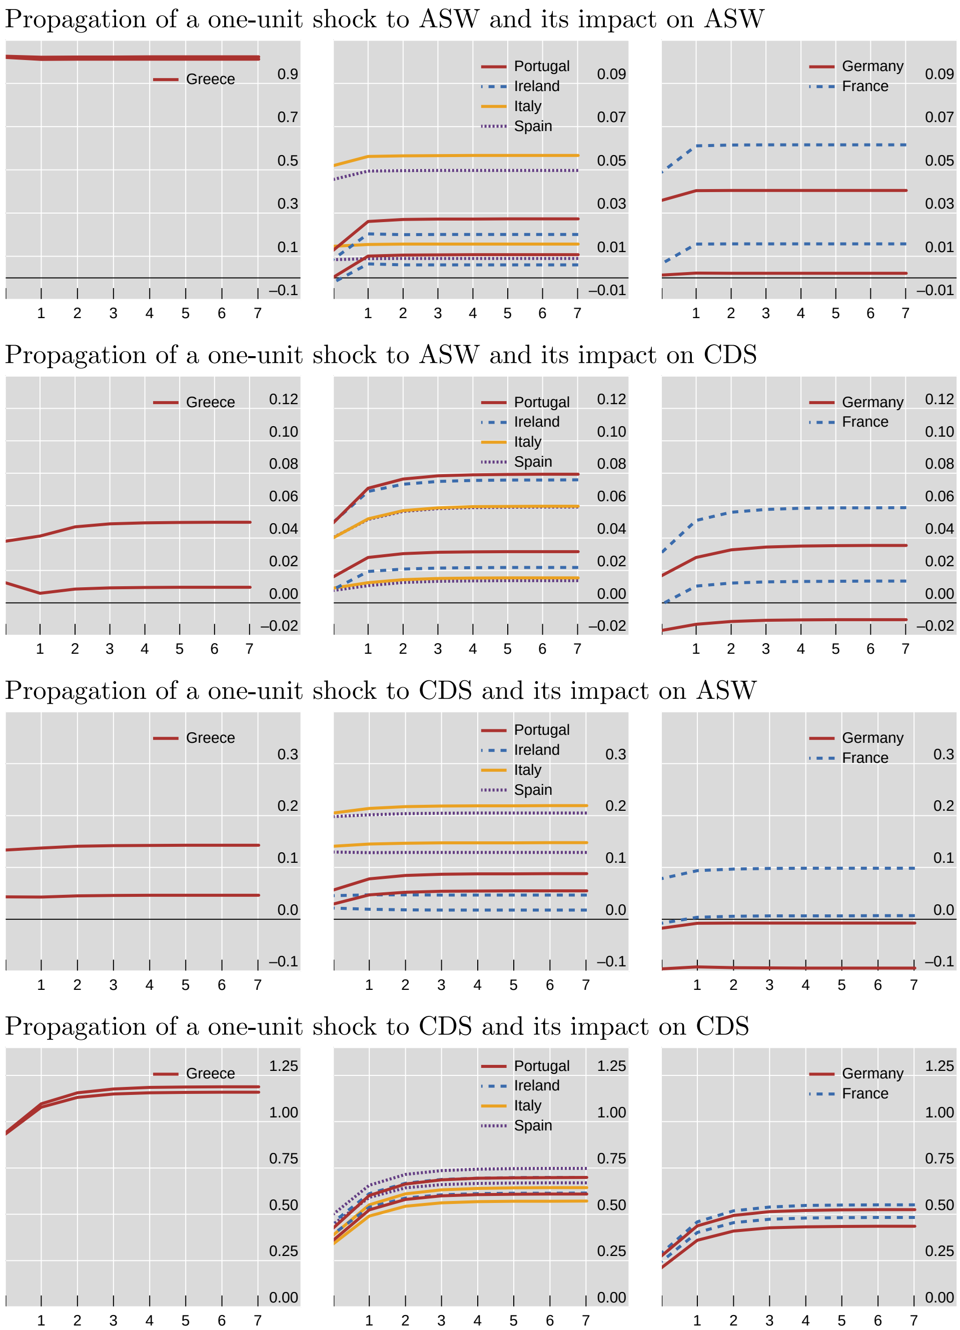

Part Ib: The importance of the bond and the CDS markets in the transmission of euro area sovereign credit risk contagion[1]

1 Introduction

The 2008-09 financial crisis caused investors to look more critically at the fiscal outlook in a number of countries, including many in the euro area. This resulted in a sharp rise in sovereign credit spreads for a number of euro area countries. At their peak, yield spreads on sovereign bonds relative to German bonds reached several hundred basis points, while before the global financial crisis these spreads had averaged only a few basis points. At the same time, market interest in trading credit risk protection on euro sovereign borrowers via credit default swaps (CDS) grew substantially and spreads on such instruments also surged. Due to these developments, policy makers and regulators paid increased attention to the market for sovereign CDS. Of particular interest was, first, the interplay between the pricing of sovereign risk in CDS and bond markets, second, the possibility that one market could be systematically leading the other, and third, the potential for sovereign credit risk contagion effects. In this paper we focus on the last-mentioned. An important motivation for providing financial support to Greece, despite the no-bailout clause in the Maastricht Treaty, was the fear on the part of policy makers that a Greek default would spill contagiously over to other highly indebted countries in the euro area (Constancio; 2011).

We analyse euro area sovereign credit risk contagion effects in GIIPS[2] countries plus France and Germany from January 2008 to end-December 2011, which we split into a pre-sovereign debt crisis and sovereign debt crisis period. Little is yet known about the transmission channels of sovereign credit risk contagion through the CDS and the bond market, and their relative importance during the euro area sovereign debt crisis. We define contagion as the strengthening in the transmission of a financial shock from one country onto another country as for example in Bekaert et al. (2005) and Ehrmann and Fratzscher (2017). We allow for the transmission to evolve through the CDS and/or the bond market. As we have data for both the CDS market and the bond market, we are able to assess contagion effects and dynamics in sovereign credit risk arising from of both transmission channels (CDS and bond market). We do not aim to find specific drivers of sovereign contagion as e.g. cross-border financial or trade linkages as in Brutti and Saure (2015) or Chinn and Forbes (2004) nor specific global shocks as drivers of contagion as for example in Adrian and Brunnermeier (2011) or Longstaff et al. (2011). Our paper intends

[1] This chapter is joint work with Kristyna Ters (University of Basel) and has been submitted to the *Journal of Central Banking* as Ters, K. and Urban, J. (2017) The importance of the bond and the CDS markets in the transmission of euro area sovereign credit risk contagion.

[2] Greece, Ireland, Italy, Portugal, Spain

to detect the financial market instrument (CDS and/or bond) that has been responsible for the transmission of contagion between sovereign entities and the interaction between the CDS and the bond market in the pricing of sovereign credit risk contagion. This is similar to the approach by Grammatikos and Vermeulen (2012) who analyse the transmission of financial and sovereign debt crises between the US and the euro area by including stocks and CDS of financial and non-financial corporate CDS.

According to the theoretical no-arbitrage condition that was introduced by Duffie (1999) the CDS and the bond yield spread of a given reference entity should both price the same probability of default. Thus, the price of the CDS and the bond yield over the risk-free interest rate should be equal for a given reference entity if markets are perfect and frictionless. However, this approximate relationship was clearly not applicable anymore during the euro area sovereign debt crisis period as the pricing and the dynamics in the CDS and the bond markets for euro area sovereign diverged (see for example Fontana and Scheicher (2016) and Gyntelberg et al. (2013)). Therefore, analysing the importance of the CDS and the bond market with respect to their role in the transmission process of sovereign contagion is an important contribution of this paper to the existing recent literature. Indeed, the ban on outright short selling of CDS in the euro area which came into force in November 2012, was introduced following overreactions in CDS prices due to speculators (Sarkozy et al.; 2010). However, not only the CDS markets but also the bond market have been blamed to be the driver of contagion dynamics during the euro area sovereign debt crisis. Constancio (2011) claims in his speech on 10 October 2011 that overreactions of market participants triggered sell-offs in euro area government bond markets. Specifically, he argues that the sell-off in Spanish and Italian government bonds following the downgrade by Moody's on 5 July 2011 was an overreaction as there were no adverse data releases or changes in the budgetary situation for Spain and Italy around that time. Constancio (2011) states that these bond market movements can be clearly specified as contagion dynamics. Still, quantifying evidence is missing on which market has been the primary source of the transmission of sovereign risk contagion. Therefore, understanding the dynamics and the interplay between the CDS and the bond market and their relative importance in the transmission of sovereign risk is not only crucial from a regulatory but also from a monetary policy viewpoint.

The use of intraday data allows us to capture the intraday patterns of credit risk contagion. Indeed, shocks that may seem to affect several countries simultaneously when viewed at a daily or lower data frequency are revealed, through the lens of intraday data, to have possible origins in one particular country with clear contagion effects on other countries. Also, Gyntelberg et al. (2013) point out that the higher accuracy of the results as compared with lower-frequency data is a clear advantage of using intraday data. They find that when using daily data, due to the dramatically smaller number of observations, the confidence bands of the estimated coefficients are extremely wide and therefore in most cases not significant. The use of intraday data to study the transmission of shocks to financial markets during the euro area sovereign debt crisis has in addition been advocated by Ehrmann and Fratzscher (2017), Ghysels et al. (2017) and Rogers et al. (2014).

Existing research has differentiated between cross-country and intra-country analysis. By using a panel VAR methodology we can control for both country-specific risk and contagion effects across countries. Panel VARs are built on the same logic as standard VARs, but, by adding a cross-sectional dimension, they become a much more powerful tool for addressing policy questions of interest related, for example, to the transmission of shocks across borders as stated in Fomby et al. (2013) and Canova and Ciccarelli (2013). By using the method of Canova and Ciccarelli (2013), we are able to shock the credit risk of an individual country and derive the individual response for each country in the panel. Linear VAR models assume by construction that the individual country shocks are uncorrelated across different countries which is particularly unrealistic in a cross-border contagion detection application in macroeconomics as also stated by Stock and Watson (2016).

One of the key contributions of our paper relative to the existing literature on sovereign risk contagion during the recent euro area sovereign debt crisis is that we do not use simple yield differences as our measure of cash spreads. Rather, we use carefully constructed asset swap spreads (ASW) based on estimated zero-coupon government bond prices. This ensures that we are comparing like with like in our empirical analysis for sovereign credit risk, by exactly matching the maturities and the cash flow structures of the CDS and the cash components. The use of ASW is also in line with the practice used in commercial banks when trading the CDS-bond basis. In addition, by calculating ASW we are able to estimate contagion impacts on Germany such as flight-to-safety effects. Germany is not included in most contagion studies since German Bund yields are used as the risk-free interest rate in the "constant maturity" approach mentioned above. Moreover, our analysis relies on intraday price data for both CDS and bonds, allowing us to estimate the contagion dynamics and the transmission channel of contagion (CDS or bond market) substantially more accurately than existing studies.[3]

Further, by extending our model to include the economic surprise index, we are able to estimate how much of sovereign risk contagion can be attributed to macroeconomic news or overreaction/lack of belief by market participants. In addition, we study the effects of the economic adjustment programmes (EAP) of the Troika, a decision group formed by the European Commission (EC), the European Central Bank (ECB), and the International

[3] Gyntelberg et al. (2013) discuss the advantages of using intraday data due to the higher accuracy of the results as compared with lower-frequency data.

Monetary Fund (IMF), in order to analyse the effect of bailouts on sovereign credit risk contagion.

Our findings will improve the understanding of the dynamics in the market for sovereign credit risk. Further, the segregation of the credit risk transmission channels will enable policy makers and regulators to better assess the relative importance of, and the risks arising from, the derivative and cash market.

The remainder of the paper is structured as follows: Section 2 discusses the related literature and Section 3 our data and the relationship between CDS and bonds. Section 4 explains the set-up and estimation of the panel VAR (PVAR) model and its extensions. Section 5 presents the empirical results and Section 6 concludes.

2 Review of the related literature and contagion identification strategy

A large body of literature concerns itself with the potential reasons and transmission channels for contagion as well as with theoretical modelling of contagion. A whole strand of this literature focuses on empirical tests for the existence of contagion in a given stress period, that is, it asks if there are stronger cross-market linkages in times of crisis. This paper belongs to the latter type, as we focus on testing for the existence of contagion during the euro area sovereign debt crisis. We extent existing research by analysing the relative importance of the CDS and the bond market as transmission channels for sovereign risk contagion.

We limit our discussion to papers that measure contagion among sovereign credit markets. Adrian and Brunnermeier (2011), Ang and Longstaff (2011), Brutti and Saure (2015), Eichengreen and Mody (2000), Kamin and von Kleist (1999), Longstaff et al. (2011), Mauro et al. (2002), and Pan and Singleton (2005) concentrate on the relationship between sovereign credit spreads and common global and financial market factors. These papers empirically identify factors that are significant variables for CDS spreads, such as the U.S. stock and bond market returns as well as the embedded volatility risk premium. Our approach in identifying the transmission of sovereign risk, defined as the propagation of a financial shock in one country onto another country, while we allow for the transmission to evolve through the CDS or the bond market (or both) is closest to the paper by Grammatikos and Vermeulen (2012). They analyse the transmission of financial and sovereign shocks between the US and the euro area by including stocks and CDS of financial and non-financial corporates.

The issues of financial shock contagion among countries in the euro area during the recent sovereign debt crisis have figured prominently in recent empirical research. Caporin et al. (2012) analyse risk contagion using the CDS spreads of the major euro area countries using different econometric approaches such as Bayesian modelling. They find that the diffusion of shocks in euro area CDS has been remarkably constant, while the risk spillover among countries is not affected by the size of the shock. Other examples are Bai et al. (2015), Neri and Ropele (2013), De Santis (2012), and Arghyrou and Kontonikas (2012). They all employ time series modelling approaches for contagion and include sovereign bond spreads (yield to maturity) to reflect pure credit risk considerations and macroeconomic variables. The results are mostly discriminated in terms of core (such as France and Austria) and peripheral countries (GIIPS). In general, they find that the bond spreads of lower-rated countries increase along with their Greek counterparts. However, their results in terms of magnitude, responses to shocks and contagion effects on core countries are somewhat mixed. Similarly to these studies, Koop and Korobilis (2016) employ an enhanced panel approach for empirical modelling of financial contagion across countries based on Canova and Ciccarelli (2013). Their findings give no clear support for euro area sovereign credit risk contagion between core and peripheral countries. Further, they do not find evidence for interdependencies within the peripheral countries (e.g. from Greece to Spain). The different results reported in these studies could be due to sample differences or to how bond spreads are calculated. Most empirical research uses the "constant maturity" approach to calculate bond yield differences (relative to Germany). Further, daily or weekly data are used for the empirical analysis, which may lead to inaccurate shock and contagion estimations, especially in periods when activity in sovereign risk markets is high during times of stress.

As pointed out by Corsetti et al. (2011), there is much disagreement among economists about the exact definition of contagion and how it should be tested. For Constancio (2011) and Forbes (2012) contagion occurs when financial or macroeconomic imbalances (shocks) create a systemic risk beyond that explained by economic fundamentals. Contagion differs from macroeconomic interdependence (comovement) among countries in that the transmission of risk to other countries is different under normal economic conditions. Forbes (2012) defines contagion as spillovers resulting from extreme negative effects. If comovements of markets are similarly high during non-crisis periods and crisis periods, then there is only evidence of strong economic linkages between these economies (Missio and Watzka; 2011). Kaminsky et al. (2003) describe contagion as an episode in which there are significant immediate effects in a number of countries following an event, such as when the consequences are fast and furious and evolve over a matter of hours or days. When the effect is gradual, Kaminsky et al. (2003) refer to it as comovement rather than contagion. Our contagion identification strategy builds upon the criteria described by the ECB where they clearly specify that contagion is present when (i) the transmission is in excess of what can be explained by economic fundamentals which is also in line with the theory provided by Eichengreen and Mody (2000); (ii) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); (iii) the events constituting contagion are negative extremes; and (iv) the transmission is sequential.

3 Data

The core data we use in our empirical analysis consists of USD-denominated five-year maturity intraday quotes on CDS contracts and government bonds for France, Germany, Greece, Ireland, Italy, Portugal and Spain. We choose this group of countries as it includes the countries most affected by the euro sovereign debt crisis, as well as Germany, which serves as the near-risk-free reference country, and France, which we consider as a low-risk control country.

According to Gyntelberg et al. (2013) when one considers the number of quotes of CDS contracts at the peak of the sovereign debt crisis in 2010, the five-year segment is the most liquid. The use of intraday data in our empirical analysis enables us to obtain much sharper estimates and clearer results with respect to market mechanisms as also shown in Gyntelberg et al. (2013). Further, they point out that sovereign credit risk dynamics follow an intraday pattern which is also advocated by Ehrmann and Fratzscher (2017), Ghysels et al. (2017) and Rogers et al. (2014).

Our sovereign bond price data come from MTS (Mercato Telematico dei Titoli di Stato). The MTS data comprise both actual transaction prices and binding bid-offer quotes. The number of transactions of sovereign bonds on the MTS platform is, however, insufficient to allow us to undertake any meaningful intraday analysis. Therefore, we use the trading book from the respective domestic MTS markets.[4] The MTS market is open from 8:15 to 17:30 local Milan time, preceded by a pre-market phase (7.30 to 8.00) and an offer-market phase (8:00 to 8:15). We use data from 8:30 to 17:30.[5]

The number of trades (see Figure 1) remain stable throughout our whole sample period (2008-2011) while in some cases they even increase (Greece is the only exception in 2011). Also, the number of quotes submitted to the MTS platform remain very stable as displayed in Figure 2. Additionally, we have far above 80% of our constructed 30 minute intervals filled with quotes (see right-hand side of Figure 2). We can therefore rule out any liquidity shocks in the cash bond market.

4 We ignore quotes from the centralised European platform (market code: EBM), as quotes for government bonds on the centralised platform are duplicates of quotes on the domestic platforms.

5 Due to our equidistant sampling frequency we have to discard either the first or the last 15 minutes of each trading day. We analysed the data quality in both intervals and found, on balance, that the first 15 minutes have a lower data quality than the interval at the end of the trading day. Consequently, we discarded the first 15 minutes of each trading day.

Figure 1: EuroMTS number of trades

The graphs show the number (in thousands) of trades per year. Italy is shown separately because the number of trades are more than an order of magnitude higher than for the other countries.

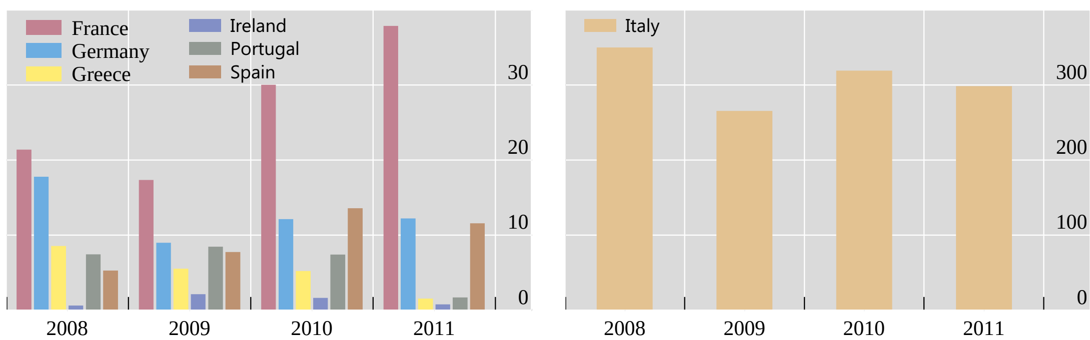

Figure 2: EuroMTS bond price data from the trading book – 5-year maturity

The left-hand figure shows the number (in millions) of data ticks in the trading book. This includes all bonds with a maturity between 4 and 6 years in the 5-year segment. The right-hand figure shows the percentage of 30 min. intervals during the trading period, which contain at least one data tick in the trading book (left-hand scale) and the the number (in thousands) of non-empty half hour intervals per year (right-hand scale). We consider 18 half hour slots per trading day, from 8:30 to 17:30 CET/CEST.

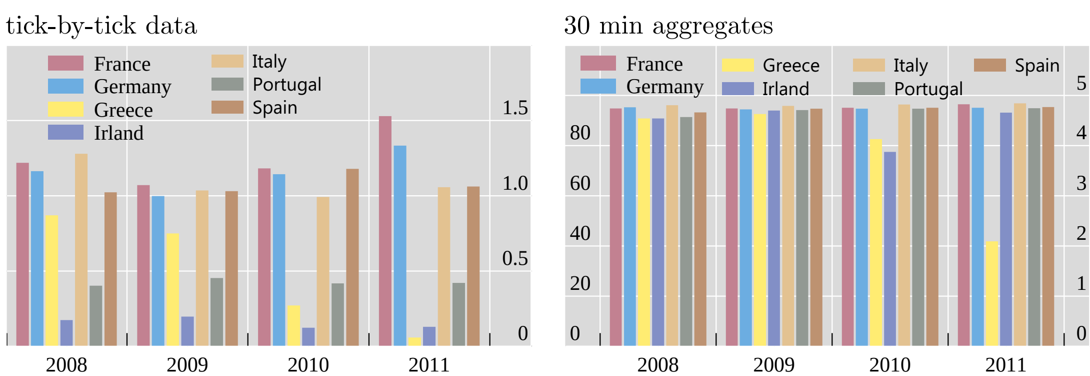

The CDS data consist of price quotes provided by CMA (Credit Market Analysis Ltd.) Datavision. CMA continuously gathers information on executable and indicative CDS prices directly from the largest and most active credit investors. After cleaning and checking the individual quotes, CMA applies a time- and liquidity-weighted aggregation, so that each reported bid and offer price is based on the most recent and liquid quotes. The CDS market, which is an OTC market, is open 24 hours a day. However, most of the activity in the CMA database is concentrated between around 7:00 and 17:00 London time. As we want to match the CDS data with the bond market data, we restrict our attention to the period from 8:30 to 17:30 CET (CEST during summer).

We construct our intraday data on a 30-minute sampling frequency on our data set, which spans from January 2008 to end-December 2011. The available number of indicative quotes for CDS does not allow a data frequency higher than 30 minutes. The euro area sovereign CDS markets were very thin prior to 2008, which makes any type of intraday analysis before 2008 impossible. Microstructural noise effects may come into play when high frequency data is used (Fulop and Lescourret (2007)). However, this does not apply to our data based on a 30-minute sampling frequency because we average the reported quotes over each 30 minute interval (for tests, robustness checks and for a more detailed discussion please refer to Gyntelberg et al. (2013)).

The overall liquidity in the CDS market increased throughout our whole sample period (Greece is the only exception in 2011), measured by the number of data-ticks provided by CMA (see Figure 3) as well as measured by the number of trades per week and net notional amounts outstanding as reported by the ISDA (see Figure 4).

Figure 3: CDS data from CMA Datavision – Liquidity

The left-hand panel shows the number (in thousands) of data ticks per year. The right-hand panel shows the number (in thousands) of non-empty half-hour intervals per year (right scale). We consider 18 half-hour slots per trading day, from 8:30 to 17:30 CET/CEST. The left scale in the right-hand panel shows the percentage of 30 min. intervals which contain at least one data tick during the 18 daily half-hour intervals we consider. Source: CMA Datavision

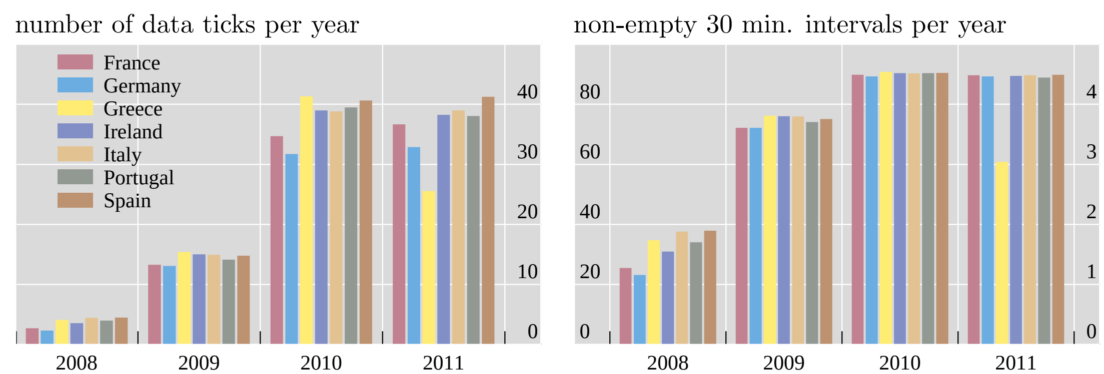

The left-hand side of Figure 3 clearly shows that the frequency of our intraday data is limited by the years 2008 and 2009 of the sovereign CDS data. While we could have chosen a sampling frequency below 30 minutes in the years 2010 and 2011, the available data does not allow any feasible analysis with a sampling frequency below 30 minutes during the pre-sovereign debt crisis period.

Figure 4: CDS net notional amount outstanding and trade count.

Panel A shows the number (in thousands) of trades per week. Panel B shows the net notional amount outstanding in USD billion. The start of the period is given by the data availability. The missing data for Greece from 2012 until 2014 is due to the Greek restructuring. Source: ISDA

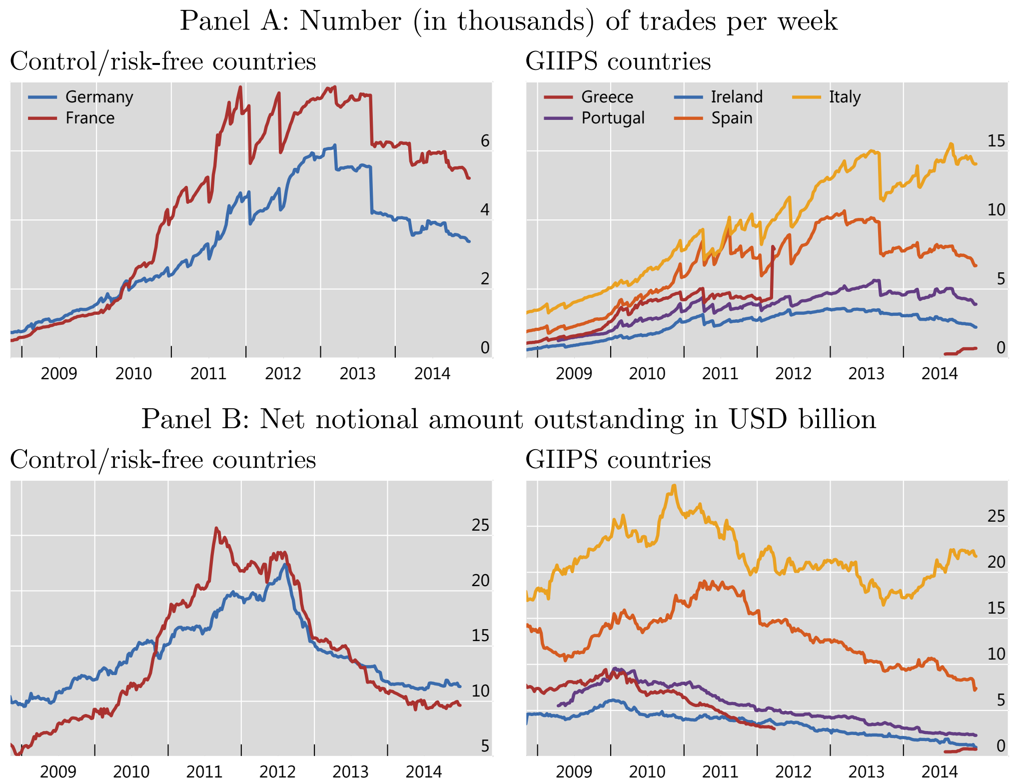

The use of a 30 minute sampling frequency results in a time series length of almost 19,000 for the period 2008 until 2011. This enables us to not only study intraday patterns of credit risk transmission but also to achieve a much higher statistical significance as compared to other studies (see Ehrmann and Fratzscher (2017), Ghysels et al. (2017), Gyntelberg et al. (2013) and Rogers et al. (2014).

We split the data into two subsamples. The first covers the period from January 2008 to 19 October 2009 and, as such, represents the period prior to the euro area sovereign debt crisis. While this period includes the most severe phase of the financial crisis, including the default of Lehman Brothers, it is relatively unaffected by market distortions stemming from concerns about the sustainability of public finances in view of rising government deficits and therefore represents the pre-sovereign debt crisis period. The second subsample covers the euro area sovereign debt crisis period and runs from 20 October 2009 to end-December 2011. As the beginning of the crisis period, we designate 20 October 2009, when the new Greek government announced that official statistics on Greek debt had previously been fabricated. Instead of a public deficit estimated at 6% of GDP for 2009, the government now expected a figure at least twice as high.

We employ CDS and bond data in our analysis in order to be able to differentiate between the transmission of sovereign risk contagion according to the credit risk channel from one country to another. Duffie (1999) argues that, since the CDS and the bond yield spread both price the same default of a given reference entity, their price should be equal if markets are perfect and frictionless. Thus, in a perfect market, due to arbitrage, the CDS spread equals the bond yield over the risk-free rate. However, for this parity to hold, a number of specific conditions must be met, including that markets are perfect and frictionless, that bonds can be shorted without restrictions or cost and that there are no tax effects, etc. According to Duffie (1999) par floating rate notes (FRN) should be used for the bond yield computation, because FRN, unlike plain vanilla bonds, carry only credit risk and no interest rate risk. However, they are relatively uncommon, in particular for sovereign entities. A further complication linked to the use of fixed-rate or plain vanilla bonds as substitutes is that it is unlikely that the maturity of these instruments exactly matches that of standard CDS contracts.

To ensure proper comparability with CDS, Gyntelberg et al. (2013) employ synthetic par asset swap spreads (ASW) for the bond leg of the basis. The use of ASW is in line with the practice used in commercial banks when trading the CDS-bond basis. By calculating ASW for our empirical analysis, we ensure accurate cash flow matching, as opposed to studies that use simple "constant maturity" yield differences for credit risk.

For an in-depth discussion on the construction of our intraday ASW please refer to Appendix A, Gyntelberg et al. (2013) and for a general discussion to O'Kane (2000).

Based on common and individual panel unit root tests our CDS and ASW price data (displayed in Figure 5) is I(1). Therefore, we estimate our subsequent models in first differences.

Figure 5: CDS and ASW spreads in basis points

The figures are based on data with a 30-minute sampling frequency. Our split into the pre- and the crisis period is indicated by the vertical line in each figure. Due to the Greek debt restructuring the data for Greece ends in September 2011.

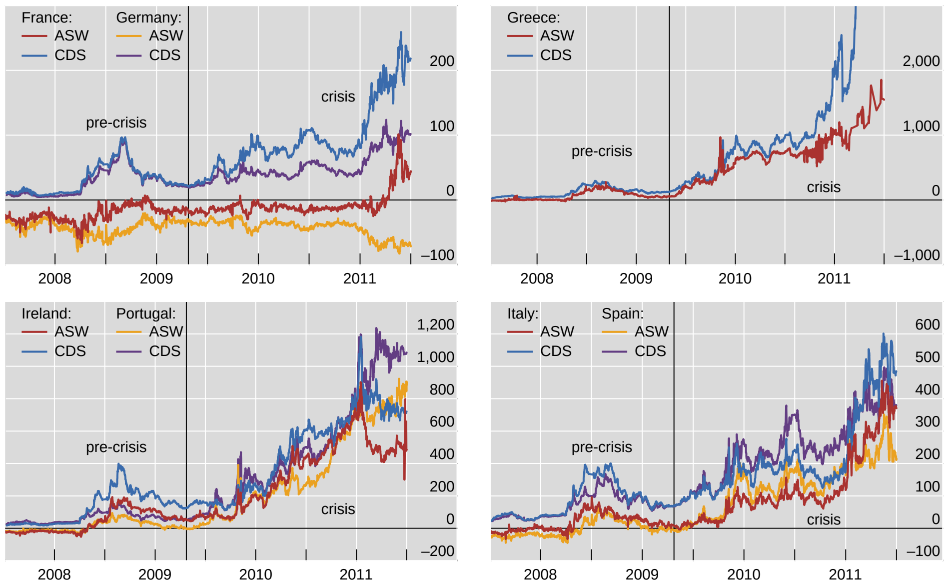

4 Modelling sovereign credit spread contagion

To empirically measure the impact of euro area sovereign credit risk contagion effects according to the credit risk channel (CDS and bond market), we employ a panel vector autoregressive (PVAR) model. PVARs have the same structure as VAR models, in the sense that all variables are assumed to be endogenous but with the difference that a cross-sectional dimension is added to the representation. We define our PVAR model following Binder et al. (2005) with fixed effects when N is finite and T is large, as $i = 1, ..., N$ is the cross-sectional dimension and $t = 1, ..., T$ is the time-series dimension in our model. According to Koop and Korobilis (2016) and Canova and Ciccarelli (2013), in this setup the PVAR is the ideal tool for examining the transmission of macroeconomic or financial shocks from one country to another.

The PVAR has several advantages over individual country VARs in a time series framework. By analysing a panel of countries, we can more accurately model contagion from one country to another since the panel approach captures country-level heterogeneity. We control for cross-sectional heterogeneity by including fixed effects in the regression. By using CDS and ASW as endogenous variables for each country in our cross-section,[6] we can differentiate the credit risk channel of contagion, which improves the understanding of the sovereign risk contagion dynamics. In contrast to linear VAR models we therefore do not have to assume that the individual country shocks are uncorrelated across different countries which is particularly unrealistic in a cross-border contagion detection application in macroeconomics as also stated by Stock and Watson (2016).

With an extension of the PVAR using a purely exogenous variable, we can assess the effect of unexpected macroeconomic news[7] on credit risk contagion for the countries in our sample.

4.1 Panel VAR

In vector autoregressive models (VAR) all variables are treated as endogenous and interdependent in both a dynamic and static sense. The VAR model is formally defined as:

$$Y_t = A_0 + A(\mathrm{L})Y_{t-1} + u_t, \tag{1}$$

where Y_t is a G x 1 vector of endogenous variables and $A(\mathrm{L})$ is a polynomial in the lag operator, A_0 is a G x 1 vector and u_t is a G x 1 vector of i.i.d. shocks.

Y_t is the stacked version of y_{it}, the vector of G variables for the cross-sectional dimension across countries $i = 1, ..., N$, i.e. $Y_t = (y'_{1t}, y'_{2t}, ..., y'_{Nt})'$ and $t = 1, ..., T$. The major difference between a VAR and the PVAR is that the covariance σ_{ij} of the residuals is zero by definition for country i different from country j in a VAR model. The PVAR is defined as follows:

$$y_{it} = A_{0i} + A_i(\mathrm{L})Y_{t-1} + u_{it}. \tag{2}$$

A_{0i} are G x 1 vectors and A_i are G x GN matrices. By construction our model allows for country specific slopes and intercepts while the intercept captures country-specific heterogeneity. Further, lags of all endogenous variables of all entities enter the equation of country i. Canova and Ciccarelli (2013) call this feature “dynamic interdependencies”. The residual u_{it} is a G x 1 vector and $u_t = (u_{1t}, u_{2t}, ..., u_{Nt})$. u_{it} is generally correlated across the cross-sectional dimension i. Canova and Ciccarelli (2013) call this feature “static interdependencies”. Thus the variance-covariance matrix for a PVAR has the following property $E(u_{it}u'_{jt}) = \sigma_{ij} \neq 0$ for $i \neq j$, i.e. static interdependencies occur when the correlations between the errors in two countries' VARs are non-zero. On the other

[6] bi-variate estimation per country

[7] By using the economic surprise index as a predetermined purely exogenous variable in the PVARX model.

hand, dynamic interdependencies occur when one country's lagged variables affect another country's variables. Hence, the PVAR is more flexible compared to a VAR ($\sigma_{ij} = 0$ for $i \neq j$).[8]

In our bivariate case, i.e. $G = 2$, we can rewrite the PVAR in Equation (2) as:

$$\begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{it} = \begin{pmatrix} A_{01} \\ A_{02} \end{pmatrix}_{i} + \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}_{ji} (\mathrm{L}) \begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{j\,t-1} + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}_{it}, \tag{3}$$

where A_{mn} are scalars and $i, j = 1, ..., N$ is the number of countries in the cross-sectional dimension. We employ a bivariate case with CDS and bond data for each country i in our analysis in order to include the dynamics of both markets for sovereign credit risk according to Duffie (1999) and Gyntelberg et al. (2013) and furthermore to be able to differentiate between the transmission of sovereign risk contagion according to the credit risk channel from one country to another.

For the estimation, we follow the approach proposed by Canova and Ciccarelli (2009) of an unrestricted PVAR which allows for the selection of restrictions involving dynamic interdependencies, static interdependencies and cross-section heterogeneities[9]. According to an empirical model comparison by Koop and Korobilis (2016), the proposed methodology by Canova and Ciccarelli (2009) shows the best properties compared to other PVAR approaches. Canova and Ciccarelli (2009) suggest adopting a flexible structure through a factorisation of the coefficients in Equation (2). Through the flexible coefficient factorisation, the PVAR can be rewritten as a reparametrised multicountry VAR and estimated using SUR (Canova and Ciccarelli; 2009). The advantage of this flexible factorisation is that the overparametrisation of the original PVAR is dramatically reduced while, in the resulting SUR model, estimation and specification searches are constrained only by the dimensionality of the estimated coefficient matrix (for a more in-depth discussion please refer to Canova and Ciccarelli (2009) and Koop and Korobilis (2016)).

We have tested different lag lengths in our PVAR model and find highly robust results. The magnitude of shock transmission is unaffected by the lag selection, while the time until the shock is absorbed can slightly vary by 1 interval (30 mins). Increasing the lag length reduces the accuracy of the results, due to the higher number of estimated parameters. However, up to a lag length of five intervals, the results remain highly robust as we are using intraday data with a minimum of 4000 observations within each subsample (pre-crisis and crisis).

Our contagion identification strategy builds upon the criteria described by the ECB. Constancio (2011) clearly specifies that the ECB identifies contagion according to: (i) the

[8] According to Canova and Ciccarelli (2013), these features distinguish a panel VAR typically used for macroeconomics and finance from a panel VAR used in microeconomics.

[9] We use demeaned and standardised first differences.

transmission is in excess of what can be explained by economic fundamentals which is also in line with the theory provided by Eichengreen and Mody (2000); (ii) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); (iii) the events constituting contagion are negative extremes; and (iv) the transmission is sequential.

Constancio (2011) states that there is no general agreement about a clear contagion identification strategy and about which one of these before mentioned four criteria are necessary or sufficient to characterise a contagion event. Therefore, we follow the same identification strategy for contagion as the ECB. As there is no general agreement concerning the above mentioned four criteria, we restrict the identification of contagion to the case where all four conditions must be fulfilled.

4.2 Panel VARX

As an extension to the previous analysis, and in order to test whether the transmission of a financial shock is in excess of what can be explained by economic fundamentals (criterion (i) of the ECB's identification strategy of contagion), we consider the response of credit risk in CDS and bond markets in our GIIPS and low-risk country sample to unexpected macroeconomic news. We follow Canova and Ciccarelli (2013) by extending the PVAR model in Equation (2) to include a predetermined purely exogenous variable X_t which results in the following PVARX model:

$$y_{it} = A_{0i} + A_i(\mathrm{L})Y_{t-1} + F_i(\mathrm{L})X_t + u_{it}, \tag{4}$$

with X_t as an M x 1 vector (M is equal to the number of exogenous variables) common to all entities i. In our case, the PVARX can also be rewritten as:

$$\begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{it} = \begin{pmatrix} A_{01} \\ A_{02} \end{pmatrix}_i + \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}_{ji} (\mathrm{L}) \begin{pmatrix} \Delta CDS \\ \Delta ASW \end{pmatrix}_{j\,t-1} + \begin{pmatrix} \tilde{F}_1 \\ \tilde{F}_2 \end{pmatrix}_i (\mathrm{L})X_t + \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}_{it}. \tag{5}$$

The extension to the PVARX model allows us to analyse whether credit risk responses can be attributed to macroeconomic fundamental news or if exaggerations in terms of lack of belief of economic agents also contributed to credit risk responses as previously defined in our identification strategy under criteria (i).

We employ the economic surprise index as predetermined exogenous variable X_t for unexpected macroeconomic news in the euro zone, i.e. $M = 1$. The Citigroup economic surprise index for the euro zone[10] (see Figure 6) is widely recognised in academia and by

[10] Bloomberg ticker: CESIEUR Index

practitioners for measuring unexpected macroeconomic news (such as in Goldberg and Grisse (2013), Scotti (2013), and Paulsen (2014)).

The Citigroup economic surprise index measures how economic news/data is developing relative to the anticipated consensus forecasts of market economists. According to Citigroup, the index captures objective and quantitative measures of economic news, defined as weighted historical standard deviations of data surprises (actual releases versus the Bloomberg survey median). A positive reading of the economic surprise index for the euro zone suggests that economic releases have on balance beaten the consensus, while a negative reading indicates the opposite. The index captures economic news on macroeconomic and fiscal variables such as employment change, the housing market, retail sales, debt-to-GDP, the budget deficit and consumer confidence in the euro zone. Thus, the Citigroup economic surprise index does not include news on monetary policy decisions.

The economic surprise index has a daily frequency that differs from the intraday data that we are analysing in this paper. However, as market participants are exposed to economic news throughout the whole day, we disperse the actual Citigroup economic surprise index data given at the end of the trading day over that entire day. We conducted several simulations with different distributions across a whole trading day to generate a pseudo-intraday economic surprise index. Our results remain extremely robust to this experiment. The design of the experiment is justified because we are interested not in the exact time line of the absorption of unexpected macroeconomic news, but rather in a qualitative picture of whether market participants react to fundamental macroeconomic news in the pricing of sovereign credit risk.

Figure 6: Citigroup Economic Surprise Index

The figure shows the daily values of the Citigroup Economic Surprise Index (Bloomberg ticker: CESIEUR Index) from the beginning of 2008 until the end of 2011. The vertical line corresponds to 19 October 2009, the end of our pre-crisis period. Source: Bloomberg.

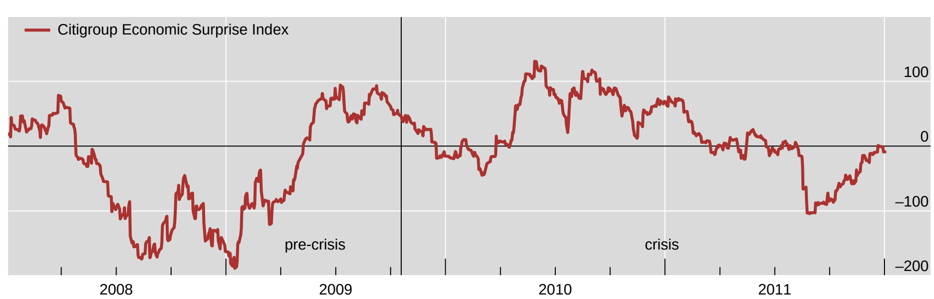

4.3 Impulse responses and shock identification

We carry out an impulse response analysis to investigate cross-border shock transmission in sovereign credit risk in the euro area countries that were most affected during the sovereign debt crisis. Further, we analyse the differences in the magnitude, the speed and the dynamics between the pre-sovereign debt crisis and crisis sample in order to identify contagion as specified by Constancio (2011) and described in Section 4.1. We focus on individual country shocks propagating from GIIPS countries and analyse the impact of an unexpected one-unit shock to credit risk in both the CDS and ASW markets from country i to j.

In standard VAR models (see Equation (1)), shock identification is performed by imposing a Choleski decomposition in all countries. To reduce the number of identification restrictions in a VAR model, it is assumed that $E(u_{it}u'_{jt})$ is block diagonal, with blocks corresponding to each country. Canova and Ciccarelli (2009) state that block diagonality implies that within a country, variables are allowed to move instantaneously, but across entities, variables can only react with one lag. Assuming block diagonality is however a strong restriction and furthermore unrealistic in a cross-border contagion application using financial market data as also stated by Stock and Watson (1988).

The identification of shocks for PVAR models as defined in Equation (2) is more complicated, given that the PVAR model allows for static interdependencies, as u_{it} is correlated across entities i. Thus, cross-entity symmetry in shock identification cannot be assumed. We compute the impulse responses following Canova and Ciccarelli (2009) as the difference between two conditional forecasts: one where a particular variable is shocked and one where the disturbance is set to zero.

Canova and Ciccarelli (2013) defines the shock identification formally as: y^t to be the history of our time series y_t up to time stamp t , Θ^t are the estimated coefficients up to t. Let $W = (\Sigma_u, \sigma^2)$ as the sum of squared residuals and the residual variance; set $\xi'_t = [u'_{1t}, u'_{2t}]$ where u'_{1t} are shocks to the endogenous variables and u'_{2t} shocks to the predetermined or exogenous variables. δ is a one standard deviation shock as we have standardised our time series.

Let $model^1_t$ contain y^t, Θ^t, W up to t and a one standard deviation shock $\xi^{\delta}_{j,t}$ to country j. The model forecast without shock is defined as $model^2_t$ which contains y^t, Θ^t, W and their variance up to t. The impulse responses at the future horizon τ to a δ impulse is the difference between these two conditional forecasts defined as:

$$IR^j_y(t,\tau) = E(y_{t+\tau}|model^1_t) - E(y_{t+\tau}|model^2_t). \tag{6}$$

5 Results

As already described in Section 4.1 our contagion identification, as specified by Constancio (2011), has to fulfil all of the following four criteria: (i) the transmission is in excess of what can be explained by economic fundamentals which is also in line with the theory provided by Eichengreen and Mody (2000); (ii) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); (iii) the events constituting contagion are negative extremes; and (iv) the transmission is sequential.

Overall our findings suggest that the CDS market was the main transmission channel for sovereign credit risk contagion during the euro area sovereign debt crisis. We can identify contagion through the CDS market during the crisis period in all GIIPS countries according to the above mentioned criteria as proposed by the ECB (Constancio; 2011). All GIIPS countries exhibit much higher impulse responses following a one-unit shock in ΔCDS from any GIIPS country. The impulse responses to a one-unit shock in ΔCDS are much higher in magnitude during the crisis period compared to the pre-sovereign debt crisis period which gives support to criterion (ii). Additionally, the impulse responses are positive in magnitude which is an increase in sovereign credit risk (higher probability of default) and thus, a negative effect (criterion (iii)). Our results also show, that during the crisis period a one-unit shock in ΔCDS from any GIIPS country resulted in a negative extreme effect in all other GIIPS countries (criterion (iv)).[11] When controlling for unexpected economic fundamental news in the PVARX framework (see Chapter 4.2) we find that during the crisis period positive unexpected news led to a higher default probability.[12] This finding gives support to the fact that market reactions can not be explained by economic fundamentals during the crisis period which supports criterion (i).

Our results suggest that sovereign risk pricing was not driven by fundamental macroeconomic news during the crisis period. Therefore, we perform an event-study to test whether market participants were reactive to the introduction of the Economic Adjustment Programmes (EAP), i.e. we test if these programmes modified the dynamics in the pricing of sovereign credit risk GIIPS countries. The rational behind this test is to analyse if the EAP changed market participants belief in the path for public finances in the countries under bailout.

Based on our findings in Section 3 on sovereign credit risk liquidity as displayed in Figures 1 to 4 we can rule out that liquidity shocks have been primarily responsible for our results during the sovereign debt crisis period (including the event studies). As the bond market liquidity remained stable during the entire sample period, the strong decline

[11] Due to our static interdependencies in our PVAR model setup we have contemporaneous shock transmissions which reflect the strong interlinkages of financial markets.

[12] For the pre-sovereign debt crisis period we find that positive unexpected news lead to lower probabilities of default.

of the importance in the transmission of sovereign risk cannot be explained by liquidity shocks in the bond market for sovereign risk. The CDS market liquidity measured by trade count and net notional volume has even strongly increased for the control countries and only slightly increased for the GIIPS countries. Hence, CDS market liquidity can only be partially made responsible for the dominance of the CDS market during the sovereign debt crisis.

5.1 Results for GIIPS and low-risk countries

As a general result, we find that, pre-crisis, the bond and CDS markets are of similar importance, i.e. the response function of country i to a one-unit shock to the ASW and CDS markets of country j is of a comparable size in the two markets (see Figure 7). These results are as expected, as both markets should price the countries' credit risk equally (Duffie; 1999). During the crisis period, the CDS market became more relevant on balance (see Figure 8). Interestingly, the inter-market shock transmission, i.e. from CDS to ASW and vice versa, is not important during the pre-crisis period. This weak connection between the two markets during the pre-sovereign debt crisis period could be explained by different market participants and their distinct investment horizons. Insurance firms active in the bond market have a longer investment horizon than, for example, hedge funds in the CDS markets. During the crisis period, shock transmission between markets becomes relatively more important, suggesting a stronger inter-market connectivity. Market participants get more vigilant to potential bad news, which may spill over from other markets.

Further, we find that the decay of a shock is faster on average in the pre-crisis period than in the crisis period (see Figures 7 and 8). The timelines of our estimated shock contagion and absorption are dramatically shorter than in existing empirical studies, such as Koop and Korobilis (2016), who find that shock contagion spreads on average within one to two months in the case of shocks that do not decay over a timeline of 10 months. We find for both sample periods that contagion propagates immediately within the first 30-minute time interval. Therefore, responses to shock contagion are typically not lagged as found, for example, in Koop and Korobilis (2016). Further, the average response for shock absorption is around one hour in the pre-crisis period and slightly longer at one to two hours on average during the crisis period. This result is clearly in line with the generally accepted notion that financial markets react very fast to new information (Gyntelberg et al.; 2013). The slower speed of shock absorption during the crisis seems to contradict our statement above that market participants are more reactive to news during crisis periods. This can be explained by the fact that the estimated timeline of shock absorption during the crisis period is strongly affected by turmoil in financial markets, while the pre-crisis period represents a relative normal market environment for European sovereign states without fast and furious shock contagion but rather with comovements across markets as specified by the ECB's contagion identification strategy and also defined by Forbes (2012). Furthermore, the ECB and Forbes (2012) also state that contagion differs from macroeconomic interdependence (comovement) among countries in that transmission of risk to other countries is different under normal economic conditions. Hence, we find comovement effects rather than contagion during the pre-crisis period as we find much smaller impulse responses to a one-unit shock compared to the crisis period.

In the pre-crisis period, a credit risk shock spreading from the ASW to the CDS market and vice versa had more or less the same impact in terms of magnitude and shock absorption. Thus, the derivatives market and the spot market were about equally significant in terms of shock contagion prior to the euro area sovereign debt crisis. However, during the crisis period we find that shock transmission from the ASW to the CDS market had a dramatically lower impact than vice versa. This leads to the assumption that the importance of the spot market as a channel of financial shock contagion decreased during the euro area sovereign debt crisis. Thus, the contagion of shocks to credit risk has been transmitted predominantly through the derivatives market.

During the pre-crisis period, a one-unit shock to either the ASW or CDS of country i results in a spread widening for all countries. However, during the crisis, we find evidence of a flight-to-safety effect to German bonds, as Germany is considered a safe haven for investors. This effect is visible in the inter-market connection, i.e. a positive shock to a GIIPS country's CDS or ASW leads to spread tightening in German ASW, while we cannot report a similar effect for German CDS. Similar behaviour is not visible for France, despite it being considered a low-risk control country.

During the pre-crisis period, we find that the magnitude of the impulse responses is similar across all countries, while during the crisis period, GIIPS countries exhibit much larger impulse responses than the rest of our sample countries do.

In contrast to the other empirical studies using this methodology, Koop and Korobilis (2016) find confidence bands for their impulse responses that all lie between positive and negative reactions to a one-unit shock to Greek bond yields relative to Germany. The advantage of our approach, using ASW and intraday data, dramatically increases the precision of the results during the crisis period.

In addition to the impulse response functions for a shock to Greek ΔASW and ΔCDS in Figures 7 and 8, we present impulse response functions for a shock to Spanish and Portuguese ΔASW and ΔCDS in Appendix B.[13]

[13] Impulse response functions for a shock to Irish and Italian ΔASW and ΔCDS show similar results and can be provided on request.

Figure 7: Impulse responses in the pre-crisis period - shock in Greece

This figure illustrates the impulse response for ΔCDS and ΔASW to a one-unit shock (increase) for the period from January 2008 to 19 October 2009. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS or ΔASW. The number of 30-minute time intervals is described by the x-axis. For each impulse response we plot the upper and lower 95% confidence bands.

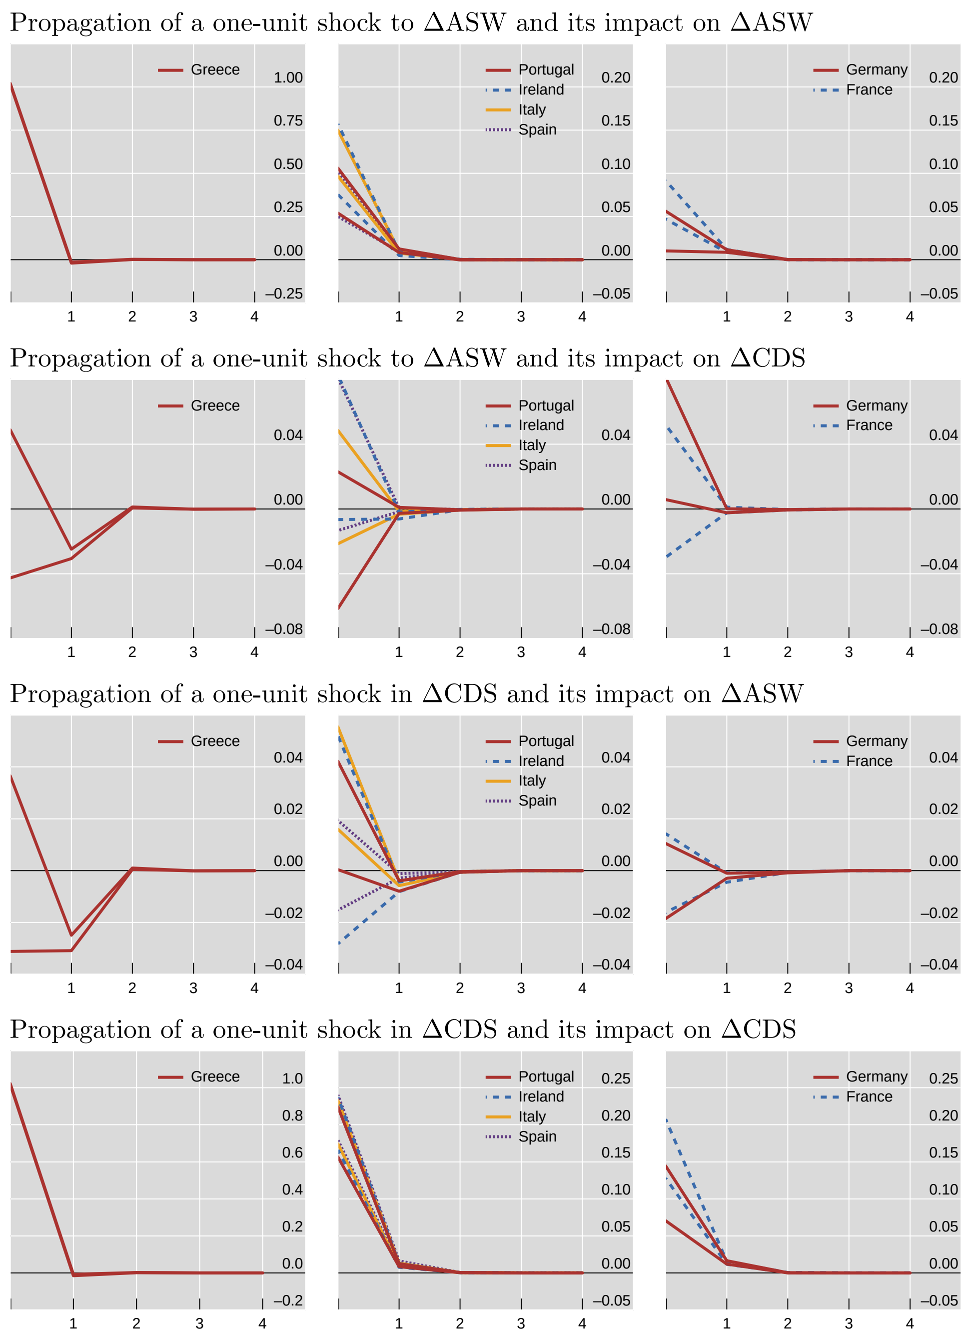

Figure 8: Impulse responses in the crisis period - shock in Greece

This figure illustrates the impulse response for ΔCDS and ΔASW to a one-unit shock (increase) for the period from 20 October 2009 to end-December 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS or ΔASW. The number of 30-minute time intervals is described by the x-axis. For each impulse response we plot the upper and lower 95% confidence bands.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

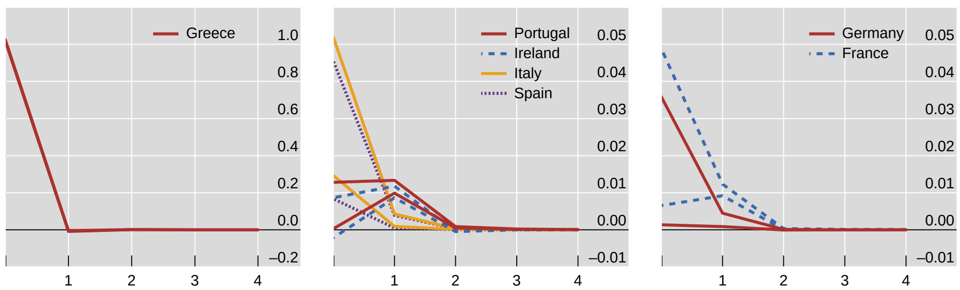

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

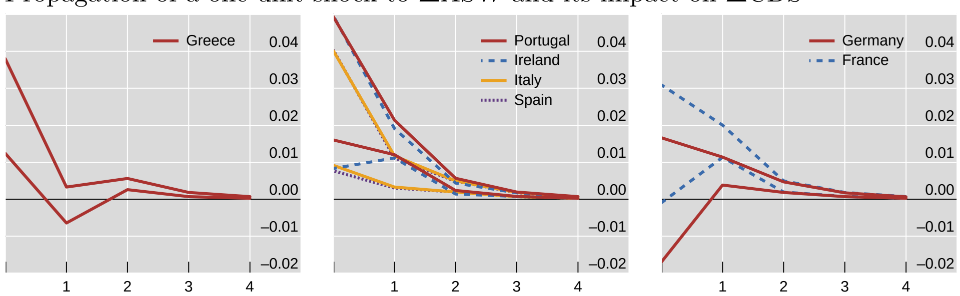

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

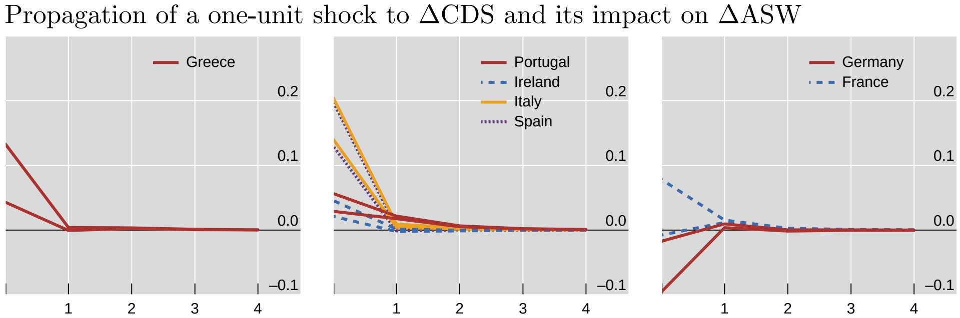

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

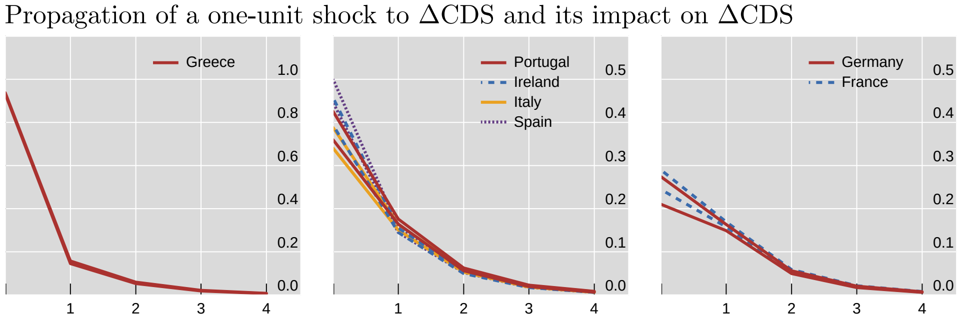

Even though policy makers may not be interested in short-lived intraday movements in sovereign credit risk, our results show that the level impacts from the short-term dynamics are persistent (see Appendix C). Hence, our results are important with regard to financial stability.

5.2 The impact of unexpected macroeconomic news on sovereign credit risk: Results from a PVARX experiment

In this section, we conduct an experiment with the aim of analysing whether responses to shocks and shock contagion can be attributed to macroeconomic fundamentals or if overreactions in sovereign credit risk during the crisis period might also be due to self-fulfilling prophecies. For Forbes (2012), contagion occurs when financial or macroeconomic imbalances (shocks) create a systemic risk beyond that explained by macroeconomic fundamentals. Gibson et al. (2012) explain the effect of self-fulfilling prophecies by interest rate spreads that were lower than justified by fundamentals prior to the crisis, owing to the role played by Greece's euro area membership in biasing investor expectations. During the crisis period, Gibson et al. (2012) define this self-fulfilling prophecy effect that interest rate spreads were higher than those predicted by fundamentals in terms of the market's disbelief that sustainable financial consolidation measures and structural reforms would be implemented. Also the ECB's contagion identification strategy is in line with the before mentioned theory as they identify contagion when the transmission is in excess of what can be explained by economic fundamentals (criterion (i) according to Constancio (2011)).

Our experiment is designed in a similar way to that of Canova and Ciccarelli (2009). We distribute the daily data of the economic surprise index over each trading day (18 time intervals). The distribution is chosen such that the maximum is reached at noon, and the sum of the 18 different intraday values is equal to the value reported by the Citigroup economic surprise index. We experimented with different distributions across a trading day and, despite the arbitrary distribution assumption, we found robust results. Next, the last seven values are removed from all time series in order to be close to the last maximum (in the case of a positive reading of the surprise index) or close to the last minimum (in the case of a negative reading of the surprise index). We then fit the PVARX model from Equation (4) and produce an out-of-sample forecast for eight intervals beyond the last data point[14], which is in the case of the pre-crisis period 15 October 2009 and in the case

[14] We have chosen the forecast length of eight intervals in order to be slightly longer than the number of removed values (seven). We experimented with different forecast lengths and found that the qualitative results remained robust. Again the choice to remove the last seven values is motivated to be close to the last maximum/minimum of the surprise index, reached at midday by construction.

of the crisis period 26 May 2011[15]. We call this forecast the "real forecast". Further, we repeat this same procedure, but now set the data of the surprise index of the last day to zero, i.e. we artificially remove the last positive or negative "shock" given by the data. We again produce an out-of-sample forecast which we call the "counterfactual forecast". The difference between the real and the counterfactual forecast captures the impact of the positive or negative values of the Citigroup economic surprise index on the last day. In other words, the experiment mimics what would have happened if the last positive or negative economic news had not occurred and thus helps answer the question of whether macroeconomic fundamental news can explain changes in sovereign credit risk.

During the pre-crisis period, we find for all countries in the sample that a positive (negative) shock from the economic surprise index on the last day (15 October 2009) leads to an expected decrease (increase) in credit risk (see Figure 9). Prior to the crisis, the magnitude of the effect following an unexpected macroeconomic news shock is similar in the bond and CDS markets. Our pre-crisis results indicate that markets reacted to macroeconomic news in pricing sovereign credit risk.

During the euro area sovereign debt crisis period, a negative reading of the economic surprise index on the last available day (26 May 2011) leads surprisingly to a decrease in credit spreads in most countries (see Figure 10). In rational markets, a negative economic news shock should lead to an increase in sovereign credit risk and thus to an increase in spreads. Our results are counterintuitive, unlike those for the pre-crisis period. For the crisis period, they show that credit markets were driven not by macroeconomic news, but most likely by monetary policy, political decisions and speculations. Figure 11 displays the individual components (the real and the counterfactual forecasts) of our unexpected economic news shock experiment for the crisis period. Subtracting the counterfactual forecast in row 2 from the real forecast in row 1 of Figure 11 produces the forecast in row 1 of Figure 10. The same applies to the remaining rows in Figures 10 and 11. Surprisingly, in most cases a negative economic shock leads to a tightening of credit spreads (rows 1 and 3 of Figure 11).

The shapes of the curves in Figures 9, 10 and 11 are due to our particular choice of decomposing the daily Citigroup economic surprise index into intraday intervals. However, other distribution choices leading to different shapes of our curves do not change the results. This gives support to the self-fulfilling crisis theory, that changes in sovereign credit risk during the euro area sovereign debt crisis were only partially driven by macroeconomic fundamentals, as markets did not react to economic news in contrast to the pre-crisis period.

[15] We chose end-May as the last time stamp in our experiment for the crisis period, because the liquidity in the Greek bond market deteriorates. The lack of pricing data from May onwards does not allow to generate a sensible intraday forecast for our experiment.

Figure 9: Positive shock to the Economic Surprise Index during the pre-crisis period

This figure illustrates a scenario of a real positive shock to the economic surprise index minus a counterfactual scenario where we assumed that the shock did not happen. The period under consideration is from January 2008 until 15 October 2009. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

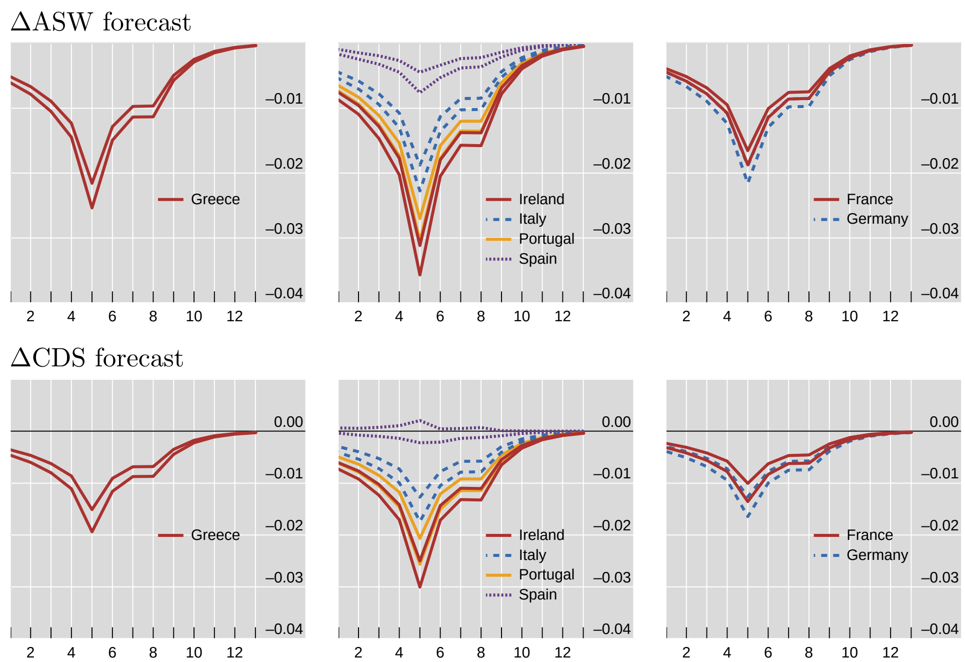

Figure 10: Negative shock to the Economic Surprise Index during the crisis period

This figure illustrates a scenario of a real negative shock to the economic surprise index minus a counterfactual scenario where we assumed that the shock did not happen. The period under consideration is from 20 October 2009 until end-May 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

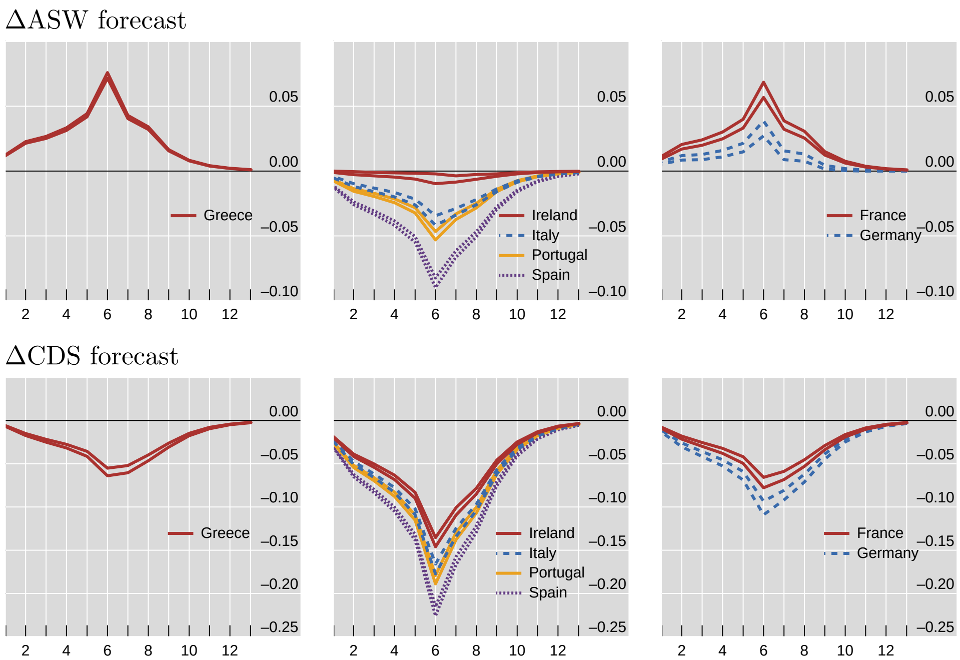

Figure 11: Real and counterfactual forecast decomposition during the crisis period

This figure presents the individual components, the real and the counterfactual forecasts, of our experiment for the crisis period 20 October 2009 until end-May 2011. The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the response of ΔASW (upper part) and ΔCDS (lower part) in basis points. The number of 30-minute time intervals is described by the x-axis. We plot the upper and lower 95% confidence bands for each country.

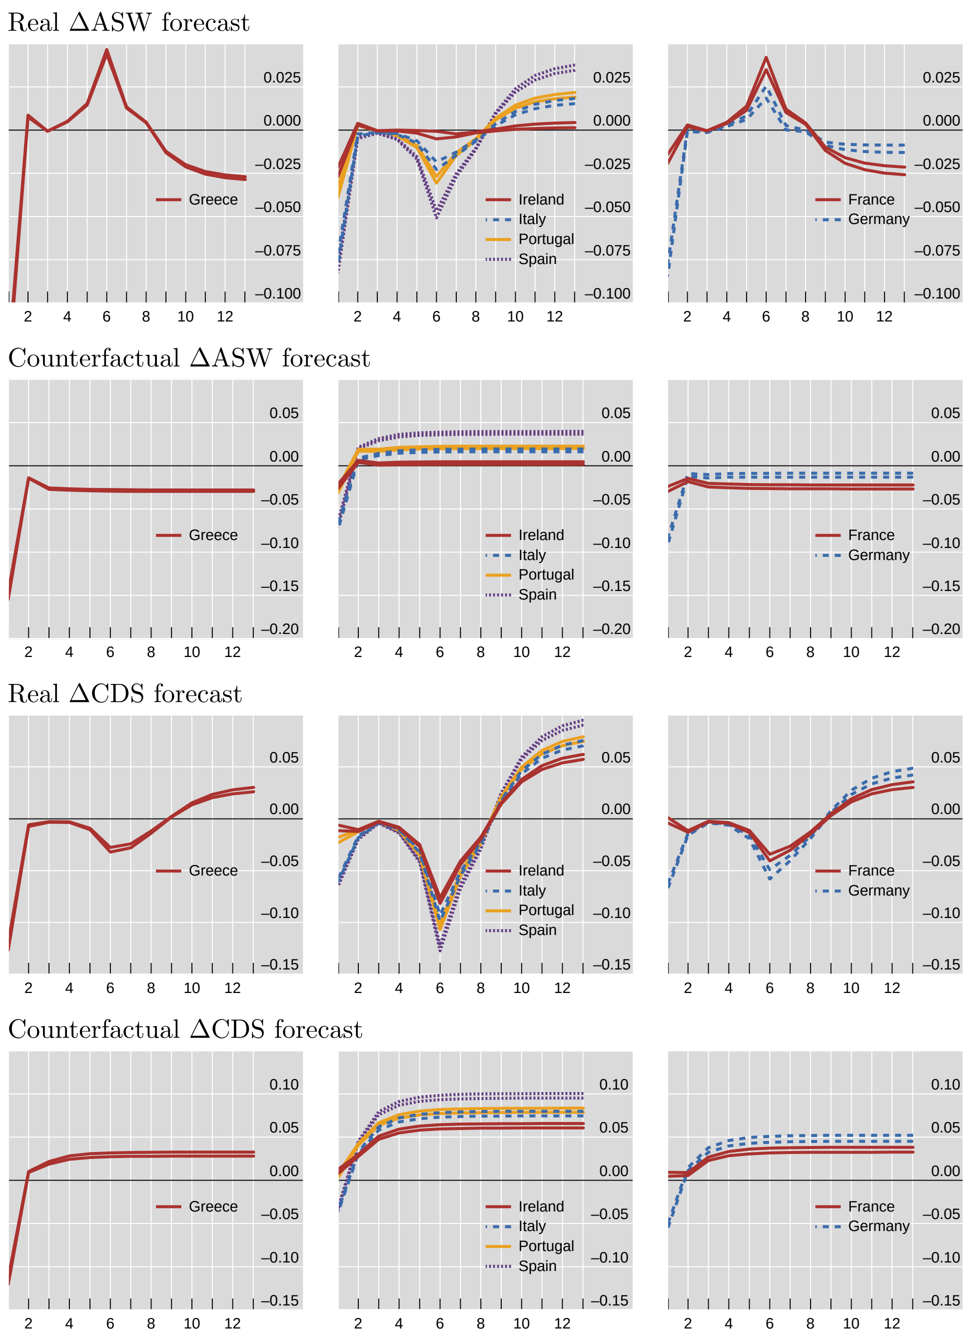

5.3 The effect of the economic adjustment programmes on contagion dynamics

The previous Section 5.2 finds that markets for sovereign credit risk were not driven by fundamental macroeconomic news during the crisis period. Thus, market participants did rather react to either monetary policy (conventional or unconventional) or were driven by their lack of belief in a sustainable path for public finances which cannot be explained by macroeconomic news.

We test for the lack-of-belief by assuming that the introduction of the economic adjustment programmes (EAP) must have changed the financial markets' belief in the future path of public finances. The EAP provided financial help to avoid a default in the most distressed countries and imposed economic and fiscal conditions. During our crisis period we observe four EAP events: two bailouts for Greece, one for Portugal and one for Ireland (see Figure 11).

Figure 12: Economic adjustment programmes

The figure shows the 30-minute CDS spreads for Greece, Ireland and Portugal. The vertical lines correspond to the four bailout events during our sample period. We define the bailout events as the announcement dates of the Memorandum of Understanding for the economic adjustment programmes in Ireland, Portugal and the first bailout in Greece. The event date of the second bailout in Greece is defined as the announcement of the preliminary draft of the second EAP.

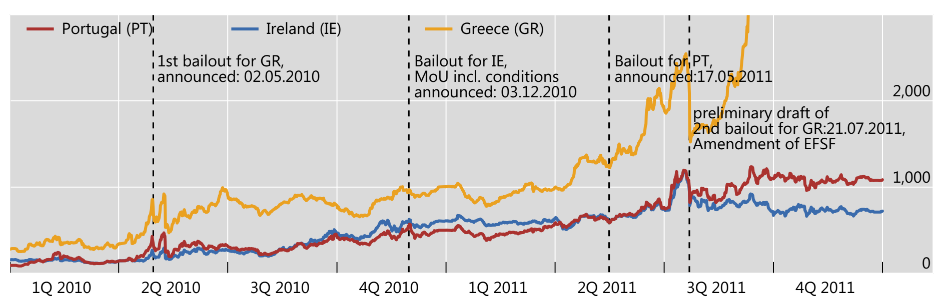

The estimation of the contagion dynamics before and after the bailouts (EAPs) in Greece, Ireland, and Portugal, is performed in a PVAR as defined in Equation (2). We define the event dates as the announcements of the Memorandum of Understanding (MoU) for the bailouts (EAPs) in Ireland, Portugal and the first bailout in Greece by the Troika, a decision group formed by the EC, the ECB, and the IMF. The MoU for the first EAP for Greece was announced on 2 May 2010. The Troika agreed on providing EUR 110bn to Greece in equity support to banks. The MoU for the EAP in Ireland was announced on 3 December 2010 with an amount of EUR 67.5bn.[16] Portugal's MoU for the EAP was announced on 17 May 2011 with an amount of 78bn EUR. Finally, on 21 July 2011 the preliminary draft of the second bailout package for Greece was approved with an amount of EUR 100bn. This preliminary draft tried to address the limitations of the first Greek EAP. The second EAP in Greece further comprised the private sector involvement (PSI) in the form of a voluntary haircut of EUR 37bn in Greek sovereign bonds held by financial institutions, the prolongation of the debt repayment with lower interest rates, and the establishment of a Task Force for Greece in order to promote economic growth. More importantly, it was also decided to enlarge the European Financial Stability Facility (EFSF) from EUR 440bn to EUR 780bn. A few days after the announcement of the preliminary draft on the second bailout in Greece, rating agencies started to downgrade Greece to junk level as they concluded that the proposed restructuring of the government debt would amount to a selective default.

All of the bailouts, except for the first EAP for Greece, were executed under the EFSF that had been created by the euro area Member States in June 2010 as a temporary crisis resolution mechanism.[17] The financial assistance to Ireland, Portugal and Greece under the EFSF was financed through the issuance of bonds and other debt instruments on capital markets.

Our event window prior to and after the announcement of a MoU is set to five trading days, with 18 intraday observations for each trading day. As our prior analysis in Section 5.1 showed strong evidence that the bond market was not the main venue for sovereign credit risk contagion during the crisis period, we estimate the PVAR in Equation (2) in an univariate framework (G=1) with CDS data.

Figure 13 displays the contagion dynamics prior to (left panel) and after (right panel) the bailout events, where all impulse response functions are significant at 95% CL. Contagion dynamics are estimated as the responses to a one-unit shock in ΔCDS arising in the country under the EAP. Before a bailout, we find strong contagion in magnitude and in terms of interlinkages of CDS markets across countries. As soon as the MoU containing the bailout conditions was published, we find that contagion was strongly reduced in terms of magnitude. The interlinkages across financial markets in the GIIPS countries decrease when the country which receives the bailout is also the origin of the shock. Furthermore, the timeline of the shock absorption is significantly shorter after a bailout, decreasing from 3 hours to 2 hours on average. When the origin of the one-unit shock comes from the country under bailout (EAP), our results suggest that the EAPs were able to reduce overall sovereign risk and contagion interlinkages across all GIIPS countries.

[16] An additional amount of EUR 17.5bn was financed via the Irish Treasury and the National Pension Reserve Fund.

[17] Since 1 July 2013, the EFSF may no longer engage in new financing programmes or enter into new loan facility agreements.

After the first bailout in Greece we find that sovereign risk contagion did dramatically decrease in magnitude in all countries, except for France and Germany. These findings indicate that financial markets' perception of the default probability for France and Germany increased after the introduction of the first EAP. This can be explained by the fact that France and Germany were the largest contributors of financial aid in the first bailout package. The shares for participation in the EAP were determined according to the GDP per capita in each country. Germany provided EUR 28bn and France EUR 16.8bn to the first EAP for Greece out of a total of EUR 110bn. Interestingly, this effect is only visible for the first bailout. In the subsequent three bailouts that were transacted under the EFSF we do not find the effect of increased credit risk in France and Germany.

A further interesting finding is that the first three bailouts did not have a significant effect on financial markets' pricing of sovereign risk in GIIPS countries that were not participating in an individual EAP as shown in Figure 11. More specifically, when we shock a country i following a bailout event in country j we do not find a decrease in magnitude of shock contagion after the event. This finding implies that overall, the first three EAP programmes were able to reduce sovereign risk in the country which received the bailout, but overall it did not affect the joint sovereign credit risk contagion in GIIPS countries.

We find a strong impact reducing the financial market interlinkages after the announcement of the preliminary draft of the second bailout for Greece in July 2011. When we shock country i, we find a significant decrease in the magnitude of shock contagion after the event. This finding is different from the before mentioned first three bailouts where we only find a decrease of shock contagion for the country that receives financial assistance. The fourth bailout shows a reduction of shock contagion for all countries (and not only when we shock the country under the bailout). Figure 14 displays the responses to a one-unit shock in Italy following the announcement of the preliminary draft in Greece. We find this effect for all countries i in our sample, that are the origin of the one-unit shock. Thus we can conclude that the fourth bailout in our sample had a stabilising impact on the joint credit risk contagion and interlinkages across all GIIPS countries. As our fourth event is also subject to a significant enlargement of the EFSF to address previous shortfalls, this stabilising effect could be attributed to the amendment of the EFSF. Further evidence of this comes from the fact that after 21 July 2011, due to the restructuring of Greek debt and a further sharp increase in the budget deficit, the CDS price of Greece sky-rocketed. However, the stabilising effect of the amended EFSF on the other countries in our sample outweighed the probability of the dramatically increased Greek default probability. This is also visible in Figure 11 as sovereign credit risk for Ireland and Portugal does not further increase after 21 July 2011.

Figure 13: Bailout events and contagion dynamics

This figure presents the pre- and the post-event windows of contagion dynamics estimated with the PVAR in Equation (2) and 30-minute CDS data. The event windows are set to five trading days, with 18 intraday observations for each trading day. The bailout events are defined in Figure 11. The y-axis represents the response of ΔCDS to a shock arising in the country of the EAP event. The number of 30-minute time intervals is described by the x-axis. We display the point estimates of the impulse responses as all estimates are significant at the 95% CI.

Economic adjustment programme for Greece (1st) 2.5.2010, shock in Greece

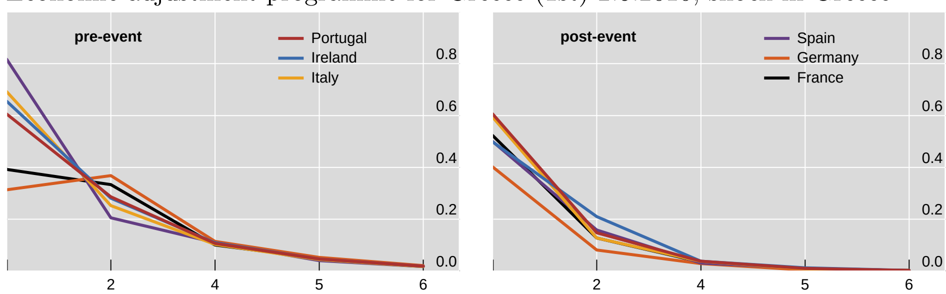

Economic adjustment programme for Ireland 3.12.2010, shock in Ireland

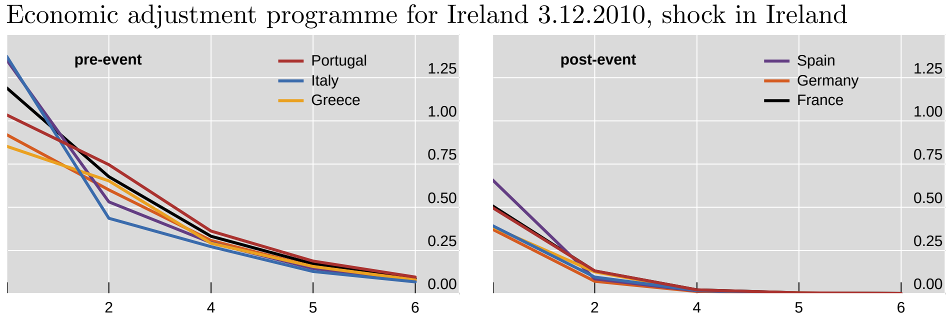

Economic adjustment programme for Portugal 17.5.2011, shock in Portugal

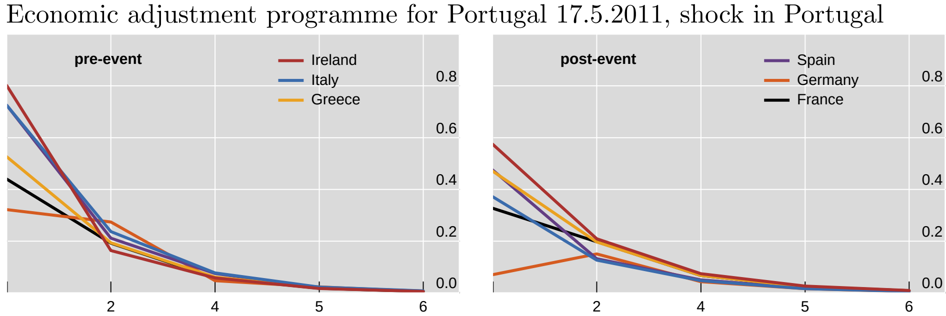

Economic adjustment programme for Greece (2nd) 21.7.2011, shock in Greece

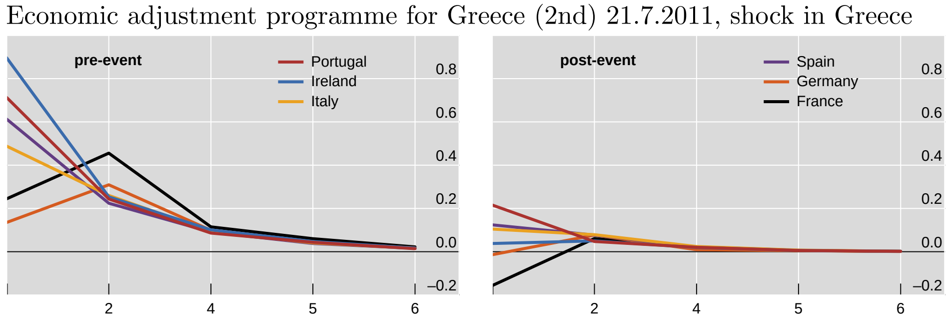

In Figure 14 we display the point estimates of the impulse responses as all estimates are significant at the 95% CL.

Figure 14: Second economic adjustment programme for Greece including the enlargement of the EFSF

This figure presents the pre- and the post-event windows of contagion dynamics estimated with the PVAR in Equation (2) and 30-minute CDS data. The event windows are set to five trading days, with 18 intraday observations for each trading day. The event date is the 21 July 2011 which is the announcement of the preliminary draft of the second bailout for Greece including the enlargement of the EFSF. The one-unit shock comes from Italy which is not the country under the EAP. The y-axis represents the response of ΔCDS. The number of 30-minute time intervals is described by the x-axis. We display the point estimates of the impulse responses as all estimates are significant at the 95% CI. Greece is not part of the sample as the CDS has been close to default following the restructuring and therefore we do not have satisfying data quality after this specific event.

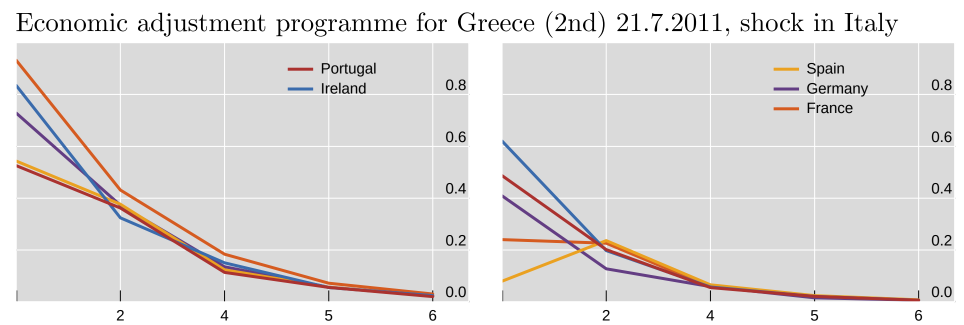

We do not find any significant effect in impulse responses when we use the governments' request for a bailout as the event date. This can be explained by the fact that bailout conditions are not yet defined and made publicly available.

6 Conclusion

The CDS market was the main venue for the transmission of sovereign credit risk contagion during the euro area sovereign debt crisis. In contrast, we find that, prior to the crisis, the two markets (CDS and bond) were similarly important in the transmission of financial contagion. The importance of the bond market decreased relative to the CDS market during the crisis period. We find evidence for sovereign credit risk contagion during the euro area sovereign debt crisis period, as our results show more drastic reactions to shocks in terms of magnitude and absorption compared to the pre-crisis period. Thus, our results on the responses to sovereign credit risk shocks during the crisis period confirm the contagion across euro area countries, as they result from extreme negative, systemic effects and are much larger in magnitude compared to the pre-crisis period, a fact which cannot be explained by macroeconomic fundamentals.[18] We find comovement effects rather than contagion during the pre-crisis period, as markets were driven by economic fundamentals, while the responses to sovereign credit risk shocks remain moderate in magnitude. The use of intraday data substantially increases the precision of the results and we find average timelines of financial shock contagion of one to two hours during the crisis period and 30 minutes to one hour prior to the crisis. This is a clear indication of the efficiency of financial markets. Liquidity in sovereign credit risk markets remained unchanged or even increased during the sovereign debt crisis period. Thus, liquidity shocks were not responsible for contagion effects during the sovereign debt crisis and comovement during the pre-crisis period.

Our results using an unexpected exogenous macroeconomic news shock suggest that, during the pre-crisis period, markets for sovereign credit risk were driven by macroeconomic news. Positive news led to a decrease in credit spreads and negative news to an increase. For the euro area sovereign debt crisis period, our results show that movements in sovereign credit spreads did not respond to macroeconomic news but were rather driven by either monetary policy or exaggerations in financial markets due to a lack of belief (a self-fulfilling crisis). These results are reinforced as we find flight-to-safety effects during the crisis period in the German bond market that are not driven by macroeconomic news. This effect is not present prior to the crisis and, interestingly, is also not visible in the French bond market.

By estimating contagion dynamics before and after the announcement of bailouts (EAP), we find that the magnitude of contagion, the interlinkages across countries, and the timeline of a shock absorption is strongly reduced when the country which received the bailout is also the origin of the shock. Our findings imply that the introduction of the EAPs had a positive effect on financial market's risk perception for the individual country under bailout (EAP). However, for the first three EAPs we do not find an effect on the joint systemic risk contagion among GIIPS countries after the bailout events. The exception is the fourth bailout which was the announcement of the preliminary draft of the second bailout for Greece. This draft addressed previous shortfalls and for the first time the EFSF was extensively enlarged. As a result, we find a stabilising effect across all GIIPS countries as the joint credit risk contagion and interlinkages are dramatically reduced.

As stated by Gyntelberg et al. (2013), the fact that CDS premia are more responsive to new information may reflect the fact that the market participants in these markets

[18] See the contagion definitions according to the ECB's contagion identification strategy as specified by Constancio (2011). Furthermore, see also Forbes (2012) and Kaminsky et al. (2003), as discussed in Section 2.

on average are more highly leveraged, are more aggressive in taking positions and hence respond more quickly to new information. Thus it is crucial for policy makers and regulators to understand the dynamics in the market for sovereign credit risk, especially in the derivative market, where contagion effects were observed to be more severe during our analysed crisis sample.

Even though policy makers may not be interested in intraday movements in credit risk, our results show that the level impacts from the short-term dynamics are persistent. Hence, our results are important with regard to financial stability.

A The construction of synthetic asset swap spreads

An asset swap is a financial instrument that exchanges the cash flows from a given security - e.g. a particular government bond - for a floating market rate[19]. This floating rate is typically a reference rate such as Euribor for a given maturity plus a fixed spread, the ASW. This spread is determined such that the net value of the transaction is zero at inception. The ASW allows the investor to maintain the original credit exposure to the fixed rate bond without being exposed to interest rate risk. Hence, an asset swap on a credit risky bond is similar to a floating rate note with identical credit exposure, and the ASW is similar to the floating-rate spread that theoretically should be equivalent to a corresponding CDS spread on the same reference entity. Specifically, the ASW is the fixed value A required for the following equation to hold[20] (O'Kane (2000))

$$\underbrace{100 - P}_{\substack{\text{Upfront payment for bond}\\ \text{asset in return for par}}} + \overbrace{\underbrace{C \sum_{i=1}^{N_{\text{fixed}}} d(t_i)}_{\text{Fixed payments}} = \underbrace{\sum_{i=1}^{N_{\text{float}}} \big(L_i + A\big) d(t_i)}_{\text{Floating payments}}}^{\text{Interest rate swap}}, \tag{7}$$

where P is the full (dirty) price of the bond, C is the bond coupon, L_i is the floating reference rate (e.g. Euribor) at time t_i, and $d(t_i)$ is the discount factor applicable to the corresponding cash flow at time t_i.

In order to compute the ASW A, several observations and simplifications have to be made. First, in practice it is almost impossible to find bonds outstanding with maturities that exactly match those of the CDS contracts and second, the cash-flows of the bonds and the CDS will not coincide. To overcome these issues, in what follows we use synthetic asset swap spreads based on estimated intraday zero-coupon sovereign bond prices. Specifically, for each interval and each country, we estimate a zero-coupon curve based on all available

[19] See O'Kane (2000) or Gale (2006) for detailed discussions of the mechanics and pricing of asset swaps.

[20] This assumes that there is no accrued coupon payment due at the time of the trade; otherwise, an adjustment factor would need to be added to the floating payment component.

bond price quotes during that time interval using the Nelson and Siegel (1987) method. With this procedure, we are able to price synthetic bonds with maturities that exactly match those of the CDS contracts, and we can use these bond prices to back out the corresponding ASW. As this results in zero coupon bond prices, we can set C in Equation (7) to zero.

A CDS contract with a maturity of m years for country j in time interval k of day t, denoted as $S_j(t_k, m)$, has a corresponding ASW $A_j(t_k, m)$:

$$100 - P_j(t_k, m) = \sum_{i=1}^{N_m} \big(L_i(t_k) + A_j(t_k, m)\big) \cdot d(t_k, t_i), \tag{8}$$

with $P_j(t_k, m)$ as our synthetic zero coupon bond price.

For the reference rate L_i in Equation (8), we use the 3-month Euribor forward curve to match as accurately as possible the quarterly cash flows of sovereign CDS contracts. We construct the forward curve using forward rate agreements (FRAs) and euro interest rate swaps. We collect the FRA and swap data from Bloomberg, which provides daily (end-of-day) data. 3-month FRAs are available with quarterly settlement dates up to 21 months ahead, i.e. up to 21×24. From two years onwards, we bootstrap zero-coupon swap rates from swap interest rates available on Bloomberg and back out the corresponding implied forward rates. Because the swaps have annual maturities, we use a cubic spline to generate the full implied forward curve, thereby enabling us to obtain the quarterly forward rates needed in Equation (8).

Given our interest in intraday dynamics, we follow Gyntelberg et al. (2013) and generate estimated intraday Euribor forward rates by assuming that the intraday movements of the Euribor forward curve are proportional to the intraday movements of the German government forward curve.[21] To be precise, for each day, we calculate the difference between our Euribor forward curve and the forward curve implied by the end-of-day Nelson-Siegel curve for Germany.[22] We then keep this difference across the entire curve fixed throughout that same day and add it to the estimated intraday forward curves for Germany earlier on that day to generate the approximate intraday Euribor forward curves. This approach makes the, in our view, reasonable assumption that the intraday variability in Euribor forward rates will largely mirror movements in corresponding German forward rates.

Finally, we need to specify the discount rates $d(t_k, t_i)$ in Equation (8). The market has increasingly moved to essentially risk-free discounting using the overnight index swap (OIS) curve. We therefore take $d(t_k, t_i)$ to be the euro OIS discount curve, which is constructed

[21] Euribor rates are daily fixing rates, so we are actually approximating the intraday movements of the interbank interest rates for which Euribor serves as a daily benchmark.

[22] Here we use the second to last 30-minute interval, because the last trading interval is occasionally overly volatile.

in a way similar to the Euribor forward curve. For OIS contracts with maturities longer than one year, we bootstrap out zero-coupon OIS rates from interest rates on long-term OIS contracts. Thereafter, we construct the entire OIS curve using a cubic spline. We use the same technique as described above to generate approximate intraday OIS discount curves based on the intraday movements of the German government curve.

To gauge the potential impact of this assumption on our empirical results, we reestimate our model using an alternative assumption that the Euribor and OIS curves are fixed throughout the day at their observed end-of-day values. Under this alternative assumption, we obviously fail to capture any movements in money market rates within the day when we price our synthetic asset swaps. Our results remain robust.

B Impulse response functions for Spain and Portugal

Figure B.1: Impulse responses in the pre-crisis period - shock in Spain

For details see Figure 7.

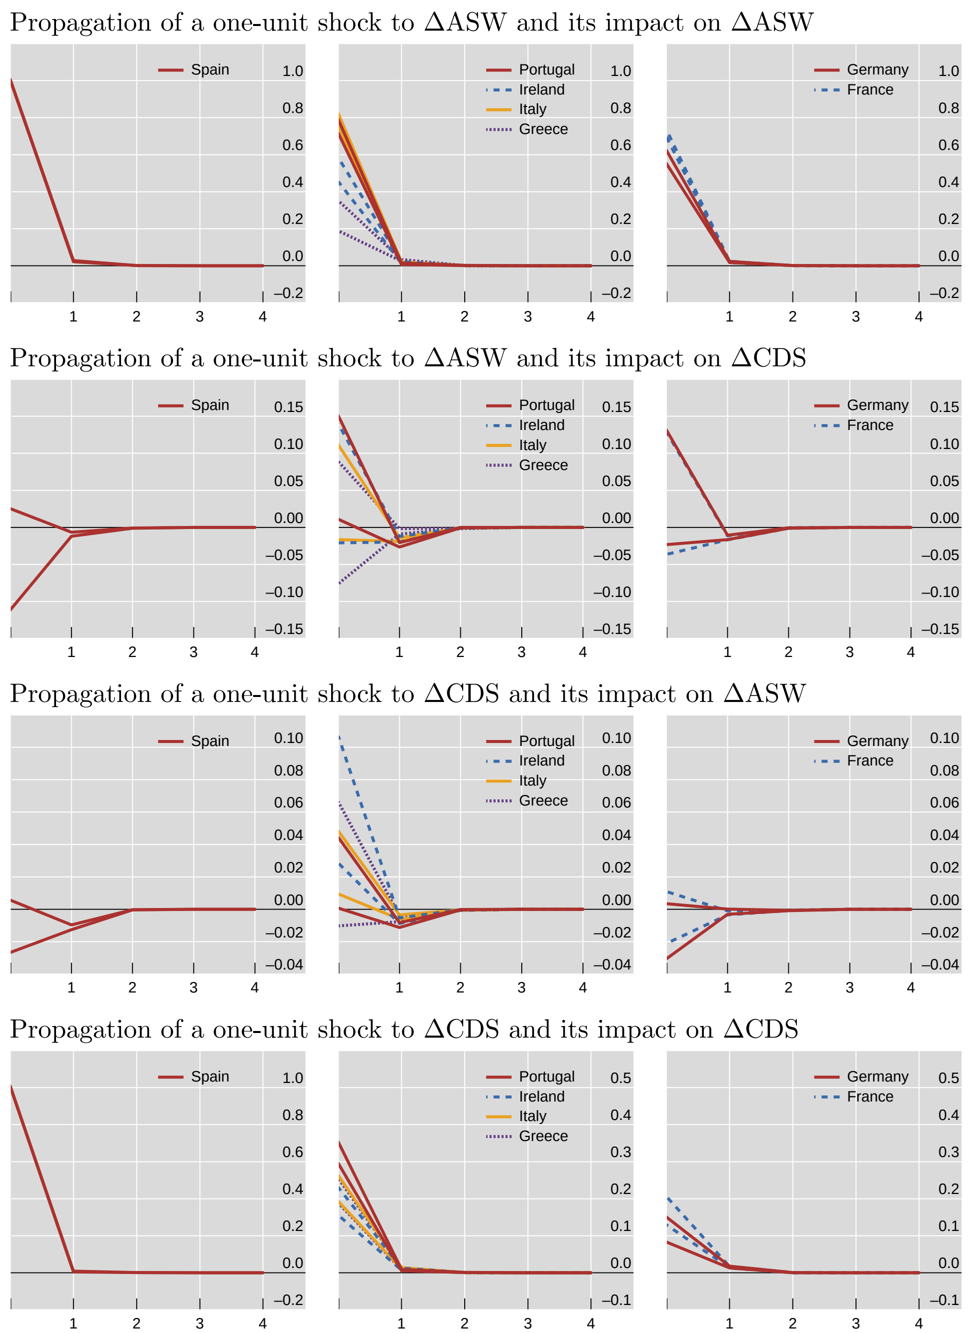

Figure B.2: Impulse responses in the crisis period - shock in Spain

For details see Figure 8.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

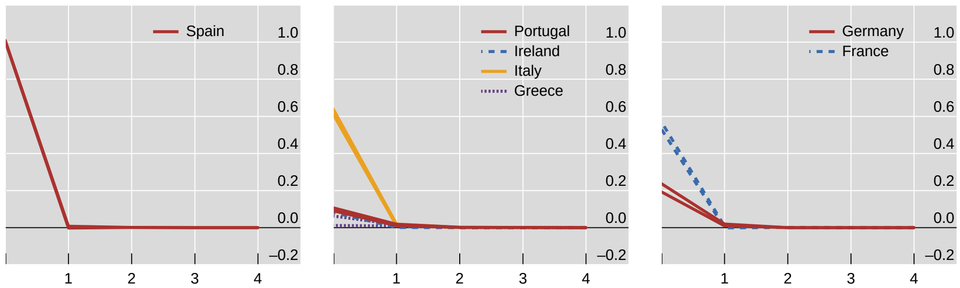

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

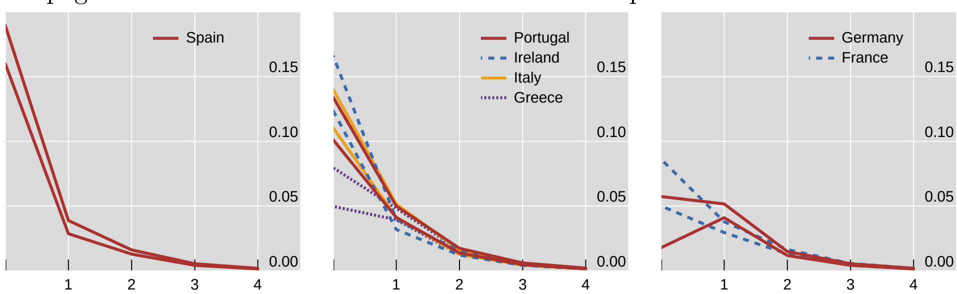

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

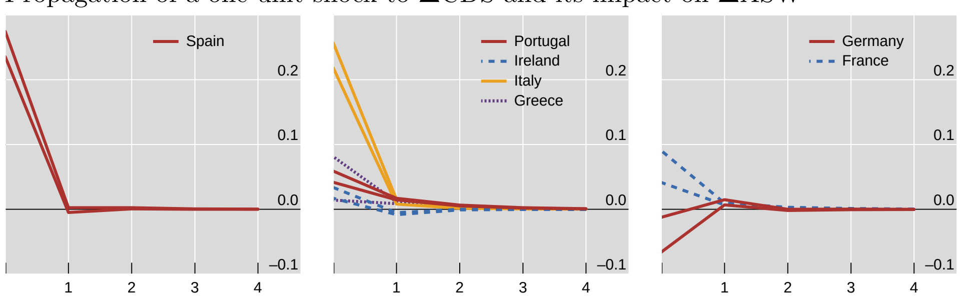

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

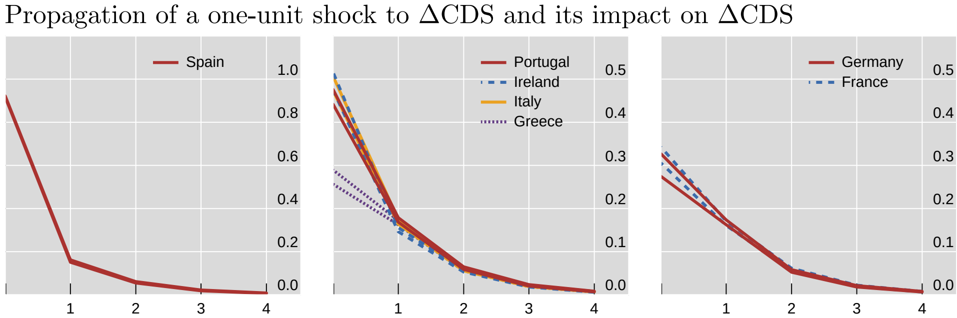

Figure B.3: Impulse responses in the pre-crisis period - shock in Portugal

For details see Figure 7.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

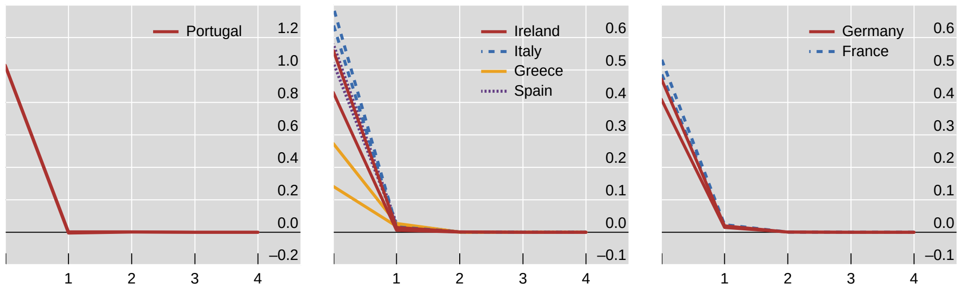

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

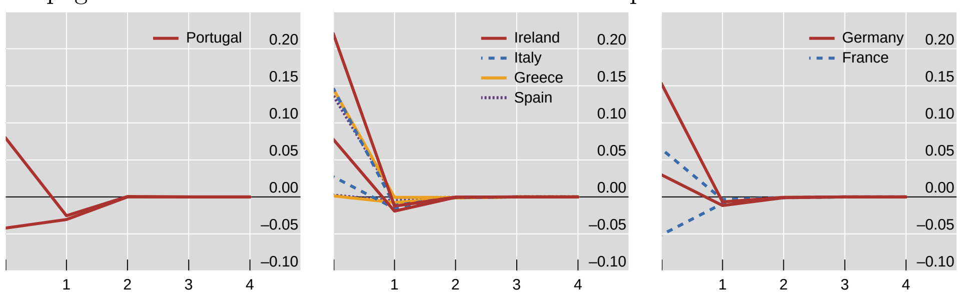

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

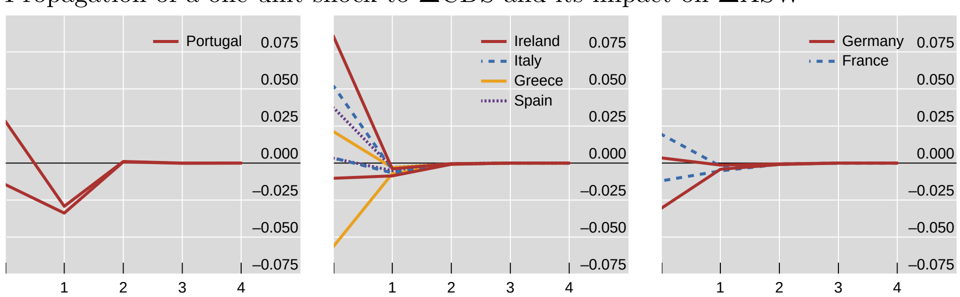

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

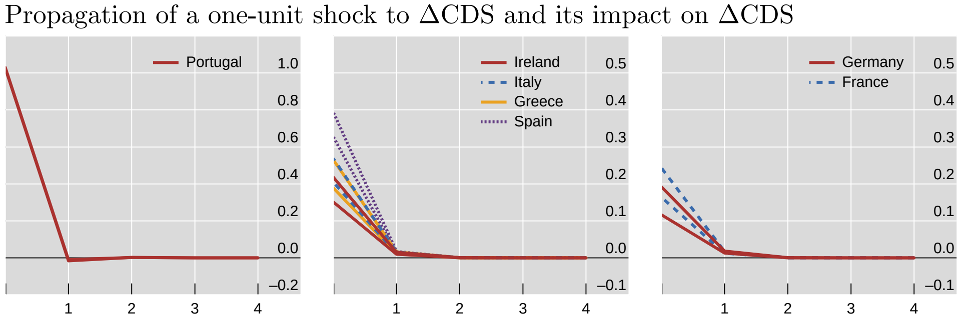

Figure B.4: Impulse responses in the crisis period - shock in Portugal

For details see Figure 8.

Propagation of a one-unit shock to ΔASW and its impact on ΔASW

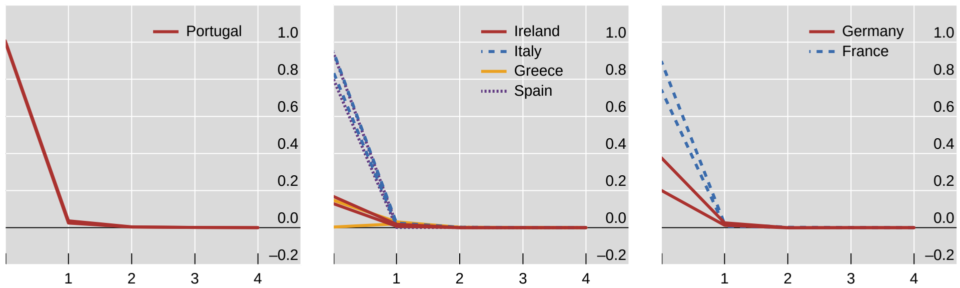

Propagation of a one-unit shock to ΔASW and its impact on ΔCDS

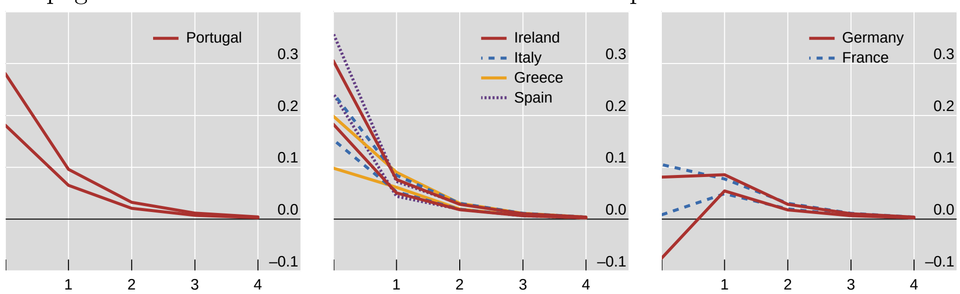

Propagation of a one-unit shock to ΔCDS and its impact on ΔASW

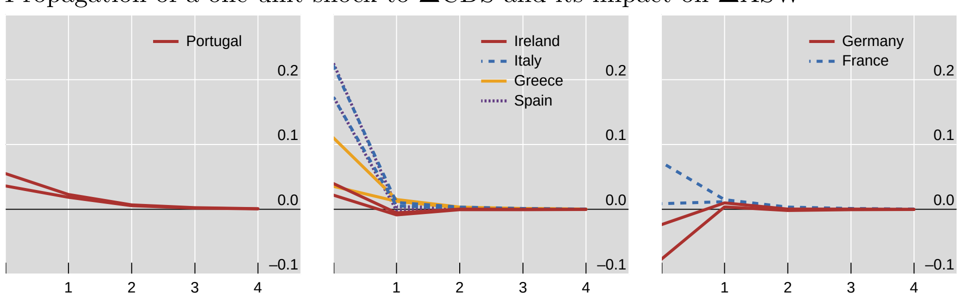

Propagation of a one-unit shock to ΔCDS and its impact on ΔCDS

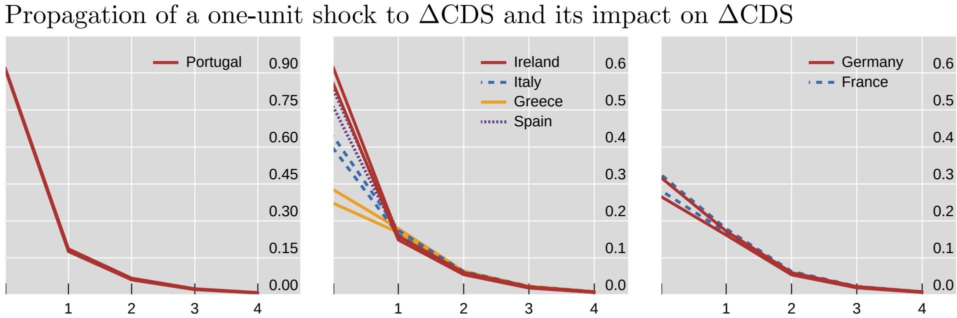

C Accumulated impulse response functions for Greece

Figure C.1: Accumulated impulse responses in the pre-crisis period - shock in Greece

This figure illustrates the accumulated impulse response for CDS and ASW to a one-unit shock (increase) for the period from January 2008 to 19 October 2009. For further details see Figure 7.

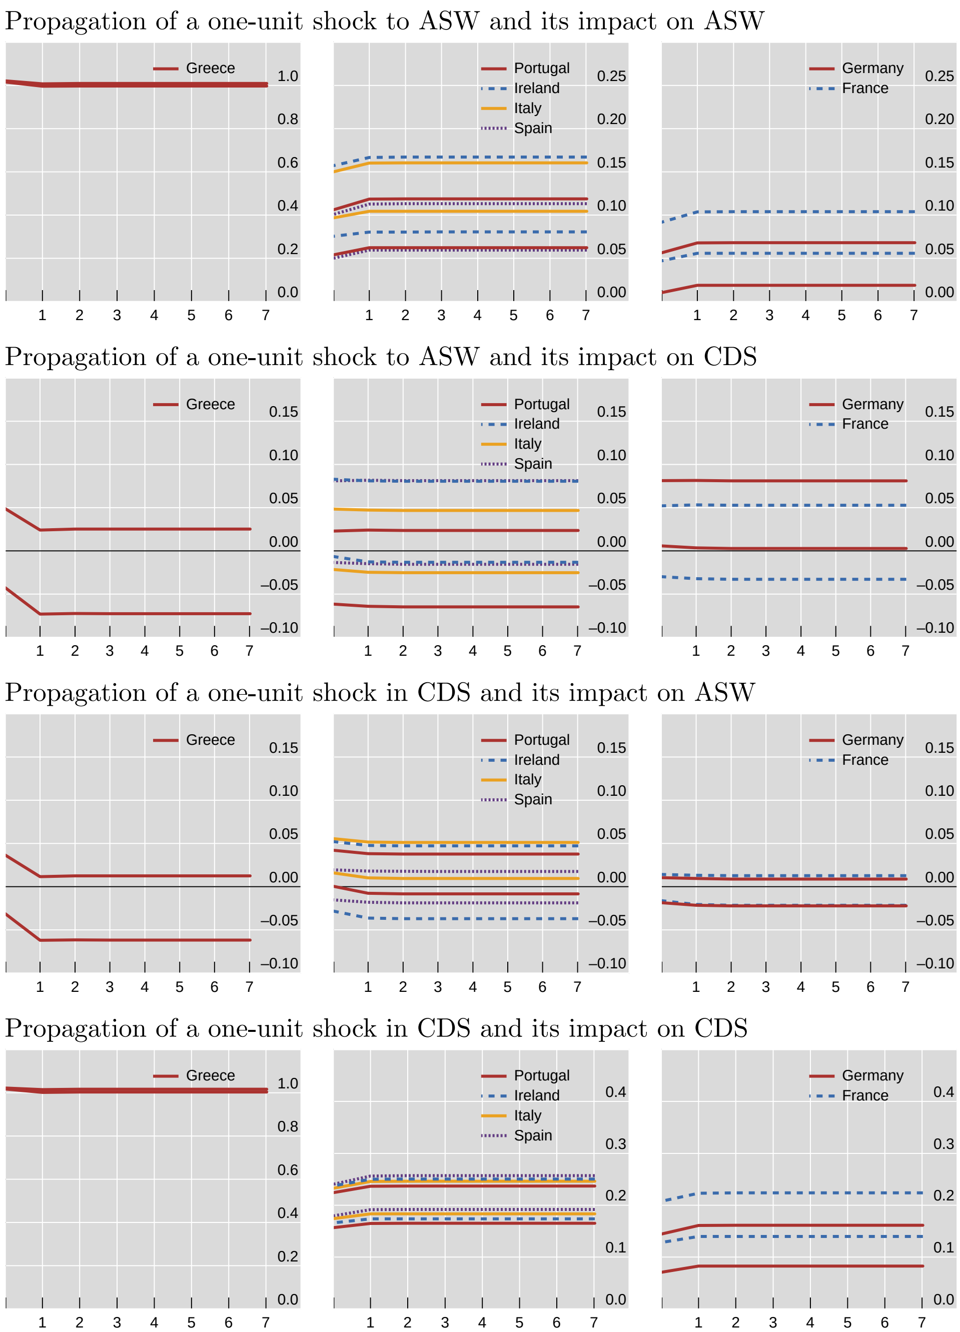

Figure C.2: Accumulated impulse responses in the crisis period - shock in Greece

This figure illustrates the accumulated impulse response for CDS and ASW to a one-unit shock (increase) for the period from 20 October 2009 to end-December 2011. For further details see Figure 8.

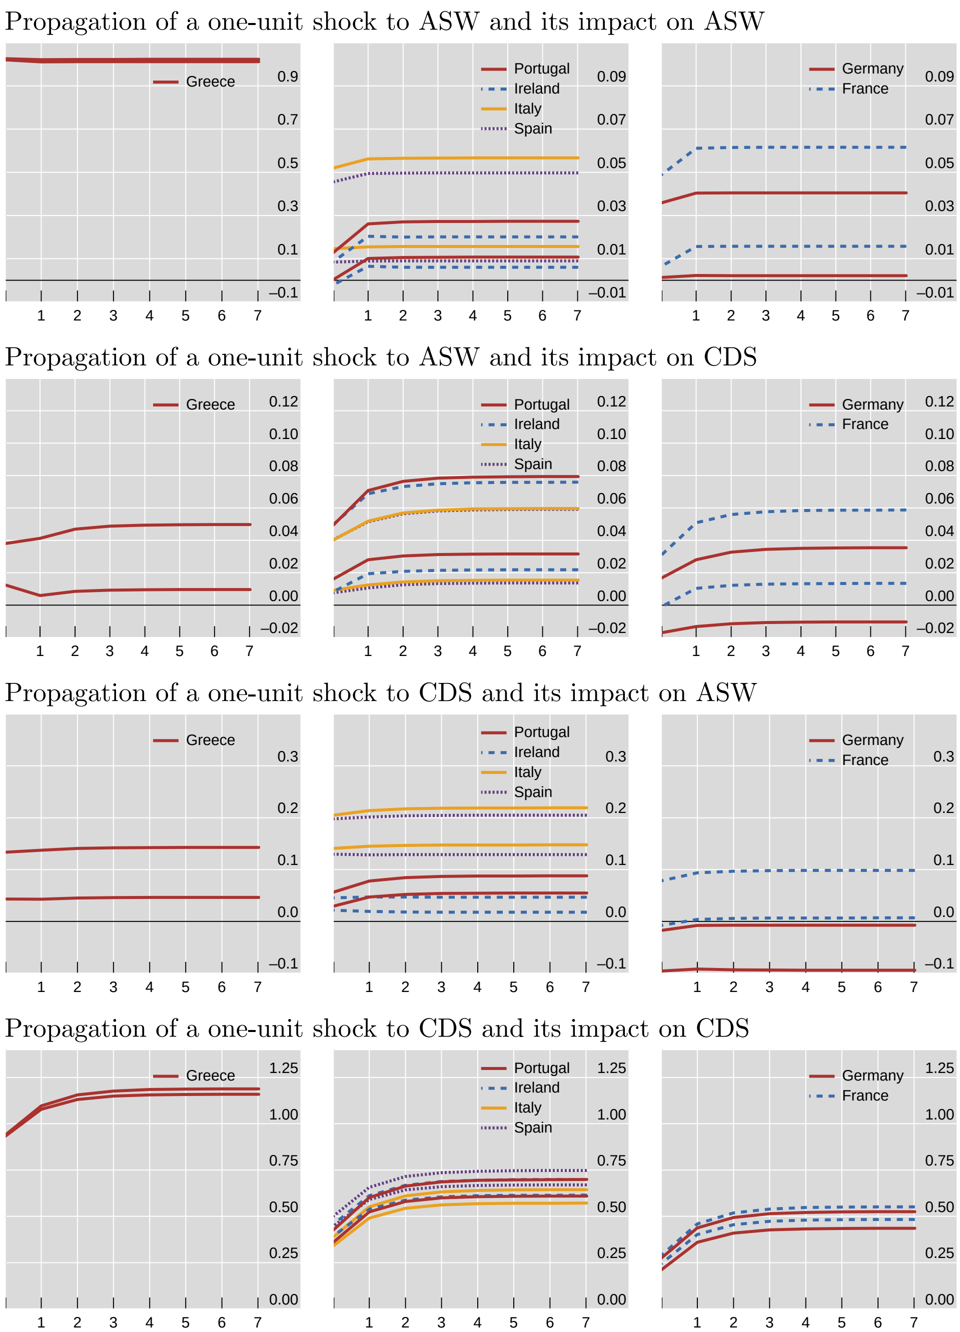

Part Ic: Credit risk contagion before and during the euro area sovereign debt crisis: Evidence from central Europe[1]

1 Introduction

The 2008-09 financial crisis as well as the subsequent euro area sovereign debt crisis created growing concerns amongst policy makers and investors of potential sovereign credit risk spillovers amongst European countries. Also in central Europe, countries such as the Czech Republic, Hungary, Poland and Slovakia (which together form the Visegrad group) experienced a growth in CDS premia with the onset of the euro area sovereign debt crisis. We ask the question whether sovereign credit risk shocks originating from the GIIPS[2] countries had an effect on the Visegrad group member countries.

The motivation for this paper is to understand how heavily central European countries were affected by the financial turmoil in peripheral Europe. In contrast to the GIIPS countries, the Visegrad group member countries were able to maintain solid public finances and high ratings (except for Hungary) throughout the euro area sovereign debt crisis. However, they still experienced a growth in CDS spreads in 2010. Papers such as Abeysinghe and Forbes (2005), Brutti and Saure (2015), Chinn and Forbes (2004) and Eichengreen et al. (1996) claim that strong trade linkages among countries increase cross-country interdependencies. As the Visegrad group member countries are important trading partners for the rest of the European Union, we should theoretically find that contagion effects from shocks in GIIPS countries lead to higher sovereign risk in the Visegrad group, due to their interdependence based on trade linkages as implied by existing research.

Most researchers so far concentrated on contagion analyses for GIIPS countries which were most affected by the euro area sovereign debt crisis. Examples of such research can be found in Bai et al. (2015), Neri and Ropele (2013), De Santis (2012) and Caporin et al. (2012). One of the few studies that focuses on central Europe is Kliber (2014) who investigates contagion effects during the Greek and Hungarian crisis on the Czech Republic, Poland and Hungary using stochastic volatility models. She reports a significant increase in volatility of Polish and Hungarian CDS as a response to the Greek crisis. Furthermore, Kliber (2014) finds that regional effects are unable to explain the periods of high volatility and correlation growth in the Czech Republic and Poland. However, her analysis does only incorporate the beginning of the Greek crisis and does not comprise the effects of the subsequent euro area sovereign debt crisis and its effect on central Europe.

---

[1] This chapter is joint work with Kristyna Ters (University of Basel) and has been accepted for publication in the *International Review of Economics & Finance* as Ters, K. and Urban, J. (2017) Intraday dynamics of credit risk contagion before and during the euro area sovereign debt crisis: evidence from central Europe.

[2] GIIPS refers to Greece, Ireland, Italy, Portugal and Spain

Another study is presented by Komarek et al. (2016) who find very high impulse response functions in terms of magnitude during the euro area sovereign debt crisis from the GIIPS to the Visegrad group member countries. They state that this finding can be most likely attributed to the low levels of debt-to-GDP ratios and the stable investment grade ratings of the Visegrad group member countries compared to the GIIPS countries. Komarek et al. (2016) and Claeys and Vasicek (2014)[3] only superficially investigate the effects on central Europe as their research focuses on GIIPS countries.

Moreover, Ters and Urban (2017b) show, using an event study, that providing financial support to Greece, Ireland and Portugal, had a stabilising effect on systemic risk for the highly indebted peripheral countries in the euro zone. In this paper we want to understand whether the same stabilisation can be seen or was even necessary for the central European countries. We are extending the studies done by Komarek et al. (2016) and Ters and Urban (2017b) by focusing on central European countries and performing an in-depth analysis on the effects of sovereign risk contagion on the Visegrad group countries. We also include Austria, France and Germany as control/riskfree countries in our analysis of sovereign credit risk contagion from the GIIPS countries to central European countries. An additional reason for the inclusion of Austria is its proximity to the Visegrad group member countries. Our data sample spans from January 2008 to end-December 2011 and we split the period in a pre-sovereign debt crisis and the sovereign debt crisis period and hence will be able to analyse changes in the dynamics of sovereign credit risk.[4]

We define the transmission of sovereign risk as the propagation of a financial shock in one country onto another country as for example in Grammatikos and Vermeulen (2012). They analyse the transmission of the financial and the sovereign debt crises between the US and the euro area by including stocks and CDS data of financial and non-financial corporate CDS. Also in line with Claeys and Vasicek (2014), Grammatikos and Vermeulen (2012) and Kliber (2014) we do not aim to find specific drivers of contagion as for example Adrian and Brunnermeier (2011) or Longstaff et al. (2011). A large body of literature concerns itself with the potential reasons and transmission channels for contagion as well as with theoretical modelling of contagion. Another strand of the literature focuses on empirical tests for the existence of contagion in a given stress period, that is, it asks if there are stronger cross-market linkages in times of crisis. Our paper belongs to the latter type, as we focus on testing for the existence of contagion during the euro area sovereign debt crisis.

[3] Claeys and Vasicek (2014) find substantial contagion effects only between GIIPS countries during the euro area sovereign debt crisis.

[4] The euro area sovereign CDS markets were very thin prior to 2008, which makes any type of intraday analysis before 2008 impossible. Further, the ban on naked sovereign CDS trading in Europe in 2012 introduces a structural break, which forces us to end our sample in December 2011.

By using a panel VAR methodology we are able to quantify the shock transmission in terms of speed and magnitude from the GIIPS to the central European countries and hence can qualify the transmission as either contagion or comovement. Existing research has differentiated between cross-country and intra-country analysis. By using a panel VAR methodology we can control for both country-specific risk and contagion effects across countries. Panel VARs are built on the same logic as standard VARs, but, by adding a cross-sectional dimension, they become a much more powerful tool for addressing policy questions of interest related, for example, to the transmission of shocks across borders as stated in Fomby et al. (2013) and Canova and Ciccarelli (2013). By using the method of Canova and Ciccarelli (2013), we are able to shock the credit risk of an individual country and derive the individual response for each country in the panel. Standard VAR models assume by construction that the individual country shocks are uncorrelated across different countries which is particularly unrealistic in a cross-border contagion detection application in macroeconomics as also stated by Stock and Watson (2016).

We are using credit default swaps (CDS) as a measure for credit risk to understand the dynamics in sovereign risk. CDS were claimed to be the primary source for contagion during the euro area sovereign debt crisis. Indeed, the ban on outright short selling of sovereign CDS in the euro area which came into force in November 2012, was introduced following overreactions in CDS prices due to speculators (Sarkozy et al.; 2010).

A substantial advantage of our analysis is the use of intraday CDS data which allows us to capture the intraday patterns of credit risk contagion and dramatically increase the accuracy of the estimated model. Indeed, shocks that may seem to affect several countries simultaneously when viewed at a daily or lower data frequency are revealed, through the lens of intraday data, to have possible origins in one particular country with clear contagion effects on other countries. Also, Gyntelberg et al. (2013) discuss the advantages of using intraday data. They find that when using daily data, due to the dramatically smaller number of observations, the confidence bands of the estimated coefficients are extremely wide and therefore in most cases not significant. Further, Gyntelberg et al. (2013) point out that sovereign credit risk dynamics follow an intraday pattern.

To our knowledge, for the first time, this paper focuses on the identification of sovereign risk contagion emanating from GIIPS onto the Visegrad group member countries. We find that shock transmission from GIIPS to the Visegrad group countries was moderate and had only the characteristic of comovement, unlike the contagion dynamics amongst GIIPS countries.

Further, as a new contribution we study the effects of the economic adjustment programmes (EAP) of the Troika, a decision group formed by the European Commission (EC), the European Central Bank (ECB), and the International Monetary Fund (IMF), in order to analyse the effect of bailouts on sovereign credit risk transmission onto the

Visegrad group countries. The pre- and post-bailout shock transmissions from the GIIPS to the Visegrad group countries were moderate. By employing an event study framework we find a stabilising effect following the bailouts on the Visegrad group as sovereign credit risk moderately declined in all Visegrad group member countries.

The results of the event studies are consistent with our findings for the sovereign debt crisis period. Slovakia was the only exception during the period around the first Greek bailout, where we find strong interlinkages in terms of magnitude of shock transmission from the GIIPS countries prior the bailout and a strong stabilisation after the bailout. This result is most likely attributed to Slovakia's euro membership.

The remainder of the paper is structured as follows: Section 2 provides an overview on the recent literature and defines contagion, as well as compares and contrasts contagion to comovement. Section 3 discusses the economic importance of the Visegrad group and Section 4 presents our data. Section 5 explains the set-up and estimation of the panel VAR (PVAR) model. Section 6 presents the empirical results and Section 7 concludes.

2 Contagion literature review and definition

The issues of financial shock contagion among countries during the sovereign debt crisis have figured prominently in recent empirical research. Caporin et al. (2012) analyse risk contagion using the CDS spreads of the major euro area countries using different econometric approaches such as Bayesian modelling. They find that the diffusion of shocks in euro area CDS has been remarkably constant, while the risk spillover among countries is not affected by the size of the shock. Komarek et al. (2016) and Ters and Urban (2017b) look at GIIPS countries and contagion effects. They use both CDS and asset swap spread data in order to differentiate the credit risk channel and find that the CDS market plays the most important role during the crisis. Severe contagion effects amongst GIIPS countries are found during the crisis. Other examples are Bai et al. (2015), Neri and Ropele (2013), De Santis (2012), and Arghyrou and Kontonikas (2012). They all employ time series modelling approaches for contagion and include sovereign bond spreads (yield to maturity) to reflect pure credit risk considerations and macroeconomic variables. The results are mostly discriminated in terms of core (such as France and Germany) and peripheral countries (GIIPS). In general, these authors find that the bond spreads of lower-rated countries increase along with their Greek counterparts. This is also in line with our findings. However, their results in terms of magnitude, responses to shocks and contagion effects on core countries are somewhat mixed. Similarly to these studies, Koop and Korobilis (2016) employ an enhanced panel approach for empirical modelling of financial contagion across countries based on Canova and Ciccarelli (2013).

As pointed out by Corsetti et al. (2011), there is much disagreement among some economists about the exact definition of contagion and how it should be tested. Kaminsky et al. (2003) describe contagion as an episode in which there are significant immediate effects in a number of countries following an event, such as when the consequences are fast and furious and evolve over a matter of hours or days. When the effect is gradual, Kaminsky et al. (2003) refer to it as comovement rather than contagion. Also Constancio (2011) states that there is no general agreement about a clear contagion identification strategy and about which criteria are necessary or sufficient to characterise a contagion event. However, he specifies that the ECB identifies contagion amongst others according to: (i) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); and (ii) the events constituting contagion are negative extremes. Our contagion identification strategy builds upon these criteria described by the ECB. We restrict the identification of contagion to the case where both conditions (i) and (ii) must be fulfilled. Otherwise, when the effect is gradual and not a negative extreme, we refer to it as comovement.

3 Economic importance of the Visegrad group

Our sample consists of 12 countries, which can be grouped as: i) the GIIPS countries: Greece, Ireland, Italy, Portugal and Spain; ii) the Visegrad group member countries: Czech Republic, Hungary, Poland and Slovakia; and iii) the control/riskfree group: Austria, France and Germany. The GIIPS countries have been most affected by the euro sovereign debt crisis and hence are a potential source of risk contagion to the Visegrad group. Austria, France and Germany are included as a control group in order to judge the relative importance of the computed results.

The Visegrad group is an important alliance of central European countries, established in 1991 in order to foster economic cooperation. The Visegrad group has a GDP (PPP) of approximately USD 1.8 trillion and 65 million inhabitants (2015). Furthermore, the GDP of the Visegrad group amounts to 75% of the combined GDP of the Visegrad group plus all Eastern European countries which are EU members[5].

Public finances of the Visegrad group member countries remained very robust over the past decade. Their debt-to-GDP ratio is far below 100%, unlike in the case of the GIIPS countries, where public debt increased dramatically since 2008, despite several economic adjustment programmes. Apart from France, the debt-to-GDP in the control group countries also remained stable over the past decade (see Figure 1).

[5] In addition to the Visegrad group countries, these are Bulgaria, Estonia, Latvia, Lithuania and Romania.

Figure 1: Public debt-to-GDP ratios

The figures show debt-to-GDP ratios in percent for the individual countries in our sample. Source: BIS and ECB.

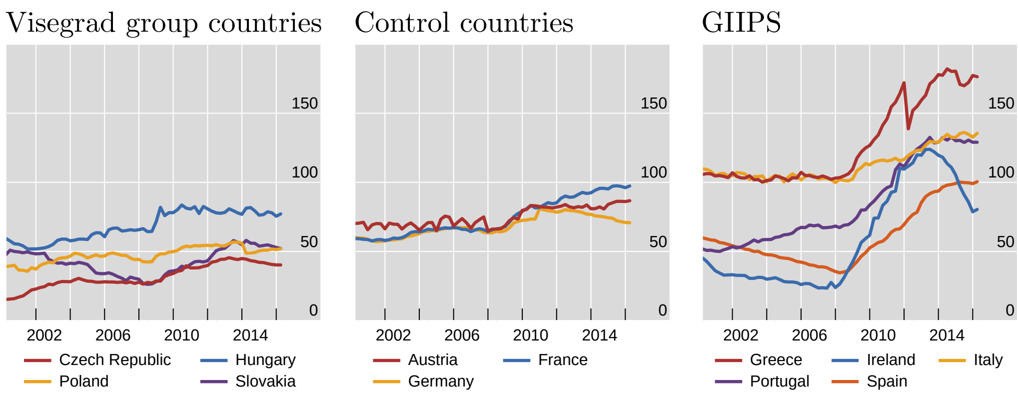

The Visegrad group is an important trading partner to the euro area. Figure 2 shows the trade linkages of each of the Visegrad group member countries with our control countries (Austria, France and Germany), with the GIIPS countries[6] and amongst each other. Even though, the trade linkages with Germany and amongst the Visegrad group are by far the strongest, the trade linkages to Austria, France and the GIIPS are around 10-20% of each Visegrad group member countries' GDP (see panel A in Figure 2).

We have included Austria in our control/riskfree group due to its geographic and economic proximity to the Visegrad group. Austria's trade linkage with each Visegrad group country is approximately USD 10-15 billion (see panel B in Figure 2) and hence Austria's trading relations with the Visegrad group are of similar importance as France and the GIIPS countries' trade linkages with the Visegrad group.

[6] Amongst the GIIPS the main trading partners of the Visegrad group are Italy and Spain.

Figure 2: Trade linkages

The figures show the total trade linkages of the Czech Republic, Hungary, Poland and Slovakia vis-a-vis Austria, Germany, France, the GIIPS countries and the Visegrad group countries (excluding the country plotted). The panel A presents the total trade linkages as percent of the countries' GDP and the panel B the total trade in USD billion. Total trade is the sum of exports and imports. Source: Bloomberg.

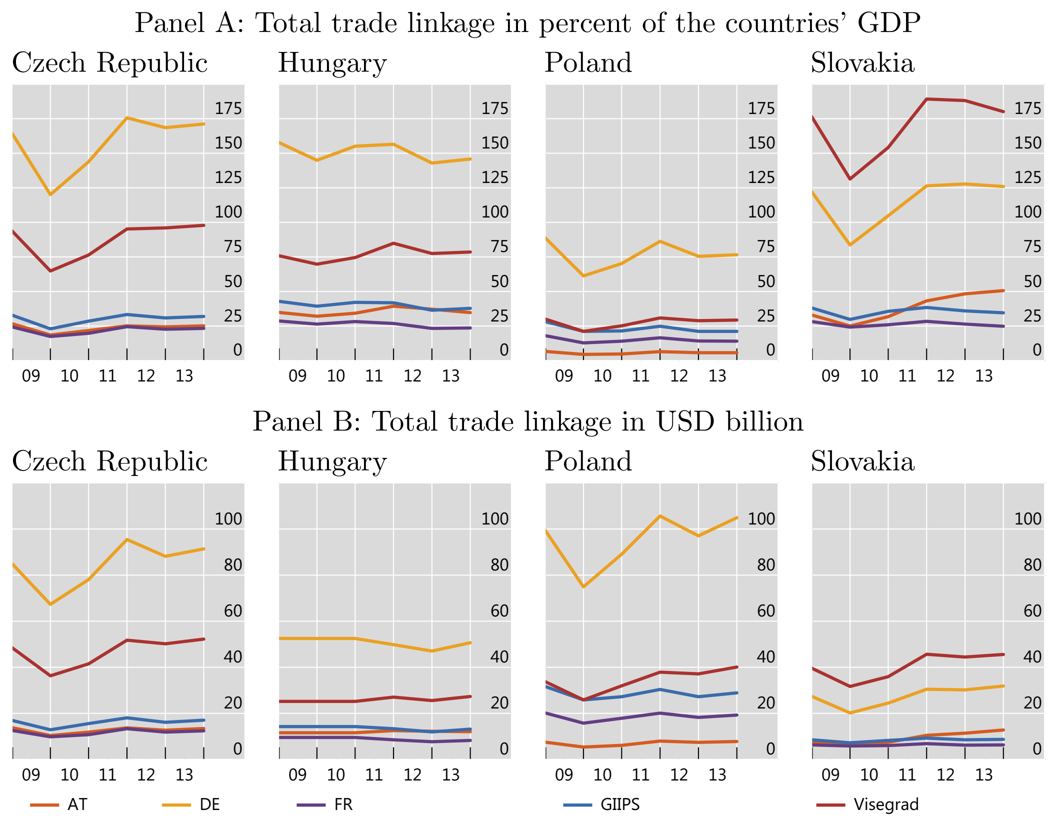

In addition, Austria is a near riskfree country (see ratings in Figure 3) and hence helps to gauge the magnitude and the timeline of the impulse responses of the individual Visegrad group member countries following a shock from the GIIPS countries (see Section 6).

Figure 3: Foreign currency long-term ratings

The figure shows average ratings from Fitch, Moody's and S&P. The Visegrad group countries (lhs) maintained (with Hungary as a slight exception) a fairly stable rating. The risk free or control countries maintained a high rating, even throughout the crisis period. The ratings of the GIIPS countries fell from an initial high level to a low medium grade and in the case of Portugal and Greece to below investment grade. Source: Fitch, Moody's, S&P, authors' calculation.

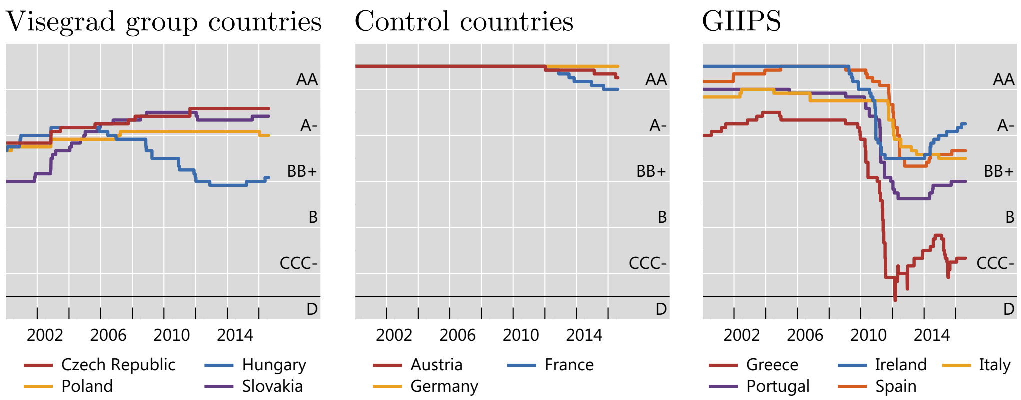

The composition of the respective groups i) to iii), mentioned at the beginning of this section, was chosen based on similar debt-to-GDP ratios (see Figure 1), similar credit ratings (see Figure 3) and similar default probabilities as priced by CDS spreads of the individual countries (see Figure 4). The control/riskfree group countries receive sustained top ratings from Fitch, Moody's and S&P. The GIIPS countries experienced dramatic downgrades since the onset of the euro area sovereign debt crisis, while the Visegrad group countries maintained their high ratings on balance throughout the crisis period (except for Hungary). This applies both to the foreign currency and the local currency long-term ratings. The foreign currency long-term ratings are presented in Figure 3. The local currency long-term ratings (not presented here) show an almost identical pattern.

In contrast to the GIIPS countries, the Visegrad group member countries were able to maintain solid public finances and high ratings (except for Hungary) throughout our sample period. Still, all Visegrad group member countries experienced a growth in CDS spreads starting with the onset of the crisis as also stated by Kliber (2014) and displayed in Figure 4. This increase in sovereign risk in the Visegrad group during the euro area sovereign debt crisis period cannot be attributed to changing fundamentals. Also the downgrade of Hungary during the Hungarian crisis in 2010 was not able to explain the increase in sovereign CDS spreads in the Visegrad group as found by Kliber (2014). She argues that the impact of the Hungarian turmoil is negligible as the Hungarian CDS market did not contribute to the changes in correlation[7] dynamics in Poland and the Czech Republic. To extend the existing research we aim to analyse whether contagion dynamics emanating from the GIIPS countries were responsible for the increase in sovereign credit risk across the Visegrad group (Figure 4). Based on existing contagion research such as Abeysinghe and Forbes (2005), Brutti and Saure (2015), Chinn and Forbes (2004) and Eichengreen et al. (1996) we should see a strong interconnectedness through trade linkages from the GIIPS countries onto the Visegrad group. This should, according to theory lead to cross-country contagion.

In addition, we aim to analyse whether the economic adjustment programmes (EAP) by the Troika had a stabilising effect on the sovereign risk of the Visegrad group. Following the framework proposed by Ters and Urban (2017b), who find that after a bailout event, systemic risk decreases in the GIIPS countries while the strongest stabilisation is visible with the second bailout for Greece.

Figure 4: CDS spreads in basis points

The figures are based on data with a 30-minute sampling frequency. Our split into the pre-sovereign debt crisis and the crisis period is indicated by the vertical line in each figure. CDS markets for Greece became increasingly illiquid at the end of 2011, hence our data for Greece ends in September 2011. Source: CMA, authors' calculations.

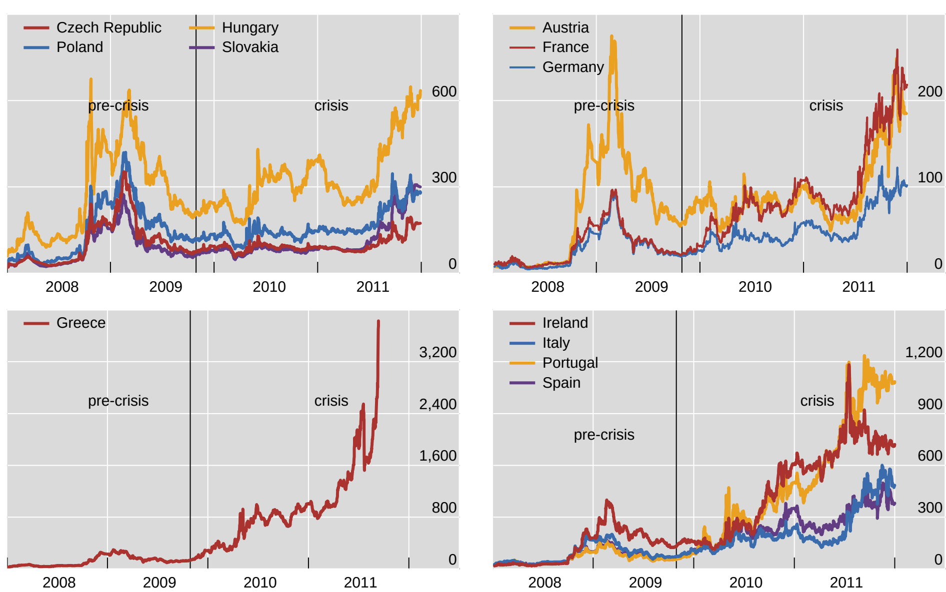

7 Hungary was, however, able to contribute to an increase in volatility in Poland and the Czech Republic.

4 Data

Our data on sovereign credit risk consists of USD-denominated five-year maturity CDS spreads for all countries in our sample. We have intraday CDS price quotes from CMA (Credit Market Analysis Ltd.) Datavision. CMA Datavision continuously gathers information on executable and indicative CDS prices directly from the largest and most active credit investors. After cleaning and checking the individual quotes, CMA applies a time and liquidity weighted aggregation so that each reported bid and offer price is based on the most recent and liquid quotes. Being an OTC market, the CDS market is in principle open 24 hours a day. In practice, however, most of the activity in the CMA database is concentrated around the period from 8:30 to 17:30 CET/CEST which we are using in our analysis.

We focus on the five-year maturity segment, which represents the most liquid CDS contracts (see Gyntelberg et al. (2013) for an in-depth discussion of our data). The use of intraday data in our empirical analysis enables us to obtain much sharper estimates and clearer results with respect to the market mechanisms as also shown in Gyntelberg et al. (2013), who also point out that sovereign credit risk dynamics follow an intraday pattern. They find that when using daily data, due to the dramatically smaller number of observations, the confidence bands of the estimated coefficients are extremely wide and therefore in most cases not significant. Therefore, we are able to capture the intraday patterns of credit risk contagion and dramatically increase the accuracy of the estimated model. Indeed, shocks that may seem to affect several countries simultaneously when viewed at a daily or lower data frequency are revealed, through the lens of intraday data, to have possible origins in one particular country with clear contagion effects on other countries. Additionally, the use of intraday data enables us to carry out an event study (presented in Section 6.2) with a short pre- and post-event window and still keeping a sufficient number of observations.

We construct our intraday data on a 30-minute sampling frequency on our data set, which spans from January 2008 to end-December 2011. Using this sampling frequency each individual time series has 18,793 time stamps, with especially in 2010 and 2011 over 90% of the 30 minute intervals containing non-missing values (see right-hand graph in Figure 5). This intraday frequency ensures estimation results with high statistical significance, while avoiding microstructural noise effects. Microstructural noise effects may come into play when high frequency data is used (Fulop and Lescourret (2007)). However, this does not apply to our data based on a 30-minute sampling frequency because we average the reported quotes over each 30 minute interval (for tests, robustness checks and for a more detailed discussion please refer to Gyntelberg et al. (2013)). We have tested that at a 5-

minute sampling frequency the first signs of microstructural noise are coming into effect, as seen by realised variance patterns.

The available number of indicative quotes for CDS does not allow a data frequency higher than 30 minutes during the earlier years of our sample period (eg 2008 and 2009). The euro area sovereign CDS markets were very thin prior to 2008, which makes any type of intraday analysis before 2008 impossible. Further, the ban on naked sovereign CDS trading in Europe in 2012 introduces a structural break, which forces us to end our sample in December 2011.

Figure 5: CDS data from CMA Datavision – Liquidity

The left-hand panel shows the number (in thousands) of data ticks per year. The right-hand panel shows the number (in thousands) of non-empty half-hour intervals per year (right scale). We consider 18 half-hour slots per trading day, from 8:30 to 17:30 CET/CEST. The left scale in the right-hand panel shows the percentage of 30 min. intervals which contain at least one data tick during the 18 daily half-hour intervals we consider. Source: CMA Datavision

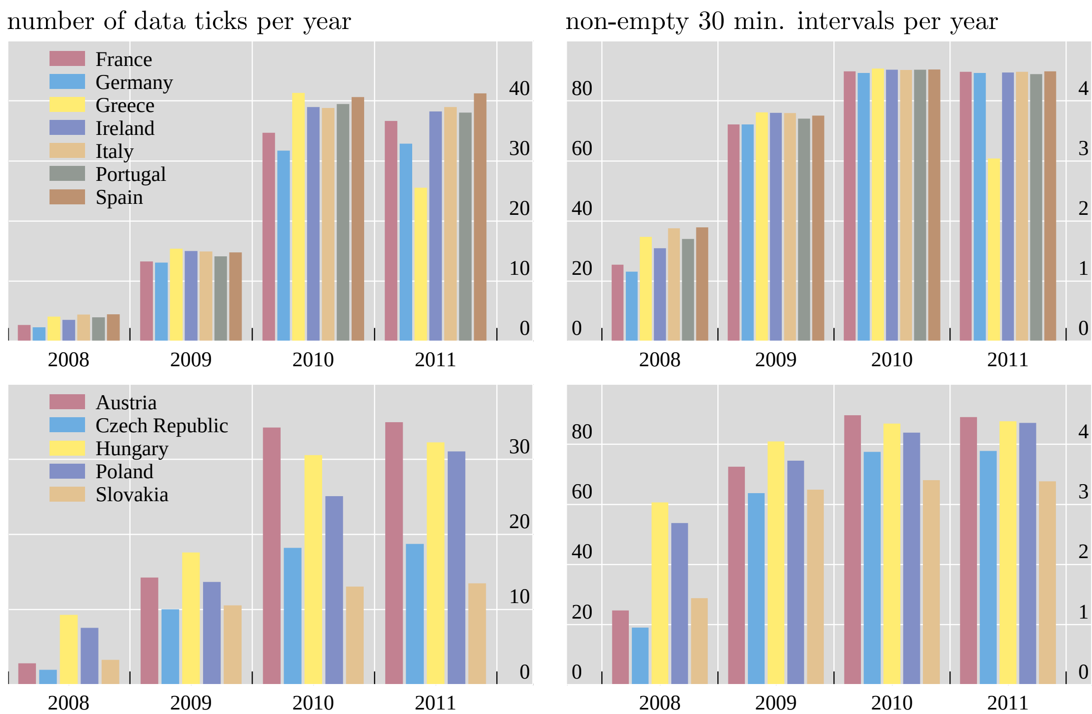

USD-denomination is standard for euro area sovereign CDS. The rational behind choosing a currency which is not the country's legal tender is that in any restructuring or default event the legal tender of the country under consideration is expected to weaken significantly.[8] Thus, the USD-denomination is an additional hedge which leads to the fact that on one hand, they are much more liquid than EUR-denominated CDS and on the other hand, they trade at a higher premium. As we only use USD-denominated CDS for all countries in our sample we do not have any distortions due to the currency denomination.

Even though the euro area sovereign debt crisis has not ended in 2012, our data period still covers the time period of the most turbulent stage of the crisis and contains the first four economic adjustment programmes (bailouts) of the Troika[9].

In Figure 5 we present information on CDS liquidity such as the number of ticks and the number of non-empty half-hour intervals. As already mentioned, our chosen sampling frequency of 30 minute intervals ensures that our data contains over 90% non-missing values as presented in the right hand panel. The left hand panel in Figure 5 presents the number of data ticks per year. Of particular interest is, that CDS liquidity increased during the crisis period for all countries in our sample. Also, number of data ticks and number of non-empty intervals for the Visegrad group member countries are approximately comparable to the GIIPS and control countries.

In addition to the liquidity measures for CDS in Figure 5 we also present the notional amount outstanding and the trade count reported by the ISDA in Figure 6. The data is publicly available from mid 2008 onwards. In most cases the notional amount outstanding remains high during the crisis and our overall sample period (2008 until 2011). The overall decline of the notional amounts outstanding starts in 2012, most likely driven by the ban on outright short selling of sovereign CDS. However, the trade count (lower panel in Figure 6) remains high still beyond 2011 and even increases for some countries throughout 2012.

[8] Vice versa EUR-denominated CDS are more liquid than USD-denominated CDS for US sovereign CDS.

[9] A decision group formed by the EC, ECB and the IMF.

Figure 6: CDS net notional amount outstanding and trade count.

Panel A shows the net notional amount outstanding in USD billion. Panel B shows the number (in thousands) of trades per week. The start of the period is given by the data availability. The missing data for Greece from 2012 until 2014 is due to the Greek restructuring. Source: ISDA

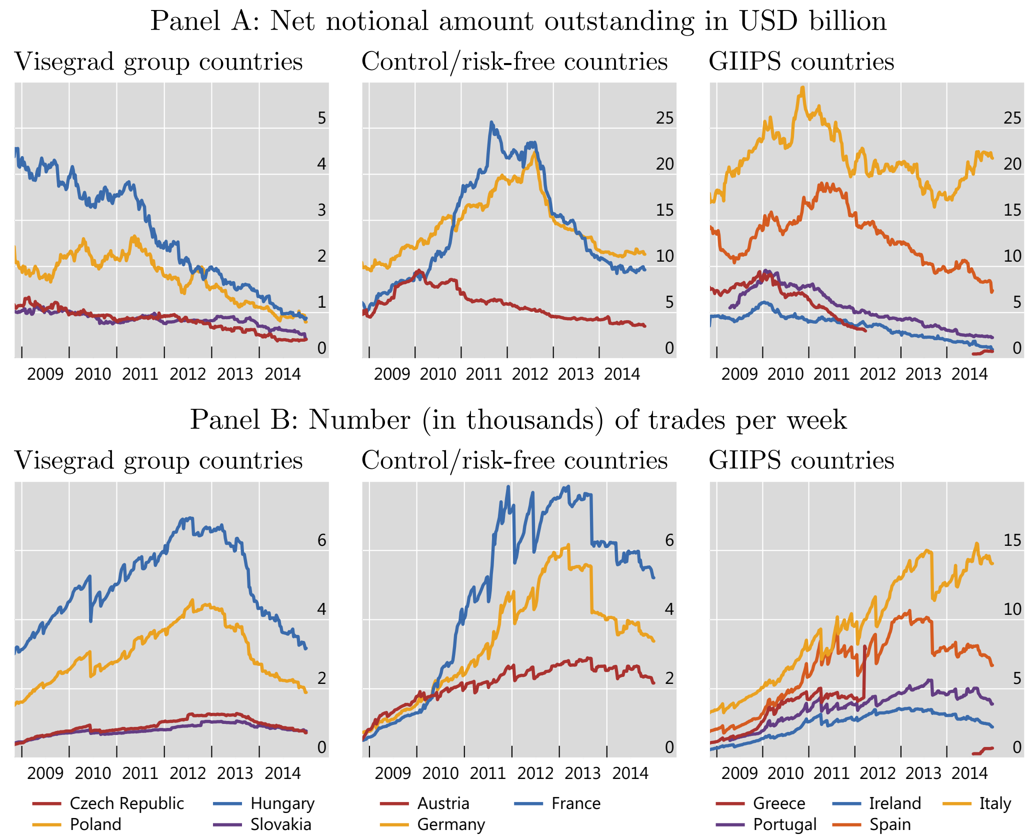

Thus, our sovereign CDS data is highly liquid during our sample period which enables us to conduct our analyses on an intraday frequency.

We split the data into two subsamples. The first subsample covers the period from January 2008 to 19 October 2009 and, as such, represents the period prior to the euro area sovereign debt crisis. While this period includes the most severe phase of the financial crisis, including the default of Lehman Brothers, it is relatively unaffected by market distortions stemming from concerns about the sustainability of public finances in view of rising government deficits and therefore represents the pre-sovereign debt crisis period. The second subsample covers the euro area sovereign debt crisis period and runs from 20 October 2009 to end-December 2011. As the beginning of the crisis period, we designate 20 October 2009, when the new Greek government announced that official statistics on Greek debt had previously been fabricated. Instead of a public deficit estimated at 6% of

GDP for 2009, the government now expected a figure at least twice as high. The split in a pre-crisis and a crisis period is relevant for our analysis in Section 6.1 in order to detect whether we find contagion or comovement emanating from the GIIPS countries onto the Visegrad group. We have also tested other dates as the start of the crisis, for example 1st April 2010, which was defined by van Rixtel and Gasperini (2013) as the begin of the sovereign debt crisis. Our results remain robust to this choice.

An interesting feature of the Visegrad group countries is that Slovakia is an euro area member, while the other countries have their own domestic currency. The GIIPS countries, which have been most affected by the sovereign debt crisis are also euro area members. Hence, our grouping allows to check whether the behaviour of the shock transmission is indeed dependent on the euro area membership or does depend on economic fundamentals.

Based on common and individual panel unit root tests our CDS price data (displayed in Figure 4) is I(1). Therefore, we estimate our subsequent models in first differences.

5 Modelling sovereign credit risk contagion by using a panel VAR model

To empirically measure the impact of euro area sovereign credit risk contagion effects we employ a panel vector autoregressive (PVAR) model. PVARs have the same structure as VAR models, in the sense that all variables are assumed to be endogenous but with the difference that a cross-sectional dimension is added to the representation. The PVAR has several advantages over individual country VARs in a time series framework. By analysing a panel of countries, we can more accurately model contagion from one country to another since the panel approach captures country-level heterogeneity. We control for cross-sectional heterogeneity by including fixed effects in the regression. According to Koop and Korobilis (2016) and Canova and Ciccarelli (2013), in this setup the PVAR is the ideal tool for examining the transmission of macroeconomic or financial shocks from one country to another. In contrast to linear VAR models we therefore do not have to assume that the individual country shocks are uncorrelated across different countries which is particularly unrealistic in a cross-border contagion detection application in macroeconomics as also stated by Stock and Watson (2016).

The PVAR is defined as follows:

$$y_{it} = A_{0i} + A_i(\mathrm{L})Y_{t-1} + u_{it}. \tag{1}$$

A_{0i} are G x 1 vectors and $A_i(\mathrm{L})$ are G x GN matrices with L being the lag operator. N represents the number of countries and G the number of variables per country. Y_t is the stacked version of y_{it}, the vector of G variables for each country $i = 1, ..., N$, ie $Y_t = (y_{1t}^{'}, y_{2t}^{'}, ..., y_{Nt}^{'})^{'}$ and $t = 1, ..., T$. Further, lags of all endogenous variables of all entities enter the equation of country i. Canova and Ciccarelli (2013) call this feature "dynamic interdependencies". The residual u_{it} is a G x 1 vector. u_{it} is generally correlated across the cross-sectional dimension i. Canova and Ciccarelli (2013) call this feature "static interdependencies". Thus the variance-covariance matrix for a PVAR has the following property $E(u_{it}u_{jt}') = \sigma_{ij} \neq 0$ for $i \neq j$, ie static interdependencies occur when the correlations between the errors in two countries' VARs are non-zero. On the other hand, dynamic interdependencies occur when one country's lagged variables affect another country's variables. Hence, the PVAR is more flexible compared to a VAR ($\sigma_{ij} = 0$ for $i \neq j$).[10]

We focus on the most dominant credit risk channel, the CDS market (for a detailed discussion of the bond and CDS market as well as their interlinkages we refer to Komarek et al. (2016)). Hence, in our case with $G = 1$, we can rewrite the PVAR in Equation (1) as:[11]

$$\Delta CDS_{it} = \begin{pmatrix} A_{01} \end{pmatrix}_i + \begin{pmatrix} A_{11} & A_{12} \end{pmatrix}_{ji} (\mathrm{L}) \Delta CDS_{j\,t-1} + \begin{pmatrix} u_1 \end{pmatrix}_{it}, \tag{2}$$

where A_{mn} are scalars and $i, j = 1, ..., N$ represents the number of countries in the cross-sectional dimension.

For the estimation, we follow the approach proposed by Canova and Ciccarelli (2009) of an unrestricted PVAR which allows for the selection of restrictions involving dynamic interdependencies, static interdependencies and cross-section heterogeneities. According to an empirical model comparison by Koop and Korobilis (2016), the proposed methodology by Canova and Ciccarelli (2009) shows the best properties compared to other PVAR approaches. Canova and Ciccarelli (2009) suggest adopting a flexible structure through a factorisation of the coefficients in Equation (2). Through the flexible coefficient factorisation, the PVAR can be rewritten as a reparametrised multicountry VAR and estimated using SUR (Canova and Ciccarelli; 2009). The advantage of this flexible factorisation is that the overparametrisation of the original PVAR is dramatically reduced while, in the resulting SUR model, estimation and specification searches are constrained only by the dimensionality of the estimated coefficient matrix (for a more in-depth discussion please refer to Canova and Ciccarelli (2009) and Koop and Korobilis (2016)).

We have tested different lag lengths in our PVAR model and find highly robust results. The magnitude of shock transmission is unaffected by the lag selection, while the time until the shock is absorbed can slightly vary by 1 interval (30 mins). Increasing the lag length reduces the accuracy of the results, due to the higher number of estimated

[10] According to Canova and Ciccarelli (2013), these features distinguish a panel VAR typically used for macroeconomics and finance from a panel VAR used in microeconomics.

[11] We use demeaned and standardised first differences of our raw CDS data.

parameters. However, up to a lag length of five intervals, the results remain highly robust as we are using intraday data with a minimum of 4000 observations within each subsample (pre-crisis and crisis).

Our contagion identification strategy builds upon the criteria described by the ECB. Even though Constancio (2011) states that there is no general agreement about a clear contagion identification strategy and about which criteria are necessary or sufficient to characterise a contagion event he however specifies that the ECB identifies contagion amongst others according to: (i) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); and (ii) the events constituting contagion are negative extremes. Thus, we restrict the identification of contagion to the case where both conditions must be fulfilled.

5.1 Impulse responses and shock identification

We carry out an impulse response analysis to investigate cross-border shock transmission in sovereign credit risk from the GIIPS countries onto the Visegrad group member countries and our control countries. Further, we analyse the differences in the magnitude, the speed and the dynamics between the pre-sovereign debt crisis and crisis sample in order to identify contagion as specified by Constancio (2011) and described in Section 5. We focus on individual country shocks propagating from GIIPS countries and analyse the impact of an unexpected one-unit shock to credit risk in both the CDS and ASW markets from country i to j.

In standard VAR models, shock identification is performed by imposing a Choleski decomposition in all countries. To reduce the number of identification restrictions in a VAR model, it is assumed that $E(u_{it}u'_{jt})$ is block diagonal, with blocks corresponding to each country. Canova and Ciccarelli (2009) state that block diagonality implies that within a country, variables are allowed to move instantaneously, but across entities, variables can only react with one lag. Assuming block diagonality is however a strong restriction and furthermore unrealistic in a cross-border contagion application using financial market data as also stated by Stock and Watson (1988).

The identification of shocks for PVAR models as defined in Equation (1) is more complicated, given that the PVAR model allows for static interdependencies, as u_{it} is correlated across entities i. Thus, cross-entity symmetry in shock identification cannot be assumed. We compute the impulse responses following Canova and Ciccarelli (2009) as the difference between two conditional forecasts: one where a particular variable is shocked and one where the disturbance is set to zero.

Canova and Ciccarelli (2013) defines the shock identification formally as: y^t to be the history of our time series y_t up to time stamp t , Θ^t are the estimated coefficients up to t. Let $W = (\Sigma_u, \sigma^2)$ be the sum of squared residuals and the residual variance; set

$\xi'_t = [u'_{1t}, u'_{2t}]$ where u'_{1t} are shocks to the endogenous variables and u'_{2t} shocks to the predetermined or exogenous variables. δ is a one standard deviation shock as we have standardised our time series.

Let $model^1_t$ contain y^t, Θ^t, W up to t and a one standard deviation shock $\xi^{\delta}_{j,t}$ to country j. The model forecast without shock is defined as $model^2_t$ which contains y^t, Θ^t, W and their variance up to t. The impulse responses at the future horizon τ to a δ impulse is the difference between these two conditional forecasts defined as:

$$IR^j_y(t, \tau) = E(y_{t+\tau}|model^1_t) - E(y_{t+\tau}|model^2_t). \tag{3}$$

6 Results

We investigate whether sovereign credit risk shocks emanating from the GIIPS countries spilled over contagiously onto the Visegrad group member countries. We rely on the identification strategy as described by the ECB with contagion effects detected when: (i) the transmission is different from regular adjustments observed in tranquil times (in line with Forbes (2012)); and (ii) the events constituting contagion are negative extremes. We restrict the identification of contagion to the case where both conditions (i) and (ii) must be fulfilled. Otherwise, when the effect is gradual and not a negative extreme, we refer to it as comovement. Thus, we will analyse the dynamics of sovereign credit risk during the pre-crisis and the sovereign debt crisis period in order to detect condition (i) and (ii) in order to analyse whether the shock transmission from the GIIPS to the Visegrad group was contagion or rather comovement. We do so by comparing the respective impulse response functions based on their magnitude, time of the shock decay and the response speed to a shock.

In our analysis of the pre-sovereign debt crisis and crisis period (Section 6.1) we apply one unit shocks to the Greek, Irish and Portuguese CDS markets, because these countries have been most severely hit by the economic crisis and had to be bailed out. By comparing the pre-crisis and crisis period we aim to detect whether the increase in sovereign credit risk in the Visegrad group was driven by contagion.

Further, we study the four bailout events of the Troika[12] between 2010 and 2011 in Section 6.2. We aim to find out whether the economic adjustment programmes (EAP) by the Troika had a stabilising effect on the sovereign risk of the Visegrad group. Following the framework proposed by Ters and Urban (2017b), we find that after a bailout event, systemic risk decreases in the GIIPS countries while the strongest stabilisation is visible with the second bailout for Greece. In our event study we will shock the country that receives the bailout and use five working day windows before and after each bailout.

[12] The troika is a decision group formed by the EC, the ECB, and the IMF.

We find clear indication that the shock transmission from the GIIPS to the Visegrad group countries is in almost all cases comovement rather than contagion. Whereas we find contagion effects amongst the GIIPS countries (see also Komarek et al. (2016), Ters and Urban (2017b) as well as Claeys and Vasicek (2014)). The magnitude of shock transmission from the GIIPS to the Visegrad group countries increases only slightly during the crisis and does therefore not represent a negative extreme. In other words, the shock transmission from the GIIPS countries to the Visegrad group countries remained moderate while condition (i) and (ii) for contagion detection are both not fulfilled. By employing an event study framework to analyse whether the bailouts for Ireland, Greece and Portugal had a stabilising effect on the Visegrad group, we find that in the post-event window of the bailout, sovereign credit risk moderately declined in all Visegrad group member countries. In general, we can see a stabilising effect due to the bailouts on sovereign credit risk in all countries and groups in our sample while the reduction of contagion dynamics is much higher amongst the GIIPS countries.

We were anticipating that Slovakia as the only euro area member amongst the Visegrad group member countries was expected to have the highest sensitivity to shocks emanating from the GIIPS countries due to its dependence on the euro. However, we find that shock transmission was independent from the euro area membership as all Visegrad group member countries behaved very homogeneously to shocks from the GIIPS countries. This points to the fact, that investors only focused on the countries that were most affected during the sovereign debt crisis. Furthermore, we also find that investors seem to price in how strong a potential future involvement of a country is in any possible future bailouts. This could explain why France and Germany were relatively more affected by shocks emanating from the GIIPS countries compared to the Visegrad group.

6.1 Credit risk contagion from GIIPS to Central European countries

The timelines of our estimated shock transmission and absorption can be found in Figures 7 to 9. We have focused on shocks in Greece (Figure 7), in Portugal (Figure 8) and in Ireland (Figure 9) as these countries have been most affected by the sovereign debt crisis and were furthermore bailed out by the Troika. We find for both subsample periods that the shock propagates immediately within 30 to 120 minutes (while 1 interval represents 30 minutes in the Figures 7 to 9). Therefore, our half-lives are dramatically shorter than in existing empirical studies and responses to shocks are typically not lagged as found, for example, in Koop and Korobilis (2016). They find that shock contagion spreads on average within one to two months in the case of shocks that do not decay over a timeline of 10 months. Our results are clearly in line with the generally accepted notion that financial markets react very fast to new information and that sovereign CDS adjust at an intraday speed during our whole sample period (Gyntelberg et al.; 2013).

The magnitude of the impulse responses is similar across all countries in each group and the respective subsamples (pre-crisis and crisis). However, the group consisting of the GIIPS countries exhibit much larger magnitudes of impulse responses compared with the rest of our sample countries. We find comovement effects in the Visegrad group and our control/riskfree countries as both conditions (i) and (ii) in our contagion identification are not fulfilled. When comparing the impulse responses following a shock in the GIIPS countries during the pre-crisis and the crisis period, we can see in Figures 7 to 9, that the impulse responses only moderately increase during the crisis period for all Visegrad group member countries. Additionally, the time until a shock decays and the response speed to a shock remained unchanged in between the two subsamples. Consequently, the Visegrad group member countries experienced comovement effects that are relatively smaller compared to the comovement effects that we find in our control/riskfree group.

The size of the impulse responses in the GIIPS countries hint to a contagion behaviour during the crisis. In other words, sovereign credit risk markets have clearly distinguished between countries with solid public finances and countries with large debt-to-GDP ratios. The slightly stronger impulse responses of the control/riskfree group relative to the Visegrad group to shocks from the GIIPS countries during the crisis cannot be explained by changes in fundamentals. Our results imply that investors might also price in how strongly the countries (mainly France and Germany) are involved in any possible future bailouts. These results might explain why France and Germany were relatively more affected by shocks emanating from the GIIPS countries compared to the Visegrad group. Interestingly, the Visegrad group member countries behave very homogeneously, despite the fact that Slovakia is an euro area member, like the GIIPS countries. A possible explanation for this finding is that financial markets might have honoured the solid public finances of Slovakia.

The impulse responses show a larger magnitude in all countries during the sovereign debt crisis period. The pre-sovereign debt crisis period represents a relative normal market environment for euro area sovereign entities without fast and furious shock contagion but rather with comovements across markets based on the ECB's identification strategy as defined by Constancio (2011). The sovereign credit risk transmission[13] during the crisis remains moderate for the Visegrad group countries while it slightly increases in the control/riskfree group and strongly increases in the GIIPS countries. This can be explained by the fact that GIIPS countries were subject to severe turmoil in their countries' government finances with subsequent ratings downgrades which led to strong contagion effects amongst each other. Hence, we find contagion effects during the sovereign debt

[13] Measured by the impulse response for each country in our sample following a one unit shock in the GIIPS countries.

crisis period for GIIPS countries as we find much stronger responses to shocks compared to the pre-crisis period (strong negative extremes during the crisis period).

Figure 7: Impulse responses - shock in Greece

This figure illustrates the impulse response for ΔCDS to a one-unit shock (increase) for the pre-crisis period (January 2008 to 19 October 2009) and the crisis period (20 October 2009 to December 2011). The figures are based on 5-year tenor data with a 30-minute sampling frequency. The y-axis represents the impulse response to a one-unit shock to Greek ΔCDS. The number of 30-minute time intervals is described by the x-axis. The shock decays after 2 hours (4 intervals), hence we plot only the time line up to 9 intervals (half trading day).

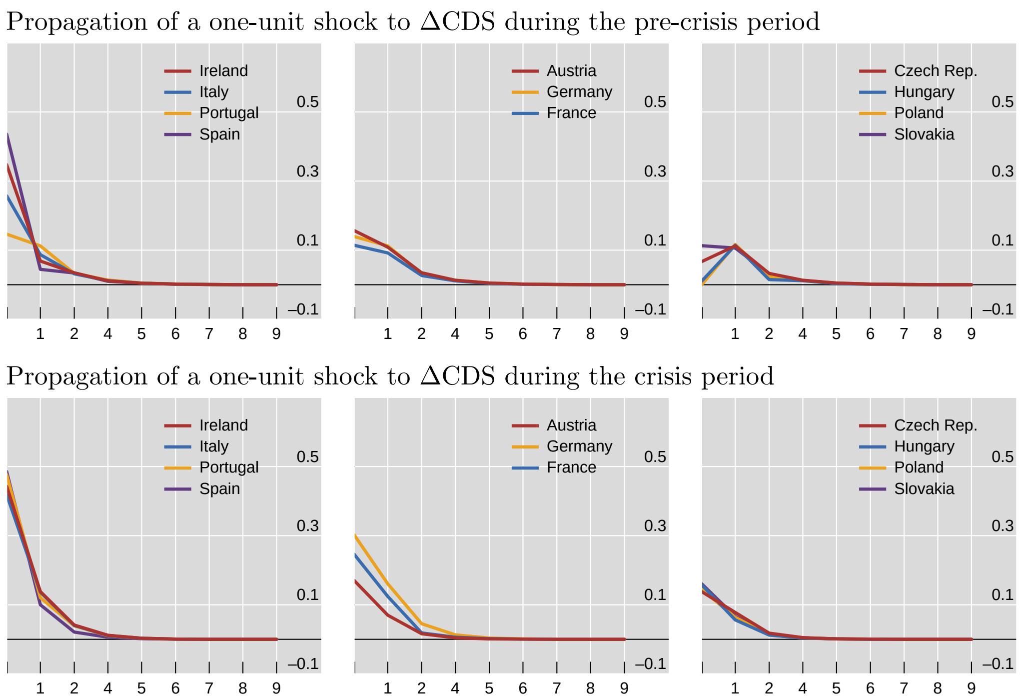

Figure 7 which displays the responses to a one unit shock in Greek ΔCDS shows no visible difference in the timeline of shock absorption between the pre-crisis period (upper panel) and the crisis period (lower panel), but the magnitude of the impulse responses increased during the crisis. The behaviour of all GIIPS countries is more homogeneous during the crisis as they all display very similar impulse responses. Germany's impulse response has the highest magnitude in the control/riskfree country group during the sovereign debt crisis period, followed by France and Austria.

Surprisingly, Hungary does not experience a negative extreme shock during the crisis period (measured by the impulse response to a one unit shock in GIIPS countries). We would have expected to see a more dramatic difference between the pre-crisis and crisis period due to the higher debt-to-GDP (see Figure 1) and less strong ratings (see Figure 3) of Hungary compared to the other Visegrad group countries.

The impulse responses in Figure 8 to a one unit shock in Portuguese ΔCDS and Figure 9 to a one unit shock in Irish ΔCDS show a very similar picture as the impulse responses in Figure 7 to a one unit shock in Greek ΔCDS. The difference between the pre-sovereign debt crisis period (upper panels) and the sovereign debt crisis period (lower panels) of the estimated impulse responses is most pronounced for the GIIPS countries. Impulse responses significantly increase in magnitude in all GIIPS countries following shocks in Greece, Ireland and Portugal while the impulse responses only moderately increase in the crisis period for our Visegrad and control/riskfree countries. Consequently, GIIPS are characterised by contagion dynamics while the Visegrad and the control/riskfree group experienced comovement effects (please refer to our identification strategy at the beginning of Section 6).

Figure 8: Impulse responses - shock in Portugal

For details see Figure 7.

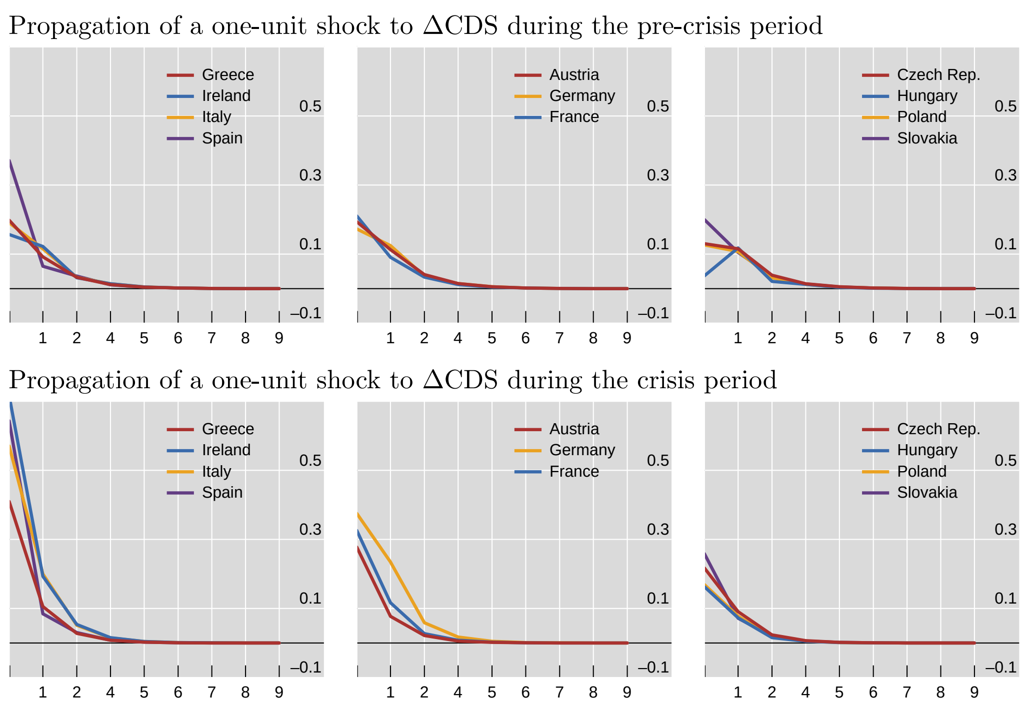

It is interesting to note that the responses of the CDS markets to shocks from Portugal, Ireland and Greece do not distinguish between the weaker public finances of Hungary or the euro membership of Slovakia. This result may lead to the conclusion that during the crisis credit risk markets focused strongly on the peripheral countries and not so much on the economic fundamentals of other less affected countries.

Figure 9: Impulse responses - shock in Ireland

For details see Figure 7.

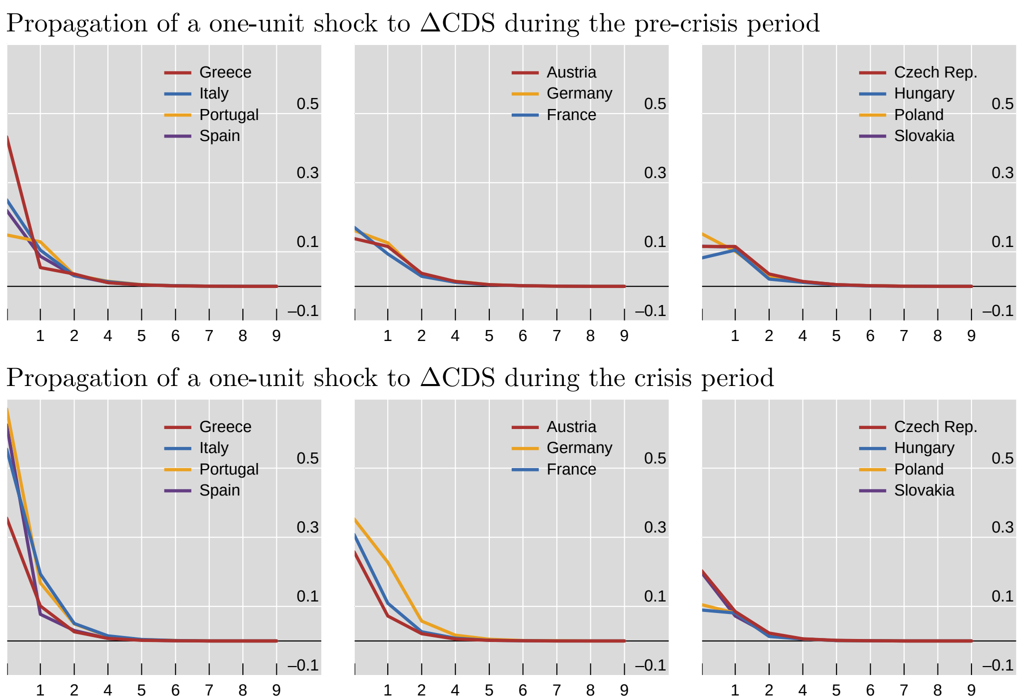

Koop and Korobilis (2016) find in their contagion analysis confidence bands for their impulse responses, that all lie between positive and negative reactions to a one-unit shock. The use of our intraday data results in a dramatic increase in the precision of the results. As a representative example we show the 95% confidence bands for a shock to Greek ΔCDS during the crisis period in Figure 10. Figure 10 needs to be compared to the lower panel in Figure 7. The key observation from Figure 10 is that the 95% error bands are very narrow, almost like a line plot. This high precision of our point estimates is based on the large amount of intraday data points in our subsample. This finding also confirms the aforementioned advantage of using intraday data as also found by Gyntelberg et al. (2013). The impulse response functions, presented in Figure 7 are statistically highly significant.

We find the same narrow error bands for all our estimated impulse response functions and hence only show our point estimates, for purely presentational purposes.

Figure 10: 95% error bands for the impulse responses - shock in Greece

This figure illustrates the 95% error bands for an impulse response for a propagation of a one-unit shock to ΔCDS during the crisis period. The mid values are presented in Figure 7. The confidence bands have been computed using a MCMC simulation with 500 draws.

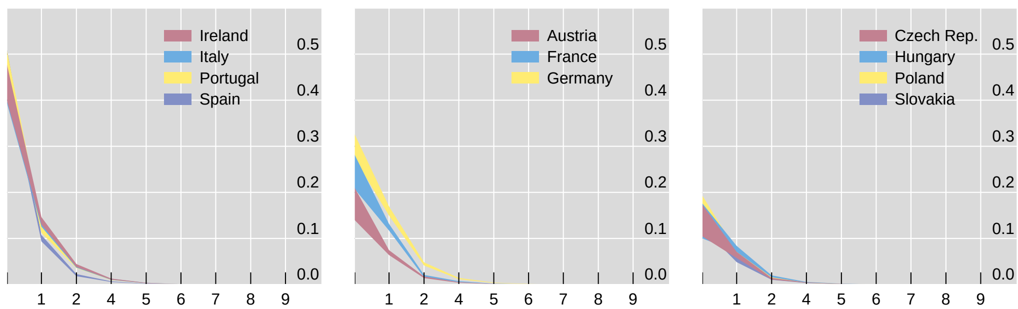

6.2 The effect of the economic adjustment programmes on contagion dynamics

In the previous Section 6.1 we have analysed the changes in the dynamics between the pre-sovereign debt crisis and the sovereign debt crisis period. Now we aim to investigate the dynamics around respective bailout events of the Troika, a decision group formed by the EC, the ECB, and the IMF. This will enable us to analyse whether the bailouts were able to improve financial stability by decreasing sovereign credit risk in the Visegrad group.

The bailouts or economic adjustment programmes (EAP) provided financial help to avoid a default of the most distressed countries and imposed economic and fiscal conditions on the country under bailout. During our crisis period we observe four EAP events: two bailouts for Greece, one for Portugal and one for Ireland (see Figure 11).

We define the event dates as the announcements of the Memorandum of Understanding (MoU) for the bailouts (EAPs) in Ireland, Portugal and the first bailout in Greece by the Troika. We do not find any significant effect in impulse responses when we use the governments' request for a bailout as the event date. This can be explained by the fact that bailout conditions have not yet been defined and made publicly available at that stage. The MoU for the first EAP for Greece was announced on 2 May 2010. The MoU for the EAP in Ireland was announced on 3 December 2010. Portugal's MoU for the EAP was announced on 17 May 2011. Finally, on 21 July 2011 the preliminary draft of the second bailout package for Greece was approved. For more details on the EAPs we refer to Ters and Urban (2017b).

All bailouts, except for the first EAP for Greece, were executed under the European Financial Stability Facility (EFSF) that have been created by the euro area Member States in June 2010 as a temporary crisis resolution mechanism. The financial assistance to Ireland, Portugal and Greece under the EFSF was financed through the issuance of bonds and other debt instruments on capital markets.

Figure 11: Economic adjustment programmes

The figure shows the 30-minute CDS spreads for Greece, Ireland and Portugal. The vertical lines correspond to the four bailout events during our sample period. We define the bailout events as the announcement dates of the Memorandum of Understanding for the economic adjustment programmes in Ireland, Portugal and the first bailout in Greece. The event date of the second bailout in Greece is defined as the announcement of the preliminary draft of the second Greek EAP.

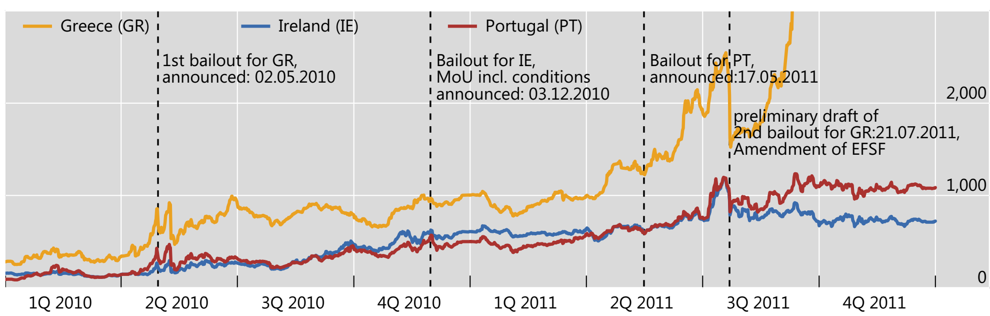

Figures 12 to 15 display the contagion dynamics prior (upper panel) and after (lower panel) the four bailout events. Contagion dynamics are estimated as the responses to a one-unit shock in ΔCDS arising in the country under the EAP. We take 5 working days, containing 18 data points per day, before and after the event and employ again the PVAR methodology as given in Equation (2). This event study makes the use of intraday data inevitable as it enables us to carry out our event study with a short pre- and post-event window and still keeping a sufficient number of observations.

The results in this event study reveal a similar behaviour as in the previous subsection: shock transmissions from the GIIPS to the Visegrad group countries were smaller (comovement) than shock transmissions among the GIIPS countries and to the control group countries. Generally, it can be said that as soon as the MoU containing the bailout conditions was published, we find that contagion/comovement was reduced in terms of magnitude. In other words, our results suggest that the EAPs were able to reduce overall sovereign credit risk and contagion interlinkages from the GIIPS to the Visegrad group countries and to the control group countries as well as amongst the GIIPS countries. However, we find a weakening efficiency of the EAP from the first to the third bailout in the ability to decrease sovereign credit risk spillovers (Figures 12 to 14).

Figure 12: 1st Greek bailout event (02.05.2010)

This figure presents the pre- and the post-event window of impulse response functions estimated with the PVAR in Equation (2) and 30-minute CDS data. The event window are set to five trading days, with 18 intraday observations for each trading day. We apply a one unit shock to the country that receives the bailout. The y-axis represents the response of ΔCDS to a shock arising from the country of the EAP event. The number of 30-minute time intervals is described by the x-axis. We display the point estimates of the impulse responses as all estimates are significant at the 95% CL.

Propagation of a one-unit shock to Greek ΔCDS before the event

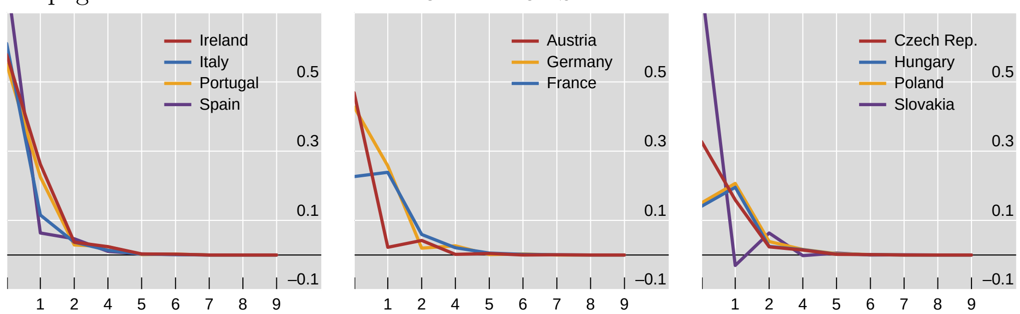

Propagation of a one-unit shock to Greek ΔCDS after the event

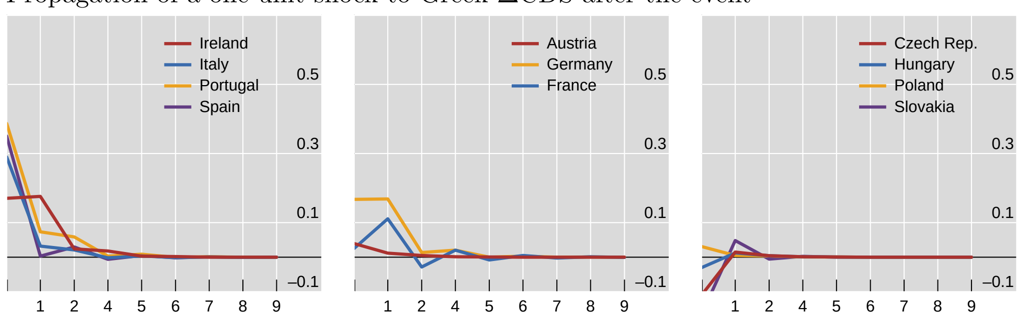

Figure 12 shows that after the first bailout in Greece we find that sovereign credit risk contagion did dramatically decrease in magnitude in all countries. Interestingly, Slovakia seems to be much stronger interconnected to the GIIPS countries prior to the first Greek bailout compared to the other Visegrad group countries. Most likely, this behaviour is linked to the euro membership of Slovakia. This effect can only be found for the very first bailout and is not visible for the subsequent bailouts. Prior to the first bailout, investors feared a severe downturn in the entire euro area, which includes Slovakia. The abandoning of the no-bailout clause of the Maastricht treaty and the increasing political intentions to avoid credit risk contagion amongst European countries, which later culminated in the speech of ECB president M. Draghi at the 26th July 2012, might have led to a change in investors expectations, with respect to sovereign credit risk of Slovakia. Consequently, the Visegrad group countries behaved more homogeneously in terms of the magnitude of the impulse responses during later bailouts (bailout for Ireland, Portugal and second bailout for Greece are found in Figures 13 to 15), which is also in line with the findings in Section 6.1.

For the second and third bailout (Irish bailout and Portuguese bailout, Figures 13 and 14) which have been transacted under the EFSF it is striking to see a reduction in the effectiveness of the bailouts compared to the first bailout in Greece. We find only a moderate reduction of the magnitude of the impulse response functions and hence sovereign credit risk within the GIIPS group (Figures 13 and 14) prior and post to the event date. Spillovers to the control/riskfree group are reduced and the interlinkage of the GIIPS countries to the Visegrad group countries decreased as well after the event.

Figure 13: Irish bailout event (03.12.2010)

For details see Figure 12.

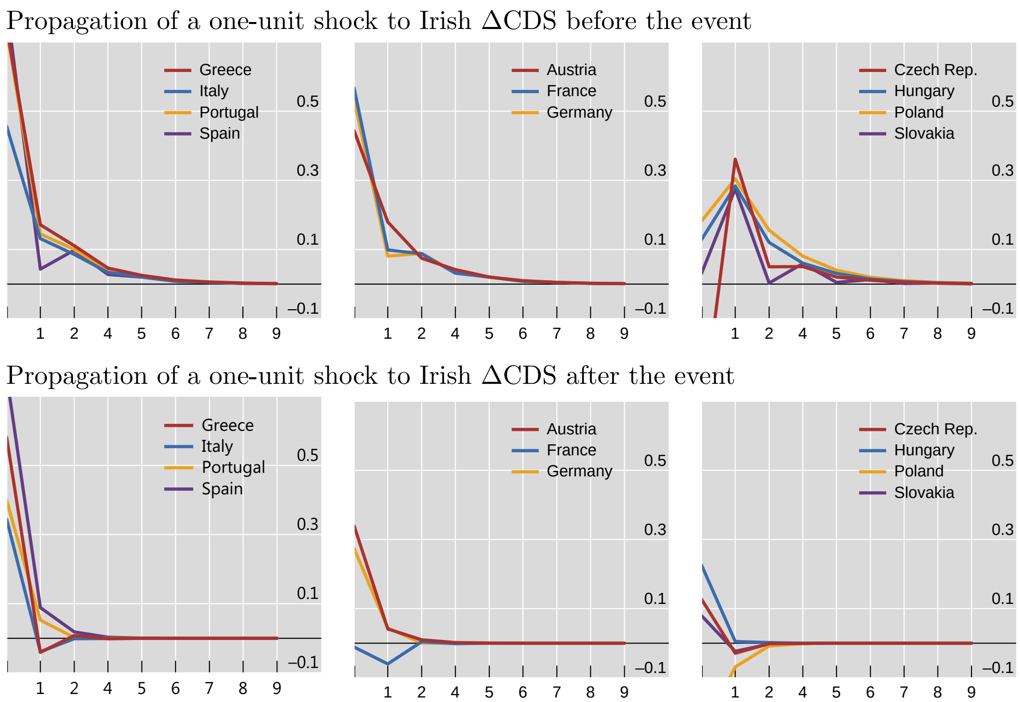

The Irish bailout in Figure 13 reveals a stabilising effect only for the control group and the Visegrad group. A similar pattern is found found the bailout in Portugal in Figure 14.

Figure 14: Portuguese bailout event (17.05.2011)

For details see Figure 12.

Propagation of a one-unit shock to Portuguese ΔCDS before the event

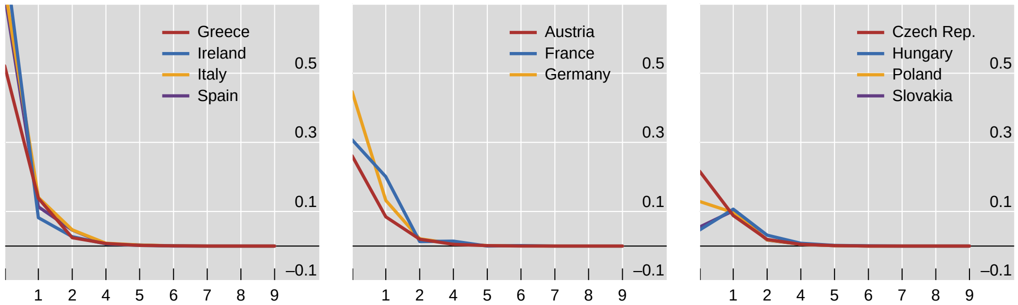

Propagation of a one-unit shock to Portuguese ΔCDS after the event

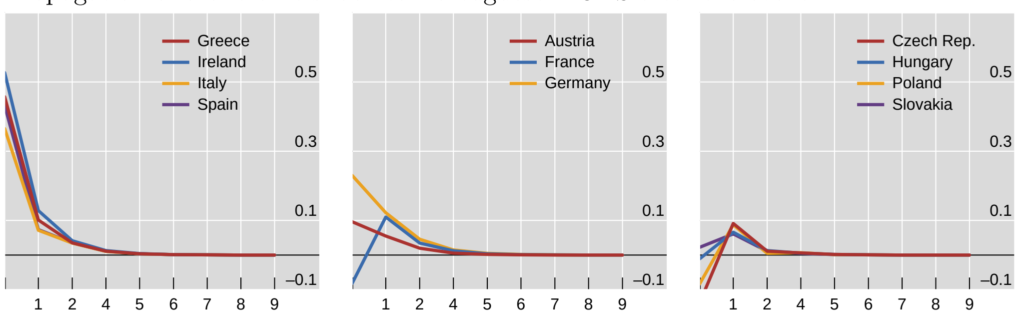

Policy makers must have recognised the declining effectiveness of the EAPs in terms of reducing systemic risk amongst GIIPS countries. Therefore, in parallel to the negotiations of the second Greek EAP it was decided to enlarge the European Financial Stability Facility (EFSF) from EUR 440 billion to EUR 780 billion to address the limitations of previous agreements. This resulted in a dramatic reduction of the magnitude of the impulse response functions amongst the GIIPS and also for Germany.[14], if one compares the pre- and post-event results around the second Greek bailout.

[14] Germany contributed most to the EAP. The shares of the participation in the EAP were determined according to the GDP per capita in each country. For example, Germany provided EUR 28 billion and France provided 16.8 billion to the first Greek bailout.

Figure 15: 2nd Greek bailout event (21.07.2011)

For details see Figure 12.

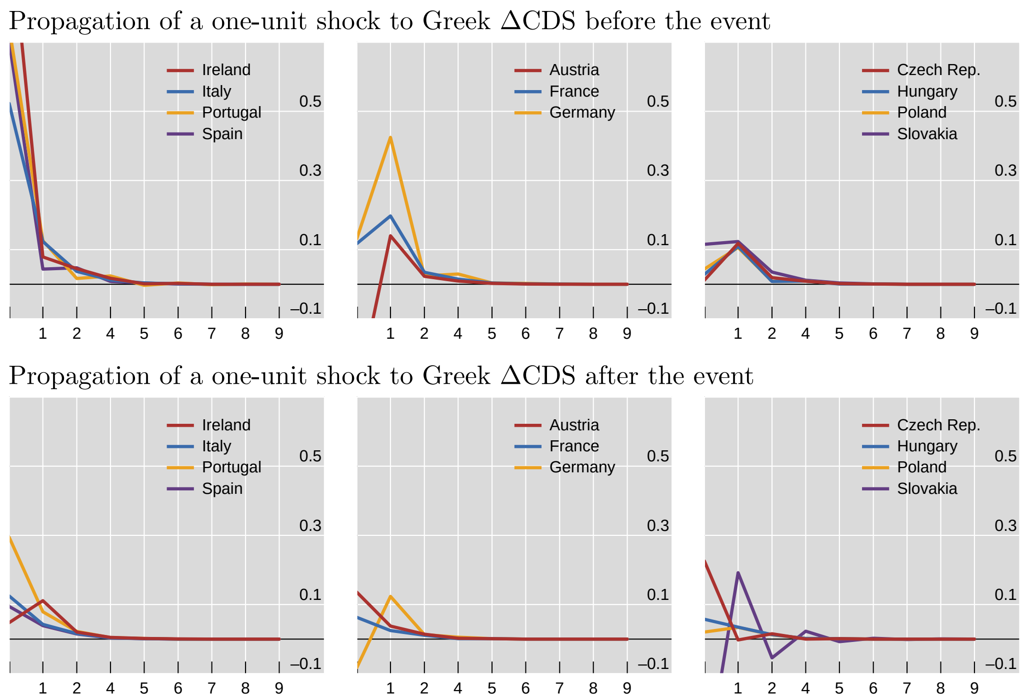

The second bailout in Greece as displayed in Figure 15 shows a clear stabilisation across all countries. This can be clearly explained by the dramatic amendment of the EFSF which confirms that expectations play an important role in sovereign credit risk markets as also discussed in Kliber (2014).

7 Conclusion

Our findings show evidence that the Visegrad group countries have been immune to strong sovereign credit risk contagion from the GIIPS countries as we find comovement and no contagion. In contrast, we find clear contagion effects amongst the GIIPS countries. This result points to the fact, that investors must have focused on the countries that were most affected by the euro area sovereign debt crisis. Additionally, the fact that even though Hungary has experienced a significant increase in its debt-to-GDP ratio in 2010 with respective ratings downgrades, its impulse responses to a shock in the GIIPS countries do not differ from the other Visegrad group member countries. Also interesting in this perspective is the case for Slovakia who is the only euro area member amongst the Visegrad group countries. Contrary to the expectation that Slovakia should have a higher sensitivity to shocks from GIIPS countries due to its dependence on the euro, we do not find differences in the impulse responses for Slovakia compared to the rest of the Visegrad group.

Our results of the pre-sovereign debt crisis period and the crisis period show differences in the impulse responses in terms of magnitude. Responses to sovereign credit risk shocks emanating from the GIIPS countries to the Visegrad group or control group countries during the pre-sovereign debt crisis and the sovereign debt crisis period can be categorised as comovement. We find contagion effects rather than comovement for shock transmission amongst the GIIPS countries as the transmission is fast and furious with strong negative effects. These results are in line with the macroeconomic fundamentals of the individual countries in terms of, amongst other, the debt-to-GDP ratio.

The use of intraday data substantially increases the precision of the results and we find average timelines of financial shock contagion of 30 to 120 minutes during the pre-crisis and crisis period. This is a clear indication of the efficiency of financial markets.

By estimating contagion dynamics before and after the announcement of bailouts (EAP), we find that the magnitude of the impulse response functions is strongly reduced when the country which received the bailout is also the origin of the shock. Our findings imply that the introduction of the EAP had a positive effect on financial market's risk perception for the individual country under bailout (EAP). The EAP had only a small effect on the Visegrad group countries as the credit risk market interlinkages to the GIIPS have been marginal even during the crisis. The only exception is Slovakia, where we found a rather strong interlinkage to the GIIPS countries prior to the first Greek bailout. However, the credit risk market interlinkage of Slovakia to the GIIPS countries waned after this first Greek bailout, which can be explained by the Slovakia's solid public finances. Furthermore, the first bailout was not yet financed under European Financial Stability Facility (EFSF), but directly by the EU member countries.

Further, we find that the individual bailouts from the first to the third EAP have had a diminishing effectiveness effect in terms of the ability to decrease systemic risk. As a result, policy makers have addressed these limitations and have enlarged the EFSF from EUR 440 billion to EUR 780 billion at the time when the 2nd Greek EAP was negotiated. This enlargement was acknowledged by the financial markets as we find strong reductions in interlinkages between credit risk markets after the announcement of the preliminary draft of the second bailout for Greece.

Conclusively, the theory of a higher interdependence through trade linkages as presented in Section 3 does not apply to the case of the Visegrad group during the euro area sovereign debt crisis. Even though the trade linkages of the Visegrad group to the GIIPS countries are around 10-20% of each Visegrad group member countries' GDP we only find that this leads to a comovement during the crisis period.

Part II
Arbitrage costs and the persistent non-zero CDS-bond basis

Part IIa: Limits to arbitrage: Estimating unknown arbitrage costs from a persistent non-zero basis[1]

1 Introduction

The theoretical no-arbitrage condition between two similar financial market instruments traded in the spot and derivative market (or between a convertible bond and its underlying stock), is a cornerstone for the empirical research on price discovery. The no-arbitrage condition requires that the pricing in the spot market must be equal to the derivative market. If not, any pricing discrepancy would present investors with an arbitrage opportunity which will disappear rapidly, as arbitrageurs will exploit any mispricing. This mispricing is measured by the so-called basis which we define in this paper as the difference between the spot and the futures price. There exists no universal definition of the basis, and different definitions are more common for different markets. In credit risk markets the basis is defined as derivative minus spot price (Gyntelberg et al.; 2013) or more concretely as CDS spread minus the spread on a par risky fixed-rate bond over the riskfree rate. Lien and Yang (2008) define the basis as the difference between spot and future prices in their application in commodity markets. Fama and French (1987) and McMillan (2005) on the other hand define the basis as future minus spot prices. Our econometric analysis and methodology is independent of which definition we chose.

The basis trading strategy, in which an arbitrageur believes that two similar financial market instruments are mispriced relative to each other, aim to take opposing long and short positions in these two securities in order to make a gain on the convergence of their values. In case of a positive basis, arbitrageurs will bet on a weakening basis (short basis position) and in case of a negative basis, arbitrageurs bet on a strengthening basis (long basis position). However, for the arbitrage condition to hold, markets must be perfect and frictionless. In practice, however, frictions and imperfections often make such arbitrage trades difficult and costly to varying degree. These imperfections include limited and time-varying liquidity across market segments, unavailability of instruments with identical maturity and payout structures, and the fact that some arbitrage trades require tying up large amounts of capital for extended periods of time.

A substantial part of the transaction costs of an arbitrage transaction is unknown when the arbitrage trade is initiated, making it risky. For index trades for example, Sutcliffe (2006) states that this risk can occur because the bid-ask spread and brokers' commission when unwinding the spot position at delivery vary with the value of the index

[1] This chapter is based on Urban, J. (2017) Limits to arbitrage: Estimating unknown arbitrage costs from a persistent non-zero basis. It is planned to extend this analysis in a joint collaboration with K. Ters and submit the resulting paper to the BIS working paper series.

basket and that there may be a transaction tax which varies in proportion to the index, eg an arbitrageur in the UK who buys the index basket at delivery must pay 0.5% stamp duty (Sutcliffe; 2006). Adams and van Deventer (1993) suggest, that in case of unknown arbitrage costs, traders should depart from the usual one-to-one ratio for the size of the spot and futures positions. In order to eliminate the transaction cost risk, arbitrageurs that are buying shares and selling futures, should buy $1/(1-p)$ index baskets for every one futures contract sold, where p is the proportion of the value of the index basket that must be paid in transaction costs at delivery. When selling shares and buying futures, the arbitrageurs should sell $1/(1+p)$ index baskets for every futures contract bought. Adams and van Deventer (1993) state, that this will remove the transaction cost risk from the arbitrage trade. However, they do not propose a methodology that can estimate the unknown transaction costs off such trades.

As a result of existing transaction costs on arbitrage trades, the difference between the prices in the spot and derivatives market for two similar financial market instruments, the so-called basis, is typically not zero. Moreover, the basis can become sizeable and persistent in times of market stress. Finally, when entering into a basis trade, the arbitrageur is exposed to the risk that the trade will move in the wrong direction. Thus, when markets are volatile, the basis trader is likely to ask for a higher compensation for the increased risk of the trade.

A persistent non-zero basis is therefore likely to reflect the unwillingness of arbitrageurs to try to exploit it, unless the pricing mismatch is greater than the cost of undertaking the arbitrage trade. Empirically, we would therefore expect to see such arbitrage forces intensifying as the magnitude of the basis exceeds some level that reflects the costs that traders face in the market. This suggests that the adjustment process towards the long-run equilibrium is nonlinear, in that it differs depending on the level of the basis. In order to capture such behaviour, we extend the linear vector error correction model (VECM) which has been the convention in existing studies (see for example Blanco et al. (2005) or Fontana and Scheicher (2016)) to a nonlinear setup using a threshold VECM (TVECM). There also exists research that specifically aims to estimate the effect of transaction costs on arbitrage such as Stevens (2015) in the market for crude oil. Stevens (2015) finds that transaction costs increase the persistence of the basis in the market for crude oil. He explains the non-zero basis by the absence of arbitrage. Forbes et al. (1999) investigate index futures arbitrage for the S&P 500 stock index and the nearest to delivery futures contract and find significant transaction costs that prevent arbitrage in the middle regime. Forbes et al. (1999) also find clear indication for arbitrage trading when the basis breaks out of the middle regime into the outer regimes. However, both Forbes et al. (1999) and Stevens (2015) employ an univariate structural change test to the cointegrating residual based on Tsay (1989). This approach would however only be valid when the cointegrating vector is known. Both of these before mentioned papers do not provide a solution for this problem. Forbes et al. (1999) also state in their conclusion, that the problem of an unknown cointegrating relationship in multiple threshold error correction models has not yet been resolved.

Hansen and Seo (2002) provide a methodology to estimate 2-regime threshold vector error correction models (TVECM) with an unknown cointegrating relationship. However, they do not provide a solution for the case beyond 2 regimes. An extension to a 3-regime or 2-threshold TVECM is important, because from an economic point of view transaction costs for a positive and a negative basis trade may exist. Furthermore, the model setup as proposed by Hansen and Seo (2002) is not adaptable to economic and financial market problems with a significant deviation from the theoretical parity relationship as they did not account for a persistent deviation in their long-term equilibrium condition.

The contribution of our paper is the development of an estimation procedure for threshold error correction models with three regimes (two thresholds) and an unknown cointegrating vector which is especially suitable to model arbitrage in markets with frictions. The estimation of an unknown cointegrating vector is particularly important for distorted parity relationships such as in financial markets and economic applications that exhibit a significant non-zero deviation from the theoretical parity relationship[2]. Hence, our proposed methodology also incorporates the possibility of a deviation from the parity relationship in the long-term equilibrium condition which is also a new contribution to the existing literature. We are able to quantify unknown transaction costs by employing our proposed estimation procedure for 3-regime error correction models with an unknown cointegration relationship. As our proposed model allows for nonlinear adjustment of prices in derivative and the spot markets towards the long-run equilibrium we can estimate the region where arbitrageurs step into the market as the trading opportunity is 'sufficiently profitable' for both, positive and negative basis trades.

The rest of the paper is structured as follows. Section 2 discusses the setup and estimation of our TVECM. It provides also a comprehensive simulation study to justify the validity of the proposed methodology. Section 3 provides some empirical applications, for two raw material markets (gold and platinum) and index trading (DAX and S&P 500), to illustrate the method and Section 4 concludes.

2 Threshold vector error correction model (TVECM)

The VECM concept implies that every small deviation from the long-run equilibrium leads instantaneously to an error correction mechanism. By extending the linear VECM

[2] The theoretical parity relationship states that the pricing in the spot market is identical to the pricing in the derivative or futures market for the same underlying

approach to a threshold vector error correction model (TVECM) we can model non-linearities in the adjustment dynamics. Threshold cointegration was introduced by Balke and Fomby (1997) as a feasible mean to combine regime switches and cointegration.

In the case of a persistent non-zero basis between the spot and the derivative market we expect to see that if the basis is lower than the cost of undertaking an arbitrage trade based on the observed basis[3], the arbitrageurs have no incentive to carry out the trade. Only when the deviation from the long-term equilibrium exceeds a critical threshold, such that the expected profit exceeds the costs, will economic agents act to move the basis back towards its long-term equilibrium. As a result, adjustments to the long-term equilibrium are likely to be regime-dependent, with no or a relatively weak adjustment mechanism in the regime where arbitrageurs have no incentive for trading as the overall transaction costs exceed the expected profit from the arbitrage trade.

The TVECM approach extends the VECM by allowing the behaviour of price quotes for spots S_t and derivatives D_t for a specific reference entity or underlying to depend on the state of the system. One can formulate a general TVECM with l regimes respectively with $l-1$ thresholds as follows[4]:

$$\Delta y_t = \sum_{j=1}^{l} \left[\mu^j + \lambda^j (S - \beta_1 D - \beta_0)_{t-1} + \Gamma^j(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1,\, \theta^{j-1},\, \theta^j) + \varepsilon_t, \qquad (1)$$

where $y_t = (S_t\ D_t)^\mathsf{T}$. All thresholds θ^j are ordered and

$$d_t(\beta_0,\, \beta_1,\, \theta^{j-1},\, \theta^j) = I(\theta^{j-1} \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^j)\,, \qquad (2)$$

with the error correction term $\mathrm{ec}_{t-1}(\beta_0, \beta_1) = (S - \beta_1 D - \beta_0)_{t-1}$. The indicator function $I(\theta^{j-1} \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^j)$ is 1 if the error correction term $\mathrm{ec}_{t-1}(\beta_0, \beta_1)$ is in the interval $[\theta^{j-1},\, \theta^j)$ and otherwise 0. Further, by definition the threshold θ^0 is $-\infty$ and θ^l is ∞ in Equation (1). We focus on the bi-variate case where $\Delta y_t = (\Delta S_t\ \Delta D_t)^\mathsf{T}$ is a 2-dimensional I(0) time series. The vector of price quotes $y_t = (S_t\ D_t)^\mathsf{T}$ is cointegrated with unknown β_0 and β_1. The error correction term $\mathrm{ec}_{t-1}(\beta_0, \beta_1)$ is therefore stationary by construction. $\varepsilon_t = (\varepsilon_t^S\ \varepsilon_t^D)^\mathsf{T}$ is a vector of i.i.d. shocks and $j \in \{1, 2, ..., l\}$ is the index denoting the l different regimes.

Equation (1) constitutes a vector autoregressive model in first-order differences with $\Gamma^j(L) = \sum_{k=1}^{m} \alpha^{j,k} L^k$ and L as lag operator, m as number of VAR lags, as well as a 2×1 constant μ^j and an additional error correction term $\lambda^j \mathrm{ec}_{t-1}(\beta_0, \beta_1)$. The speed of

[3] We define the basis as spot price minus derivative or futures price.
[4] For a derivation of the TVECM see for example Balke and Fomby (1997).

adjustment parameters $\lambda^j = (\lambda_1^j \quad \lambda_2^j)^\mathsf{T}$, the constant μ^j, and the lagged VAR terms are regime-dependent conditioned on the state of the error correction term $\mathrm{ec}_t(\beta_0, \beta_1)$.

The Schwarz (Bayesian) information criterion (SIC) should be used to determine the VAR order. Lütkepohl (2006) states that in large samples for multivariate models when $T \to \infty$ only the SIC criterion is seen to be strongly consistent for any K-variate system.

The error correction term represents the long-term equilibrium of the two time series which has to be an AR(1) process by construction (Johansen; 1988). The VAR-term represents the short-run dynamics coming from market imperfections (Baillie et al.; 2002).

Hansen and Seo (2002) define their model with a constant 2×1 vector μ^j and the error correction term as $\mathrm{ec}_{t-1}(\beta_1) = (S - \beta_1 D)_{t-1}$, ie they have set β_0 to zero. Contrary to the TVECM of Hansen and Seo (2002) we set the global constant $\mu^j = 0$ and keep the intercept β_0 in the error correction term. β_0 denotes the deviation from the long-term equilibrium, which is motivated by our no-arbitrage discussion in Section 1 and the local constant β_0 represents the persistent non-zero basis. We will however also briefly discuss the Hansen and Seo (2002) specification.

In markets with no frictions, the error correction term in Equation (1) is equal to the observed basis $(S - D)_t$, ie $\beta_0 = 0$ and $\beta_1 = 1$.

The speed of adjustment parameters characterize to what extent the price changes in $\Delta y_t = (\Delta S_t \; \Delta D_t)^\mathsf{T}$ react to deviations from the long-term equilibrium. In case price discovery takes place only in the derivatives market we would find a negative and statistically significant λ_1^j and a statistically insignificant λ_2^j, as the spot market would adjust to correct the pricing differentials from the long-term relationship. In other words, in this case the derivatives market would move ahead of the spot market as relevant information reaches investors. Conversely, if λ_1^j is not statistically significant but λ_2^j is positive and statistically significant, the price discovery process takes place in the spot market only - that is, the spot market moves ahead of the derivatives market. In cases where both λ's are significant, with λ_1^j negative and λ_2^j positive, price discovery takes place in both markets.

The costs for a basis trade prevent a complete adjustment towards a zero basis. As such, in markets with frictions there may be a neutral band between the derivative and the spot market in which the error correction term in Equation (1) may fluctuate without incentives for market participants to switch funds between the spot and derivatives market. Outside of that neutral band there might however be strong incentives for market participants to switch funds, which results in an adjustment towards the long-term equilibrium. We expect to find the speed of adjustment parameters to indicate that arbitrageurs engaging in $S_t - D_t$ basis trades as soon as the basis exceeds a threshold. In a market with a positive basis $(S_t > D_t)$, arbitrageurs bet on a declining basis and will short in the spot market and go long in the derivative market. In case of a negative basis ($S_t < D_t$), arbitrageurs bet on an increasing basis while carrying out the reverse trade.

According to arbitrage theory we would in general expect to find a 3-regime TVECM when the basis fluctuates between positive and negative figures with sizeable and persistent deviations from zero. The lower regime is defined as $\mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^1$, the middle regime as $\theta^1 \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^2$, and the upper regime is defined as $\theta^2 \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1)$. The middle regime is the neutral band where no arbitrage trading occurs. There may also be certain markets or time periods with a persistent positive basis. In that case we expect to find at most two regimes ($l = 2$) with only one threshold θ^1. The lower regime (neutral regime) is defined as $\mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^1$, and the upper regime as $\theta^1 \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1)$. The regimes are reversed in case of a negative basis market.

Therefore, we will discuss three classes of nested models: $\mathcal{H}_1$, which is a 1-regime VECM, with no statistical significant threshold, ie where markets are efficient enough to not allow the basis to deviate too far from zero. In markets described by $\mathcal{H}_1$ we have no transaction costs, hence arbitrageurs will step in as soon as the basis deviates from zero. $\mathcal{H}_2$ is a 2-regime TVECM and $\mathcal{H}_3$ is a 3-regime TVECM.

Figure 1: Model classes $\mathcal{H}_i$

The VECM model $\mathcal{H}_1$ represents markets or periods, where the basis does not deviate too strong from zero. The markets are perfect and frictionless for arbitrageurs to step in immediately to correct the pricing differential between the spot and derivative market. $\mathcal{H}_2$ and $\mathcal{H}_3$ are classical multiregime models, where outside the neutral regime arbitrageurs engage in basis trades. Note that a model belonging to class $\mathcal{H}_2$ can also have a negative threshold. In that case there would be arbitrage on basis strengthening.

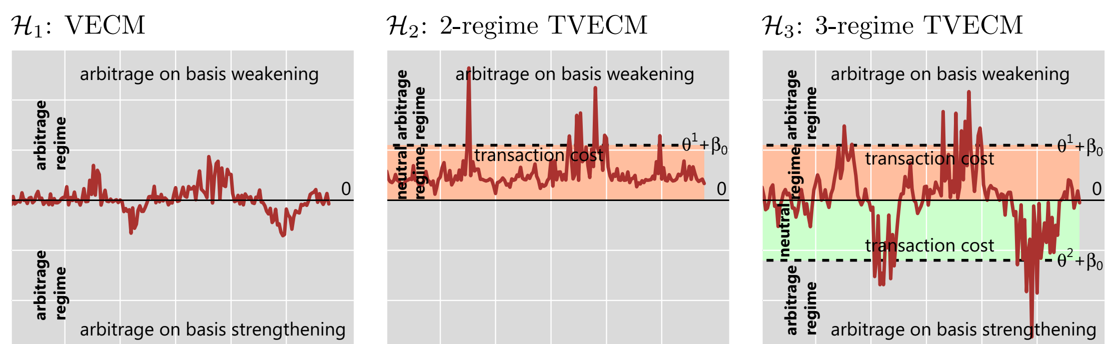

The threshold θ^j is computed relative to the estimated basis, which for our no-arbitrage model is shifted by β_0, because $\mathrm{ec}_{t-1}(\beta_0, \beta_1) = (S - \beta_1 D - \beta_0)_{t-1}$. Therefore, transaction costs, which are relative to the observed basis, are the sum $\theta^j + \beta_0$. In the case of $\theta^j + \beta_0 < 0$, we have transaction costs for an arbitrage trade on basis strengthening and in the case of $\theta^j + \beta_0 > 0$ we have costs for an arbitrage trade on basis weakening.

In the further discussion we will mainly focus on our no-arbitrage model, which is motivated by the no-arbitrage arguments discussed in Section 1, ie we set $\mu^j = 0$ and $\beta_0 \neq 0$. To be explicit we can write the VECM (class $\mathcal{H}_1$) as:

$$\begin{aligned}\Delta y_t &= \left[\lambda_1^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_1^1(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, -\infty, \infty) + \varepsilon_t \\ &= \lambda_1^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_1^1(L)\Delta y_t + \varepsilon_t\,, \end{aligned} \tag{3}$$

where the last line simply takes into account that for this special case the function d_t is identical to 1 along the entire time-axis. The subscript k in λ_k^i and Γ_k^i indicates explicitly to which model class the parameters belong, whereas the superscript i denotes the regime.

The model in Equation (1) may for a 2-regime TVECM (class $\mathcal{H}_2$) be written as:

$$\begin{aligned}\Delta y_t = &\left[\lambda_2^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^1(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, -\infty, \theta^1) \\ &+ \left[\lambda_2^2 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^2(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, \theta^1, \infty) + \varepsilon_t\,. \end{aligned} \tag{4}$$

In the case where the basis fluctuates between positive and negative values we may have to allow for two thresholds, respectively three regimes (class $\mathcal{H}_3$):

$$\begin{aligned}\Delta y_t = &\left[\lambda_3^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^1(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, -\infty, \theta^1) \\ &+ \left[\lambda_3^2 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^2(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, \theta^1, \theta^2) \\ &+ \left[\lambda_3^3 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^3(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, \theta^2, \infty) + \varepsilon_t. \end{aligned} \tag{5}$$

All parameters are allowed to switch between regimes except for β_0 and β_1. Following Hansen and Seo (2002) we estimate the model while imposing the following additional constraint for each regime:

$$\pi_0 \leq P(\theta^{j-1} \leq \mathrm{ec}_{t-1}(\beta_0, \beta_1) < \theta^j) \leq 1 - \pi_0\,, \tag{6}$$

where $\pi_0 > 0$ is a trimming parameter and P is the share of observations in each regime. This constraint allows us to identify a threshold effect only if the share of observations in each regime is large enough, ie is greater than π_0. If this condition is not met, the model reduces to a model with less regimes. Andrews (1993) argues that setting π_0 between 0.05 and 0.15 are typically good choices. We chose as a baseline setup $\pi_0 = 0.1$, but perform robustness checks also for $\pi_0 = 0.05$ and $\pi_0 = 0.15$.

2.1 Estimating the model

The most important statistical issue for threshold models is estimating the thresholds and testing their significance. Balke and Fomby (1997) suggest transforming the TVECM into a univariate arranged autoregression, while Tsay (1989) reformulates the problem into a univariate structural change test to the cointegrating residual. However, these approaches are valid only in the univariate case[5] when the cointegrating vector is known.

We follow the method proposed by Hansen and Seo (2002) who extend the literature by examining the case of an unknown cointegrating vector and we generalize their approach to a 3-regime case. They implement a maximum likelihood estimation of a bivariate TVECM with the assumption of i.i.d. Gaussian error terms. The likelihood function to be maximized for a l regime model takes the form:[6]

$$\mathcal{L}_n(\lambda_l^1,\, \ldots,\, \lambda_l^l,\, \Gamma_l^1,\, \ldots,\, \Gamma_l^l,\, \beta_0,\, \beta_1,\, \theta^1,\, \ldots,\, \theta^{l-1},\, \Sigma) = -\frac{n}{2}\ln|\Sigma| - \sum_{t=1}^{n}\frac{1}{2}\varepsilon_t'\Sigma^{-1}\varepsilon_t\,, \tag{7}$$

with $\Sigma = \mathrm{E}(\varepsilon_t\,\varepsilon_t')$ and n represents the sample size. ε_t and Σ are functions of λ_l^i, Γ_l^i, β_0, β_1 and θ^j, where $j = 1$ to $l-1$ and $i = 1$ to l.

Hansen and Seo (2002) suggest, that it is computationally convenient to hold β_0 and β_1 as well as θ^j fixed and compute the concentrated maximum likelihood estimations for λ_l^i, Γ_l^i and Σ. Due to the linearity of the model, this is simply an OLS regression. As shown in Hansen and Seo (2002) the concentrated likelihood function for β_0, β_1 and thresholds θ^j for a l regime model is:

$$\begin{aligned}\mathcal{L}_n(\beta_0,\, \beta_1,\, \theta^j) &= \mathcal{L}_n(\hat{\lambda}_l^i(\beta_0,\, \beta_1,\, \theta^j),\, \hat{\Gamma}_l^i(\beta_0,\, \beta_1,\, \theta^j),\, \hat{\Sigma}(\beta_0,\, \beta_1,\, \theta^j),\, \beta_0,\, \beta_1,\, \theta^j)\\ &= -\frac{n}{2}\ln|\hat{\Sigma}(\beta_0,\, \beta_1,\, \theta^j)| - \frac{n\cdot p}{2}\,,\end{aligned} \tag{8}$$

where again $j = 1$ to $l-1$, $i = 1$ to l in a l regime case and all variables with a hat are OLS estimators. We consider a bi-variate case with $p = 2$. The remaining task of finding the maximum likelihood estimation of β_0, β_1 and the thresholds is therefore to minimize $\ln|\hat{\Sigma}(\beta_0,\, \beta_1,\, \theta^j)|$, subject to the constraint in Equation (6).

Unfortunately, the function in Equation (8) is not smooth (see for example the left-hand panel of Figure 2), hence conventional hill-climbing algorithms cannot be used to find the extrema, therefore Hansen and Seo (2002) suggest a joint grid search. We present evidence of the good performance of the proposed grid search strategy in Section 2.3.

---

5 Engle and Granger (1987) representation of the VECM with a cointegrating vector of $\beta = [1 - 1]$.

6 We will focus on $l = 1$, 2 and 3 in our further analysis.

Two issues remain to be discussed with respect to the parameter estimation. The 2-regime model used by Hansen and Seo (2002) requires a two-dimensional grid search over (β_1, θ^1). Our model, which is motivated by no-arbitrage arguments, requires a search over a three dimensional grid in the 2-regime case and a four-dimensional grid search (β_0, β_1, θ^1, θ^2) in the 3-regime case. The 3-regime model requires the evaluation of Equation (8) at 100 Million grid points, if we evaluate each variable at 100 grid points. In order to keep the computation feasible, we suggest a sequential threshold search as it was proven to be consistent by Bai (1997) and Bai and Perron (1998).

The sequential search requires 2×100^3 grid point evaluations for the two thresholds. We will show that it is efficient to fix β_1 in the second threshold search to the value found in the first threshold search and hence reduce the search to a two dimensional space. This reduces the computational burden dramatically to 100^3 grid points for the first search and 100^2 grid points for the second search. We will discuss and justify this proposal in several comprehensive simulations in Section 2.3. In the same section we will also show that we cannot fix β_0 in the threshold search for θ^2, unlike β_1, because the β_0 estimate suffers from a large uncertainty.

The second remaining issue to be addressed is the setup of the "correct" search area for each parameter. The search region $[\theta_L, \theta_U]$ for the thresholds is straightforward as it must be identical to the interval $[\min(\mathrm{ec}_t(\beta_0, \beta_1)), \max(\mathrm{ec}_t(\beta_0, \beta_1))]$ given by the error correction term for β_0 and β_1. The region for the β_0 and β_1 parameters can be calibrated based on the estimates of the linear VECM model and the theoretical values $\beta_1 = 1$ and $\beta_0 = 0$, which would constitute the observed basis $(S - \beta_1 D - \beta_0)_t = (S - D)_t$. It is important to keep the search area for β_0 and η_1 large enough to include the minimum, but not too large to reduce precision of the grid search. The grid search for θ^2 will be reduced by the constraint $\theta^2 + \beta_0 > 0$ if the first search resulted in $\theta^1 + \beta_0 < 0$ and vice versa if $\theta^1 + \beta_0 > 0$. This is purely based on our no-arbitrage argument, where we expect to find at most two transaction costs, one for a positive basis trade and one for a negative basis trade.

Obviously the precision of the estimated parameters will depend on the distance between two neighbouring grid points. We will test in Section 2.3 various grid sizes, such as 10, 50 or 100 grid points for each search dimension as well as a dynamic grid setting. The dynamic grid setting has no fixed number of grid points, but a precision parameter fixes the distance between two neighbouring grid points. For our simulations we have chosen $\Delta\beta_0 = 0.5$, $\Delta\beta_1 = 0.01$ and $\Delta\theta^j = 0.5$. This choice is motivated by the applications discussed in Section 3 and the data generating process used in Section 2.3, which does not allow the time series to become larger in absolute terms than 500.[7] That means our basis can become as large as $\pm 1,000$ in extreme cases, leading to approximately 4,000

[7] The choice of precision parameters depends on the behaviour of the analysed time series.

grid points for β_0 and θ^j. β_1 is in a range of around 0 and 10, leading to maximal 1,000 grid points based on $\Delta\beta_1 = 0.01$.[8] The dynamic grid setting is computationally more expensive but we will show that as expected it yields the best results.

2.2 Testing for a threshold

The next step is to determine whether the estimated thresholds $\hat{\theta}^j$ are statistically significant. We know that the model class $\mathcal{H}_1$ is nested in $\mathcal{H}_2$ and $\mathcal{H}_2$ is nested in class $\mathcal{H}_3$. We start with the discussion of a 1-threshold model $\mathcal{H}_2$. In that case, under the null hypothesis, there is no threshold, so the model reduces to a conventional linear VECM where $\lambda_2^1 = \lambda_2^2 = \lambda_1^1$ and $\Gamma_2^1(L) = \Gamma_2^2(L) = \Gamma_1^1(L)$. The 1-threshold TVECM is detected under the alternative hypothesis with $\lambda_2^1 \neq \lambda_2^2$ under the constraint in Equation (6). As the models are linear, a regular LM test with an asymptotic $\chi^2(N)$-distribution can be calculated from a linear regression on the model in Equation (4). However, the LM-like test statistic can only be applied if β_0, β_1 and the threshold variable θ^1 are known a priori (Hansen and Seo; 2002). The point estimates of β_0 and β_1 under the null hypothesis are $\tilde{\beta}_0$ and $\tilde{\beta}_1$ from the linear model. However, there is no estimate of θ^1 under the null hypothesis. This implies that there is no distribution theory for the parameter estimates and no conventionally defined LM-like statistic.

We follow Hansen and Seo (2002) who derived the LM-like statistic for the one threshold case. We test for a linear VECM $\mathcal{H}_1$ under the null and a 1-threshold model $\mathcal{H}_2$ under the alternative hypothesis. The model $\mathcal{H}_1$ is defined in Equation (3) and here repeated:

$$\Delta y_t \quad = \quad \lambda_1^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_1^1(L)\Delta y_t + \varepsilon_t \tag{9}$$

and the model $\mathcal{H}_2$ is defined in Equation (4) and here also restated:

$$\begin{aligned}\Delta y_t \quad = \quad & \left[\lambda_2^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^1(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, -\infty, \theta^1) \\ & + \left[\lambda_2^2 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^2(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1, \theta^1, \infty) + \varepsilon_t\,.\end{aligned} \tag{10}$$

The subscript k in λ_k^i, Γ_k^i and in the below defined functionals denote explicitly to which model class $\mathcal{H}_k$ they belong.

We can derive the LM-like statistic as:

$$\mathrm{LM}(\beta_0,\, \beta_1,\, \theta^1) = \mathrm{vec}(\hat{A}_2^1 - \hat{A}_2^2)^{\mathsf{T}}(\hat{V}_2^1 + \hat{V}_2^2)^{-1}\mathrm{vec}(\hat{\mathrm{A}}_2^1 - \hat{\mathrm{A}}_2^2), \tag{11}$$

[8] The β_1 grid is gauged around the estimate of the VECM and the theoretical value of the observed basis, which is 1.

with the OLS estimator $\hat{A}_2^i$

$$\begin{aligned}\hat{A}_2^i(\beta_0,\,\beta_1,\,\theta^1) = \quad & \left(\sum_{t=1}^{n} X_{t-1}(\beta_0,\,\beta_1) X_{t-1}(\beta_0,\,\beta_1)^{\mathsf{T}} d_t(\beta_0,\,\beta_1,\,\theta^{i-1},\,\theta^i)\right)^{-1} \\ & \times \left(\sum_{t=1}^{n} X_{t-1}(\beta_0,\,\beta_1) \Delta y_t^{\mathsf{T}} d_t(\beta_0,\,\beta_1,\,\theta^{i-1},\,\theta^i)\right) \end{aligned} \tag{12}$$

where $i \in \{1,\,2\}$. n denotes the length of the time series. $\hat{A}_2^i$ are $(1 + p\,m) \times p$ matrices, with m denoting the number of VAR lags in our model and $p = 2$ in our bi-variate case. $\hat{V}_2^i$ is defined via the moment functionals M_2^i and Ω_2^i:

$$\hat{V}_2^i(\beta_0,\,\beta_1,\,\theta^1) = M_2^i(\beta_0,\,\beta_1,\,\theta^1)^{-1} \Omega_2^i(\beta_0,\,\beta_1,\,\theta^1) M_2^i(\beta_0,\,\beta_1,\,\theta^1)^{-1}\,. \tag{13}$$

The moment functionals are defined as:

$$M_2^i(\beta_0,\,\beta_1,\,\theta^1) \;=\; \mathbb{1}_p \otimes X^i(\beta_0,\,\beta_1,\,\theta^1)^{\mathsf{T}} X^i(\beta_0,\,\beta_1,\,\theta^1), \tag{14}$$

$$\Omega_2^i(\beta_0,\,\beta_1,\,\theta^1) \;=\; \xi^i(\beta_0,\,\beta_1,\,\theta^1)^{\mathsf{T}} \xi^i(\beta_0,\,\beta_1,\,\theta^1), \tag{15}$$

which are both $p \cdot (1 + p \cdot m) \times p \cdot (1 + p \cdot m)$ matrices.[9] X^i is a short form of the matrices of the stacked rows $X_{t-1}(\beta_0,\,\beta_1) \circ d_t(\beta_0,\,\beta_1,\,\theta^{i-1},\,\theta^i)$, with

$$X_{t-1}(\beta_0,\,\beta_1) = \begin{pmatrix} \mathrm{ec}_{t-1}(\beta_0, \beta_1) \\ \Delta y_{t-1} \\ \vdots \\ \Delta y_{t-m} \end{pmatrix}. \tag{16}$$

Hence, X^i is a $t \times (1 + p \cdot m)$ matrix of the following form:

$$X^i(\beta_0,\,\beta_1,\,\theta^1) = \begin{pmatrix} & & \vdots & \\ \mathrm{ec}_{t-1}(\beta_0, \beta_1) & \Delta y_{t-1}^{\mathsf{T}} & \dots & \Delta y_{t-m}^{\mathsf{T}} \\ & & \vdots & \end{pmatrix} \circ d_t(\beta_0,\,\beta_1,\,\theta^{i-1},\,\theta^i)\,, \tag{17}$$

[9] For the model defined in Hansen and Seo (2002) the moment functionals have the dimensions $p \cdot (2 + p \cdot m) \times p \cdot (2 + p \cdot m)$.

where $\circ$ denotes elementwise multiplication. X^i contains only non-zero entries if $\theta^{i-1} \leq \text{ec}_{t-1} < \theta^i$. For the here considered 2-regime or 1-threshold TVECM we have $i \in \{1, 2\}$ with $\theta^0 = -\infty$ and $\theta^2 = \infty$.

ξ^i is defined as $\tilde{\varepsilon}_t \otimes X_{t-1}(\beta_0, \beta_1) \circ d_t(\beta_0, \beta_1, \theta^{i-1}, \theta^i)$, with $\tilde{\varepsilon}_t$ is the OLS estimate of the residual vector from the linear model and $\otimes$ is the Kronecker-product.

We could evaluate Equation (11) at the point estimates of the null, which is the model $\mathcal{H}_1$, if the parameters β_0, β_1 and θ^1 would be known. However, there is no estimate of θ^1 for model class $\mathcal{H}_1$. Based on the union-intersection principle, Davies (1987) proposes:

$$\text{SupLM} = \sup_{\theta_L^1 \leq \theta^1 \leq \theta_U^1} \text{LM}(\tilde{\beta}_0, \tilde{\beta}_1, \theta^1) \tag{18}$$

with $\tilde{\beta}_i$ being the point estimates obtained under the null hypothesis (linear VECM).

According to the constraint in Equation (6) we set the search region $[\theta_L, \theta_U]$ such that θ_L is the π_0 percentile of the error correction term, and θ_U is the $(1 - \pi_0)$ percentile. This grid evaluation over $[\theta_L, \theta_U]$ is necessary to implement the maximisation defined in Equation (18) as the function $\text{LM}(\tilde{\beta}_0, \tilde{\beta}_1, \theta^1)$ is non-differentiable in θ^1 and hence conventional hill-climbing algorithm can not be used to find the extremum.

The value of θ^1 which maximizes Equation (18) is different from the MLE $\hat{\theta}^1$ in Section 2.1, as Equation (18) are LM tests that are based on parameter estimates obtained under the null hypothesis, ie $\mathcal{H}_1$. Also, the test statistic is calculated with HAC-consistent covariance matrix estimates which leads to differing estimates compared to the estimate in Section 2.1 (see also the discussion in Hansen and Seo (2002)).

For three sample simulations we present the maximum likelihood function and the supremum LM estimator in Figure 2, as well as the corresponding estimators of β_0 and β_1 used to compute the functions presented there. The left-hand side of Figure 2 shows the β_0 and β_1 estimates under the alternative hypothesis, ie in this case the TVECM (class $\mathcal{H}_2$), and therefore the minimum is very close to the theoretically expected threshold. The right-hand side shows the β_0 and β_1 estimates obtained from the linear VECM ($\mathcal{H}_1$) estimation, which are far off from the values used in the simulation ($\beta_0 = 10$, $\beta_1 = 1.1$), and therefore the maxima of the LM estimator are also far of the theoretical threshold, which is 3.

This issue was also discussed by Hansen and Seo (2002). The displayed difference in the estimated thresholds is generic and not special to threshold cointegration.

Just like in the one threshold case, the next step is to determine whether the estimated 2-threshold TVECM is statistically significant. Under the null hypothesis, there is one threshold θ^1, so the model reduces to a TVECM with two regimes ($\mathcal{H}_2$) described in Section 2.1 with $\lambda_3^1 \neq \lambda_3^2 = \lambda_3^3$. The two-threshold TVECM is detected under the alternative hypothesis $\mathcal{H}_3$ with $\lambda_3^1 \neq \lambda_3^2 \neq \lambda_3^3$. The constraint in Equation (6) is again applied. To be explicit, now the null hypothesis is our model $\mathcal{H}_2$, as defined also in Equation (4) and repeated:

$$\begin{aligned}\Delta y_t \quad = \quad & \left[\lambda_2^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^1(L)\Delta y_t\right] d_t(\beta_0,\ \beta_1, -\infty, \theta^1) \\ & + \left[\lambda_2^2 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_2^2(L)\Delta y_t\right] d_t(\beta_0,\ \beta_1, \theta^1, \infty) + \varepsilon_t \end{aligned} \tag{19}$$

Figure 2: Maximum likelihood estimator versus supremum LM estimator

The graphs show the maximum likelihood estimation and the supremum LM estimation for three sample simulations. Each graph presents also the corresponding estimators of β_0 and β_1 for each simulation. The left-hand side of Figure 2 shows the β_0 and β_1 values estimated under the alternative hypothesis, ie in this case the TVECM (class $\mathcal{H}_2$), and the minimum is very close to the theoretically expected threshold. The right-hand side shows β_0 and β_1 obtained from the VECM estimation, which are far off from the values used in the simulation ($\beta_0 = 10$, $\beta_1 = 1.1$), and so are the maxima of the supremum LM estimator far of the theoretical threshold, which was chosen to be 3.

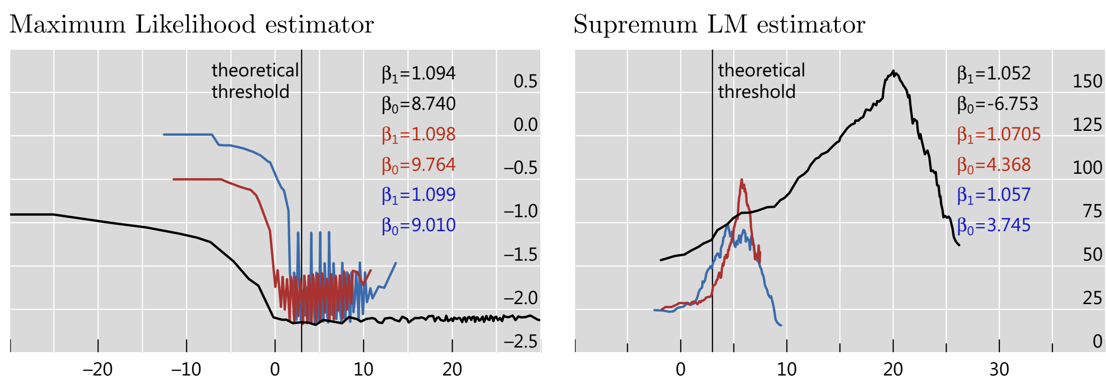

and the alternative hypothesis is our model $\mathcal{H}_3$ (as in Equation (5)):

$$\begin{aligned}\Delta y_t \quad = \quad & \left[\lambda_3^1 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^1(L)\Delta y_t\right] d_t(\beta_0,\ \beta_1, -\infty, \theta^1) \\ & + \left[\lambda_3^2 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^2(L)\Delta y_t\right] d_t(\beta_0,\ \beta_1, \theta^1, \theta^2) \\ & + \left[\lambda_3^3 \mathrm{ec}_{t-1}(\beta_0, \beta_1) + \Gamma_3^3(L)\Delta y_t\right] d_t(\beta_0,\ \beta_1, \theta^2, \infty) + \varepsilon_t. \end{aligned} \tag{20}$$

Following the same steps and arguments discussed above we find the following LM-like statistic:

$$\mathrm{LM}(\beta_0,\ \beta_1,\ \theta^1,\ \theta^2) = \mathrm{vec}(\hat{A}_3^2 - \hat{A}_3^3)^{\intercal}(\hat{V}_3^2 + \hat{V}_3^3)^{-1}\mathrm{vec}(\hat{A}_3^2 - \hat{A}_3^3). \tag{21}$$

Following again the proposal by Davies (1987) based on the union-intersection principle we get:

$$\text{SupLM} = \sup_{\theta_L^2 \leq \theta^2 \leq \theta_U^2} \text{LM}(\tilde{\beta}_0,\, \tilde{\beta}_1,\, \tilde{\theta}^1,\, \theta^2) \tag{22}$$

with $\tilde{\beta}_i$ and $\tilde{\theta}^1$ the point estimates obtained under the null (1-threshold TVECM). There is no point estimate of θ^2 under the null hypothesis. We perform again a grid search with the search region for θ^2 subject to the constraint in Equation (6). Furthermore, based on our no-arbitrage assumption and the assumption, that we have at most one positive and one negative level of transaction costs, we further impose for the grid search of θ^2 and β_0 the constraints that $\theta^2 + \beta_0 > 0$ if the first search resulted in transaction costs $\theta^1 + \beta_0 < 0$ and $\theta^2 + \beta_0 < 0$ if the first search resulted in transaction costs $\theta^1 + \beta_0 > 0$.

As there is no formal distribution theory in the case under discussion we follow the proposition by Hansen and Seo (2002) and perform two different bootstrap methodologies in order to estimate the asymptotic distribution for our model specification in Equation (1).

2.2.1 Fixed regressor bootstrap

We implement a non-parametric bootstrap of the residuals, called the "fixed regressor bootstrap", which resamples (Monte-Carlo) the residuals from the estimated linear VECM or 1-threshold TVECM, in the case of the threshold search for θ^1 or θ^2, respectively.

We follow the discussion of Hansen and Seo (2002), Hansen (2000) and Hansen (1996). We take estimates under the null hypothesis of β_0 and β_1, denoted as $\tilde{\beta}_0$ and $\tilde{\beta}_1$, respectively, and define $\tilde{\text{ec}}_{t-1} \equiv \text{ec}_{t-1}(\tilde{\beta}_i)$ and $\tilde{X}_{t-1} \equiv X_{t-1}(\tilde{\beta}_i)$, whereby ec_{t-1} denotes the error correction term (see definition below Equation (2)) and X_{t-1} is defined in Equation (16). Further, $\tilde{\varepsilon}_t$ are the residuals of the null. The name "fixed regressor bootstrap" conveys the message that $\tilde{\beta}_i$, $\tilde{\varepsilon}_t$, $\tilde{\text{ec}}_{t-1}$ and $\tilde{X}_{t-1}$ are kept fixed at their sample values.

Next, we compute a large number of times (eg 1,000) $y_{bt} = \tilde{\varepsilon}_t e_{bt}$, whereby e_{bt} is i.i.d. N(0,1) and in each draw e_{bt} is independently chosen. For each draw (identified by the index b), we perform an LM test. $\tilde{\varepsilon}_{bt}$ is computed by regressing y_{bt} on $\tilde{X}_{t-1}$. $\hat{A}_2^j(\tilde{\beta}_i,\, \theta^1)_b$ (see Equation (12) for the 1-threshold case)[10] and $\hat{\varepsilon}_{bt}(\tilde{\beta}_i,\, \theta^1)$ are computed by regressing y_{bt} on $\tilde{X}_{t-1} d_{1t}(\tilde{\beta}_i, \theta^1)$ and $\tilde{X}_{t-1} d_{2t}(\tilde{\beta}_i,\, \theta^1)$, whereby $\tilde{\beta}_i$ is kept fixed. Further, for each draw

[10] For the 2-threshold TVECM we need to compute $\hat{A}_3^j(\tilde{\beta}_i,\, \tilde{\theta}^1, \theta^2)_b$, whereby the values of the null $\tilde{\beta}_i$ and $\tilde{\theta}^1$ are kept fixed.

b we compute $\hat{V}_2^j(\tilde{\beta}_i, \theta^1)_b$ (see Equation (13) for the 1-threshold TVECM)[11], whereby $\tilde{\beta}_i$ is kept fixed again. Similar to Equations (11) and (18)[12] we compute for each draw b:

$$\text{SupLM}^* = \sup_{\theta_L^1 \leq \theta^1 < \theta_U^1} \text{vec}(\hat{A}_{2b}^1 - \hat{A}_{2b}^2)^\intercal (\hat{V}_{2b}^1 + \hat{V}_{2b}^2)^{-1} \text{vec}(\hat{\mathrm{A}}_{2\mathrm{b}}^1 - \hat{\mathrm{A}}_{2\mathrm{b}}^2), \tag{23}$$

where $\hat{A}_{2b}^i$ and $\hat{V}_{2b}^i$ are functions of the fixed $\tilde{\beta}_i$ and θ^1.

Hansen (1996) has shown that SupLM* is a valid first-order approximation to the asymptotic null distribution of SupLM. Despite having the computational cost of a bootstrap, it only approximates the asymptotic distribution

The p-value is than calculated as the percentage of SupLM* values which exceed the actual SupLM value of the original time series.

2.2.2 Residual bootstrap

This method is fully parametric with respect to the data generating process, that means for the one threshold case we use the complete specification for the null as given by $\mathcal{H}_1$ (single regime VECM versus 1-threshold TVECM as alternative) and for the 2-threshold TVECM we use the complete specification for the null as given by $\mathcal{H}_2$ (1-threshold TVECM versus 2-threshold TVECM as alternative). We further assume ε_t to be i.i.d. from an unknown distribution and fixed initial conditions. To be specific, random draws are made from $\tilde{\varepsilon}_t$, which are the residuals under the null. Using the given initial conditions from the data, and the parameters estimated under the null (in our case $\tilde{\lambda}^i$, $\tilde{\beta}_0$, $\tilde{\beta}_1$ and $\tilde{\Gamma}^i$) we recursively generate the bivariate vector series x_{bt} for the given model (in our case either $\mathcal{H}_1$ or $\mathcal{H}_2$). For each draw the SupLM* value is computed and than again the percentage of the SupLM* values which exceed the actual SupLM value (computed from the original time series) gives the p-value.

Hansen and Seo (2002) conjecture that this bootstrap method gives better finite sample performance at the computational cost of being fully parametric with respect to the data generating process.

2.3 Simulation

There is practically no empirical research using the proposed methodology of Hansen and Seo (2002) for TVECMs with one or two thresholds. Therefore, we perform numerous simulations in order to test the power of the proposed LM-like tests using different data generating processes for the VECM and TVECM model specifications in Equation (1), using $\text{ec}_{t-1}(\beta_0, \beta_1) = (S - \beta_1 D - \beta_0)_{t-1}$ and the specification according to Hansen and Seo

[11] For the 2-threshold TVECM we need to compute $\hat{V}_3^j(\tilde{\beta}_i, \tilde{\theta}^1, \theta^2)_n$, whereby the values of the null $\tilde{\beta}_i$ and $\tilde{\theta}^1$ are kept fixed.

[12] In the 2-threshold TVECM we need to compute Equation (22)

(2002). In particular, we aim to test the power of the two different bootstrap methodologies, the fixed regressor bootstrap and the residual bootstrap as discussed in Section 2.2, in terms of α- and β-errors (see Figure 3).

Figure 3: Null versus alternative hypothesis in the sequential search

The fixed regressor bootstrap and the residual bootstrap test assume that there is no threshold in the threshold search for θ^1, ie for the first threshold search the null hypothesis is a VECM ($\mathcal{H}_1$). For the second threshold search (θ^2), the null is a 1-threshold TVECM ($\mathcal{H}_2$), with the 2-threshold TVECM ($\mathcal{H}_3$) as the alternative hypothesis.

Threshold search for θ^1 Threshold search for θ^2

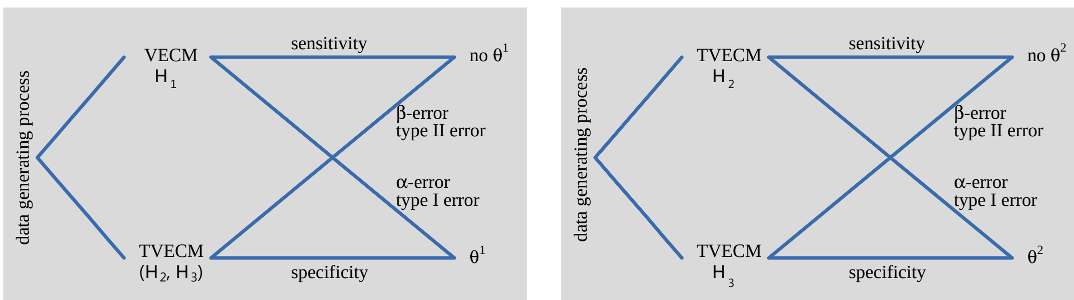

In this subsection we use the following notation: β_i and θ^i are the parameters fixed in the data generating process, while $\hat{\beta}_i$ and $\hat{\theta}^i$ are the estimates of our simulation. For each simulation we generate 1,000 estimators $\hat{\beta}_i$ and $\hat{\theta}^i$ and compute the mean as well as the standard deviation. The mean is compared to β_i and θ^i and the standard deviation is used as a measure of the quality of the estimation.

The aim of the various simulations is to understand how precise we can estimate β_i and θ^i, as well as to understand the power of the fixed regressor and the residual bootstrap.

2.3.1 The Hansen Seo model

We start our simulations with the specification used by Hansen and Seo (2002). Following Equation (1), our data generating process is defined as (see Equation (1) with $\beta_0 = 0$):

$$\Delta y_t = \sum_{j=1}^{l} \left[\mu^j + \lambda^j (1, -\beta_1)^\mathsf{T} y_{t-1} + \Gamma^j(L)\Delta y_t\right] d_t(\beta_1,\, \theta^{j-1},\, \theta^j) + \varepsilon_t. \tag{24}$$

We focus on a model with one and two regimes, ie $l = 1$ and $l = 2$, respectively. Hence, we generate 1,000 time series with no threshold and 1,000 time series with one threshold θ^1. All parameters in Equation (24) are generated with a random number generator, except for β_1 and the threshold θ^1 (in case of a threshold model). We have fixed the VAR lag in the data generating process to one. Relaxing this restriction to more VAR lags is straight forward and yields the same results, however with lower precision due to the larger number of parameters to be estimated.

The generated time series has a length of 1,000 periods and we use an equidistant search grid for β_1 and for the threshold θ^1 of 10, 50 and 100 grid points. Without loss of generality we present results for $\beta_1 = 1.2$. We further impose additional restrictions regarding the generated time series: firstly, we expect that the data generating process in Equation (24) produces a time series that is I(1) and the basis $(1, -\beta_1)^\mathsf{T} y_t$ is I(0) at 90% confidence level using the augmented Dickey-Fuller test. We do also not allow for the time series to become large in absolute terms, ie we set the maximal allowed absolute value of each generated data point to 500. For cross checking purposes we also test a simulation where we do not enforce these restrictions. We label this simulation in the following tables as "unrestricted" as opposed to the "restricted" cases. We present the results for the Hansen and Seo model specification in Table 1.

Table 1: VECM - model class $\mathcal{H}_1$, Hansen and Seo specification

The VECM model belongs to the model class $\mathcal{H}_1$, ie there exists only one regime. Hence, we do not have a threshold θ^1. The table shows the means and the standard deviations (figures in brackets) of 1,000 estimates $\hat{\beta}_1$ and $\hat{\theta}^1$ for different grid settings. We also report how often the alternative hypothesis (one threshold) is incorrectly accepted at different confidence levels by the fixed regressor and the residual bootstrap.

	grid	$\hat{\beta}_1$ (stdev)	$\hat{\theta}^1$ (stdev)	fixed regressor α-error, CL			residual α-error, CL		
theor. value		1.20	n/a	90%	95%	99%	90%	95%	99%
restricted	10	1.20	1.27	0.10	0.06	0.01	0.09	0.05	0.01
		(0.01)	(4.20)						
restricted	50	1.20	1.29	0.14	0.08	0.03	0.12	0.07	0.02
		(0.01)	(3.96)						
restricted	100	1.20	1.36	0.12	0.06	0.01	0.10	0.05	0.01
		(0.01)	(4.11)						
unrestricted	100	1.20	1.88	0.14	0.07	0.01	0.12	0.06	0.01
		(0.01)	(120.28)						

The β_1 estimate is very precise and independent of the grid size. The null hypothesis is rejected at different confidence levels (ie a threshold θ^1 is accepted) similarly by the fixed regressor and the residual bootstrap. We find most cases in Table 1 a 1% α-error at 99% confidence level. The estimate $\hat{\theta}^1$ is as expected scattered around zero, with the biggest standard deviation for the unrestricted case.

The independence of the quality of the β_1 estimate from the grid is not surprising, because the model under discussion is a VECM. The grid is constructed around the VECM estimate and the theoretical expected value $\beta_1 = 1$ (observed basis is $S_t - D_t$), which is by construction already close to the "true" VECM estimate. In other words, a grid search is not necessary for a model of class $\mathcal{H}_1$.

Further, we generate 1,000 1-threshold TVECM (2-regime TVECM) time series of length 1,000, with one threshold $\theta^1 = 30$. The results of the 1,000 MC simulations are presented in Table 2.

Table 2: 1-threshold TVECM - model class $\mathcal{H}_2$, Hansen and Seo specification

The TVECM model belongs to the model class $\mathcal{H}_2$, ie there exist two regimes and one threshold θ^1. The table shows the means and the standard deviations (figures in brackets) of the 1,000 estimates $\hat{\beta}_1$ and $\hat{\theta}^1$, as well as how often the null hypothesis (VECM) could not be rejected at different confidence levels by the fixed regressor and the residual bootstrap.

	grid	$\hat{\beta}_1$ (stdev)	$\hat{\theta}^1$ (stdev)	fixed regressor β-error, CL			residual β-error, CL		
theor. value		1.20	30	90%	95%	99%	90%	95%	99%
restricted	10	1.20 (0.03)	28.17 (3.23)	0.000	0.000	0.000	0.000	0.000	0.000
restricted	50	1.20 (0.01)	29.81 (1.18)	0.000	0.000	0.000	0.000	0.000	0.000
restricted	100	1.20 (0.01)	29.96 (0.59)	0.000	0.000	0.000	0.000	0.000	0.000
unrestricted	100	1.20 (0.01)	30.02 (7.12)	0.000	0.000	0.000	0.000	0.002	0.002

The performance of the estimation procedure and the test of the significance of the threshold is magnificent. We get practically no false negative results. The estimation of β_1 is very precise with the mean of $\hat{\beta}_1$ identical to the theoretical value 1.20 and a standard deviation below 0.03. The precision of the estimation of θ^1 gets lower in the unrestricted case or for smaller grid sizes. The larger standard deviation of $\hat{\theta}^1$ for the unrestricted case is an expected result, as the grid points have to be dispersed over a larger range, resulting in a lower precision. Unlike in Table 1 we see now a weak dependence of the quality of the

β_1 estimate on the grid choice. The model discussed in Table 2 is a TVECM, hence the VECM estimate of β_1, which is used to gauge the β_1 grid, may potentially be incorrect. This leads to the dependence on the chosen grid setup.

It is important to note that for most economic applications (see Section 3), the restricted case is important, because the time series and the basis is moving usually in a relatively narrow range.

Figure 4 shows the negative log-likelihood function for three out of our 1,000 simulations, where we have chosen a TVECM as a data generating process. From the graphs it is obvious that the function is not smooth and hence the proposed grid search method is necessary (conventional hill-climbing algorithms are inappropriate (Hansen and Seo; 2002)).

Figure 4: Negative log-likelihood function for the Hansen and Seo model specification

The figures show the negative log-likelihood function for three of our 1,000 simulations for a 1-threshold TVECM. It is obvious that the functions are not smooth. The graphs also show why our estimator $\hat{\beta}_1$ has a very high precision and $\hat{\theta}^1$ has a slightly lower precision.

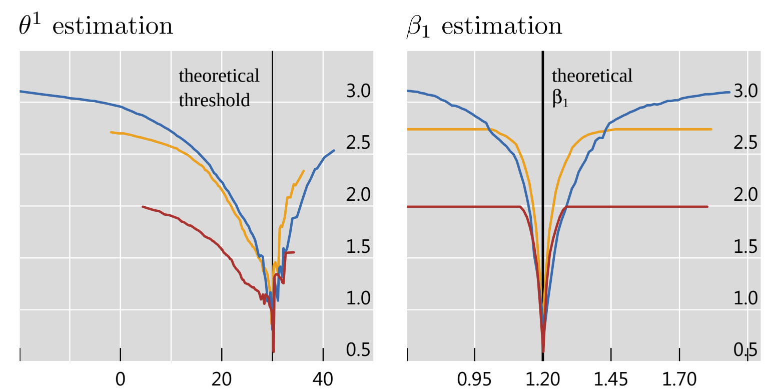

2.3.2 No-arbitrage model class $\mathcal{H}_1$

The Hansen and Seo model specification served as a benchmark or starting point which we want to compare now to the model specification (see Equation (1) with $\mu^j = 0$):

$$\Delta y_t = \sum_{j=1}^{l} \left[\lambda^j (S - \beta_1 D - \beta_0)_{t-1} + \Gamma^j(L)\Delta y_t\right] d_t(\beta_0,\, \beta_1,\, \theta^{j-1},\, \theta^j) + \varepsilon_t, \tag{25}$$

which is motivated by the no-arbitrage argument discussed in Section 1. In our further analysis as well as in Section 3 we will therefore focus entirely on the (T)VECM as specified in Equation (25). Hence, we are going to simulate more variation of grid settings and time series length, as we have done for the Hansen and Seo model. The imperative for a more in-depth simulation study is motivated by the facts that, firstly we have to search over a three dimensional grid for β_0, β_1 and θ^1 and secondly, our aim is to extend the model to include a second threshold θ^2.

Initially we start with a pure VECM model $\mathcal{H}_1$ ($l = 1$ in Equation (25)) , where we choose as an example $\beta_1 = 1.10$ and $\beta_0 = -10$. Again, we compute a time series with a sample size of 1,000 periods and use a grid size of 10, 50 and 100 as well as a dynamic grid setting. The dynamic grid setting is determined by a minimum distance, between two individual grid points. The dynamic setting is of course potentially very expensive, as it may lead to a large number of grid points. In our various simulations, we have chosen 0.01 as the distance between two grid points for the β_1 grid and 0.5 for the β_0 and θ^1 grids. In the restricted case, where we have the spot and derivative time series confined in absolute terms to 500, the basis can vary between -1,000 and 1,000 (assuming β_1 around 1). In such a case we can expect a grid size of maximal 2,000/0.5=4,000 for the β_0 grid as well as for the θ^1 grid. The grid size for the β_1 does usually not reach such extreme values. We have chosen for economic reasons the value of β_1 in our data generating process to be close to 1 and construct the grid around the theoretical value 1 (observed basis) and the value found in the initial VECM estimation (see discussion at the end of Section 2.1). Assume that the VECM estimation is 4, than the grid is gauged from around 0.8 and 4.2, to include the VECM estimate and the theoretical value of 1. This leads to grid size of (4.2-0.8)/0.01=340. The advantage of the dynamic grid setting is that the precision of the estimation process is predefined.

The results of our simulation using a VECM as a data generating process are shown in Table 3.

The estimations of β_1 and β_0 are very precise and practically independent of the grid size. Both parameters are estimated with similar precision in the restricted and unrestricted case. The reason for the independence of the results on the grid setup is straightforward, because we have used the standard VECM estimates and have gauged the grid around these values. Our data generating process is a VECM, therefore we have generated already the "correct" estimates without a grid search. The estimate of θ^1, is as expected scattered around zero, with the highest standard deviation in the unrestricted case. This behaviour is also obvious, because the grid for the threshold search cannot be inferred from the VECM model, but must be inferred from the generated basis, which is larger by construction in the unrestricted case.

Table 3: VECM - model class $\mathcal{H}_1$

The VECM model belongs to the model class $\mathcal{H}_1$, ie there exists only one regime and the data generating process has no threshold θ^1. The table shows the means and the standard deviations (figures in brackets) of the estimates $\hat{\beta}_i$ and $\hat{\theta}^1$ for our 1,000 simulations, as well as the theoretical values fixed in the simulation. The last 6 columns show how often a threshold (alternative hypothesis) is incorrectly accepted at different confidence levels by the fixed regressor and the residual bootstrap. The abbreviation 'dyn' stands for a dynamic grid point setting, where the number of grid points is determined via a precision parameter. The estimate $\hat{\theta}^1$ is as expected statistically zero.

	grid	$\hat{\beta}_1$ (stdev)	$\hat{\beta}_0$ (stdev)	$\hat{\theta}^1$ (stdev)	fixed regressor α-error, CL			residual α-error, CL		
theor. value		1.10	-10	n/a	90%	95%	99%	90%	95%	99%
restricted	10	1.10	-10.01	0.06	0.10	0.06	0.01	0.06	0.04	0.00
		(0.01)	(0.36)	(2.17)						
restricted	50	1.10	-9.99	0.07	0.13	0.06	0.01	0.09	0.04	0.01
		(0.02)	(0.45)	(2.91)						
restricted	100	1.10	-10.02	-0.03	0.10	0.05	0.02	0.06	0.03	0.01
		(0.03)	(0.52)	(3.02)						
restricted	dyn	1.10	-10.00	0.01	0.09	0.05	0.01	0.07	0.03	0.01
		(0.02)	(0.40)	(3.01)						
unrestricted	50	1.10	-10.01	0.58	0.09	0.04	0.00	0.07	0.03	0.00
		(0.02)	(0.24)	(16.99)						

2.3.3 No-arbitrage model class $\mathcal{H}_2$, searching for one threshold

We investigate now a 1-threshold or 2-regime TVECM, ie we set $l = 2$ in Equation (25). In the following 1-threshold TVECM specification we chose $\beta_1 = 1.1$, $\beta_0 = 10$ and $\theta^1 = 3$ and present the results in Table 4.

The results of the simulation yield two immediate observations, the mean of $\hat{\beta}_1$ is practically independent of the grid setup and the outcomes of the bootstrap methodologies are also independent. The fixed regressor and residual bootstrap have very small β-errors, below 2% at 99% confidence level.

As expected the standard deviation for $\hat{\beta}_1$ are getting larger for coarser grid settings. In other words, we find a clear dependence of the quality of the estimator $\hat{\beta}_1$ from the grid size for a TVECM unlike in the VECM case presented in Table 3. This behaviour is expected, because now the VECM estimators in the error correction term, which are used to construct the grid search area (see discussion near the end of Section 2.1), are poor estimators of the "true" TVECM values. Hence, depending on how far-off the VECM estimators are from the true TVECM values, the grid can be considerable large. The distance between two neighbouring grid points has a strong influence on the quality of the estimation results. Surprisingly, we do not find this behaviour for $\hat{\beta}_0$ and $\hat{\theta}^1$, where the standard deviations for these estimators do not get smaller for finer grid settings. Obviously, it is difficult to simultaneously estimate these two parameters, as both represent a sort of vertical shift along the y-axis.

Table 4: 1-threshold TVECM - model class $\mathcal{H}_2$

The 1-threshold TVECM belongs to model class $\mathcal{H}_2$, ie there exist two regimes and a threshold θ^1. For further details we refer to Table 3.

	grid	$\hat{\beta}_1$ (stdev)	$\hat{\beta}_0$ (stdev)	$\hat{\theta}^1$ (stdev)	fixed regressor β-error, CL			residual β-error, CL		
theor. value		1.10	10	3	90%	95%	99%	90%	95%	99%
restricted	10	1.10	7.49	4.79	0.002	0.004	0.007	0.004	0.005	0.005
		(0.29)	(50.56)	(47.48)						
restricted	50	1.09	5.42	7.4	0.002	0.002	0.004	0.003	0.006	0.007
		(0.10)	(14.74)	(14.92)						
restricted	100	1.10	5.42	7.49	0.003	0.005	0.007	0.004	0.005	0.012
		(0.04)	(16.63)	(16.00)						
restricted	dyn	1.10	8.39	4.46	0.002	0.002	0.007	0.004	0.005	0.012
		(0.02)	(32.99)	(33.00)						
unrestricted	50	1.05	-1.67	14.28	0.002	0.005	0.008	0.002	0.004	0.010
		(0.84)	(111.16)	(121.17)						

This conjecture is also in line with the fact that independent of the grid size and even for the dynamic grid setting, the mean values of $\hat{\beta}_0$ and $\hat{\theta}^1$ from our 1,000 simulations are imprecise, ie far-off from their theoretical values, fixed in the data generating process. A comparison of Table 2 (Hansen and Seo model) with Table 4 suggests that the Hansen and Seo model seems to have a better behaviour. One possible reason is, that the dimensionality of the grid in the no-arbitrage model setting is higher as compared to the Hansen and Seo model.

However and most importantly, if we compute the sum of $\hat{\beta}_0$ and $\hat{\theta}^1$ for each simulation and compute the average and standard deviation we find highly promising results, which are presented in Table 5.

Table 5: Transaction costs in the 1-threshold TVECM - model class $\mathcal{H}_2$

The table shows for the same simulations as in Table 4 the sum $\hat{\beta}_0+\hat{\theta}^1$, which corresponds to the transaction costs based on our no-arbitrage argument.

	grid	$\hat{\beta}_0+\hat{\theta}^1$	stdev
theor. value		13	
restricted	10	12.28	4.37
restricted	50	12.82	3.14
restricted	100	12.91	1.33
restricted	dyn	12.85	1.17
unrestricted	50	12.25	7.75

The mean of the sum $\hat{\beta}_0+\hat{\theta}^1$ of the 1,000 estimates is very precise, with the best performance found for the dynamic grid setting for the important restricted case. This result is also graphically shown in Figure 5. This is a very convenient result, because for arbitrageurs the sum $\hat{\theta}^1+\hat{\beta}_0$, which represents the transaction costs, is important and not the individual pieces $\hat{\beta}_0$ and $\hat{\theta}^1$. We have tested this result for arbitrary parameters. We can conclude, that the results of the simulation of our model setup are as good as the results achieved for the Hansen and Seo model.

Figure 5: Distribution of 1-threshold TVECM parameter estimates

The distributions of the parameter estimates are based on 1,000 MC simulations. The first three graphs show the distribution of $\hat{\beta}_1$, $\hat{\beta}_0$ and $\hat{\theta}^1$ for our model in Equation (10) for the dynamic grid setup. The graph on the very right hand side presents the sum $\hat{\beta}_0+\hat{\theta}^1$, which is the estimator of the transaction costs. The distribution of $\hat{\beta}_1$ is very narrow, whereas we find large outliers for $\hat{\beta}_0$ and $\hat{\theta}^1$. The distribution of $\hat{\beta}_0+\hat{\theta}^1$ is again very narrow.

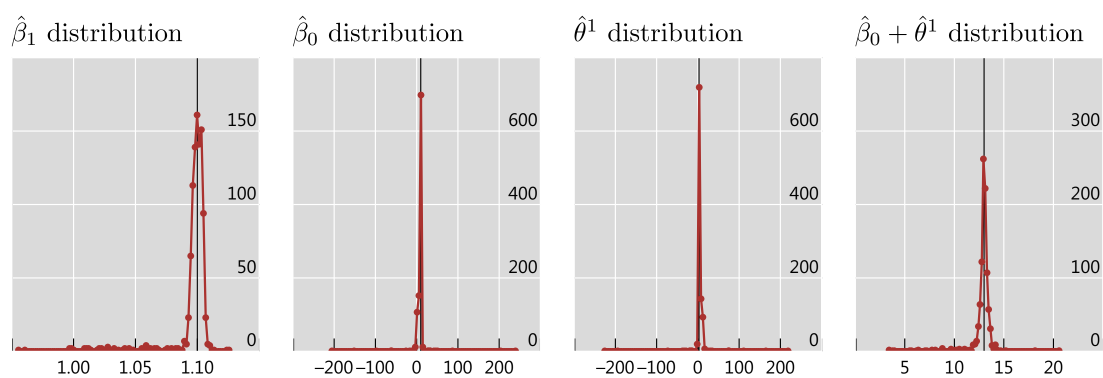

The graphs in Figure 5 contain, despite their similarity, several interesting features. For the $\hat{\beta}_0$, $\hat{\beta}_1$ and $\hat{\theta}^1$ estimators we find that the weight of the distribution is heavily centred around the values of the data generating process. However, we find several outliers in our 1,000 simulations for $\hat{\beta}_0$ and $\hat{\theta}^1$ far away from that theoretical value fixed in the data generating process. The range of the distribution is partly determined by the chosen grid. That is why the plot for the estimator of β_1 is shifted to the right, because we have calibrated the grid to include the value $\hat{\beta}_1 = 1$ (observed basis). The most important and interesting finding is, that if we look at the very right-hand graph in Figure 5, namely the sum of $\hat{\beta}_0 + \hat{\theta}^1$, we see that the dispersion is dramatically lower compared to its individual parts (two graphs in the middle).

The negative log-likelihood functions in Figure 6, which need to get minimized, show a non-smooth behaviour, hence the proposed grid search is necessary. The graphs give also an immediate explanation why the estimation of β_1 is so precise (assuming a sufficiently fine grid is chosen) and why the estimators of β_0 and θ^1 are poor.

Figure 6: Negative log-likelihood functions for the 1-threshold TVECM

The figures show the negative log-likelihood functions for three out of our 1,000 simulations. It is obvious that the functions are non-smooth. It also shows why we find a very high precision for the estimator $\hat{\beta}_1$ and a poor precision for the estimators $\hat{\beta}_0$ and $\hat{\theta}^1$.

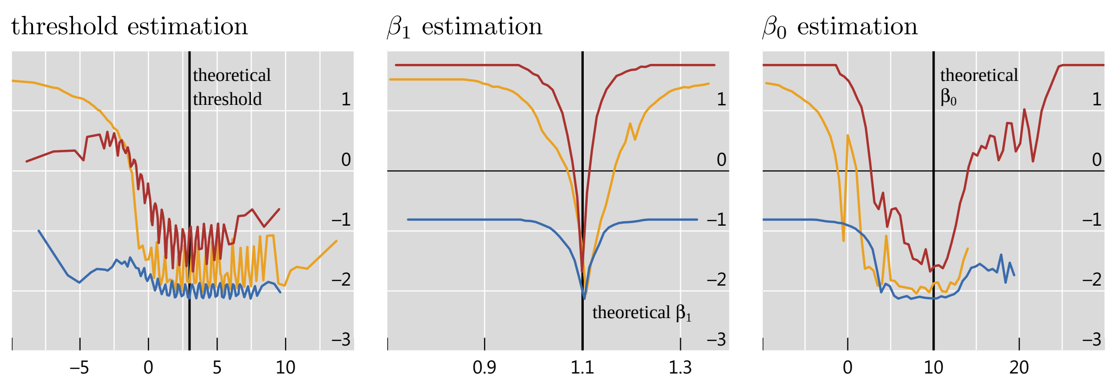

2.3.4 No-arbitrage model class $\mathcal{H}_2$, searching for two thresholds

As a next and important step we generalize the simulations and are going to search and test for two thresholds. We are making one assumption, based on purely economic reasoning, namely that, we have at most one positive and/or one negative threshold, which correspond to the two transaction costs for a positive basis trade and a negative basis trade. This assumption relates to our no-arbitrage discussion in Section 1 which in turn links straight to the three classes of models $\mathcal{H}_i$ (see Figure 1).

The previous analysis has shown that the dynamic grid point setting is the most precise methodology to estimate the transaction costs (see Table 5), even though computationally more costly. Therefore, we will from now on use this grid setting only.

We use the same model parameters and the same model, a 1-threshold TVECM (model class $\mathcal{H}_2$), as in the previous section, however, test if our two bootstrap methodologies reject the existence of a second threshold. The data generating process has no second threshold, ie θ^2 does not exist and the estimate $\hat{\theta}^2$ should be distributed around zero and the existence of θ^2 should be rejected by our bootstrap methods.

Our previous simulation has shown that β_1 is estimated with high precision if we use a dynamic grid setup. In order to save compute time we fix β_1 during the second threshold search to the parameter value found in the first threshold search and hence reduce the search for the second threshold to a 2 dimensional grid for β_0 and θ^2.

In order to improve the understanding of the reliability of the methodology, which includes now a second threshold search, we vary also the generated time series length. We present results for several different setups in Table 6, however we will use only the dynamic grid point setting.

Table 6: 1-threshold TVECM - model class $\mathcal{H}_2$, β_1 fixed

This table shows the parameter estimates for the threshold search for θ^1 and θ^2 as well as the theoretically expected values. All results are generated with a dynamic grid setup and β_1 is fixed in the threshold search for θ^2.

		$\hat{\beta}_1$	$\hat{\beta}_0$	$\hat{\theta}^1$	$\hat{\beta}_0 + \hat{\theta}^1$	$\hat{\beta}_0$	$\hat{\theta}^2$
	theor. value	1.10	10	3	13	10	n/a
type	periods						
restricted	1,000	1.10	8.39	4.46	12.85	8.64	-10.38
		(0.02)	(32.99)	(33.00)	(1.17)	(19.90)	(19.82)
unrestricted	1,000	1.10	11.05	2.00	13.05	16.29	-17.83
		(0.01)	(26.29)	(26.33)	(1.86)	(49.91)	(49.78)
restricted	2,000	1.10	9.35	3.57	12.93	8.09	-9.51
		(0.01)	(7.93)	(7.91)	(0.72)	(13.83)	(13.72)

As in the previous simulations, the estimator $\hat{\beta}_1$ as well as $\hat{\theta}^1 + \hat{\beta}_0$ are very precise, with the highest precision achieved in the restricted case and especially in the case where the tested time series has a length of 2,000. We use a dynamic grid setup, which is the reason why we find standard deviations which are similar for the estimators $\hat{\beta}_0$ and $\hat{\theta}^1$ in the unrestricted case compared to the restricted case in the first threshold search. In Table 4 the relevant standard deviations of the unrestricted case are much bigger compared to the figures of the restricted case.[13]. The estimator $\hat{\theta}^2$ is distributed around zero, as expected, because we do not have a second threshold.

The results of the reliability of the bootstrap algorithms for both threshold searches are given in Table 7. Again, we find that both tests, the fixed regressor and the residual bootstrap, produce consistent and acceptable results in terms of α- and β-errors.

The β-errors are very small, as expected from the results in Section 2.3.3. The α-errors of the second threshold search are higher than the α-errors in the first threshold search which are presented in Table 3, however still acceptable.

Table 7: Bootstrap tests for the 1-threshold TVECM - model class $\mathcal{H}_2$, β_1 fixed

This table shows the β errors for the threshold search for θ^1 and the α errors of the threshold search for θ^2. We use the same ordering of model setups as in the previous Table 6.

first threshold bootstrap						second threshold bootstrap					
fixed regressor			residual			fixed regressor			residual		
90%	95%	99%	90%	95%	99%	90%	95%	99%	90%	95%	99%
0.002	0.002	0.007	0.004	0.005	0.012	0.21	0.15	0.07	0.25	0.15	0.06
0.003	0.003	0.006	0.004	0.004	0.014	0.35	0.30	0.23	0.36	0.30	0.18
0.000	0.000	0.000	0.001	0.001	0.001	0.23	0.16	0.09	0.27	0.16	0.09

In order to test the robustness of the method, where we keep β_1 fixed in the threshold search for θ^2, we repeat the simulation and search now also for an optimal $\hat{\beta}_1$ in the threshold search for θ^2 using the log-likelihood method. The threshold search for θ^1 is performed in the usual manner, hence we only need to report the figures for the threshold search for θ^2. As an illustration, we focus on the restricted case and a time series length of 1,000 in our data generating process. We present the results in Table 8.

[13] Table 4 shows results for the restricted and unrestricted case for the example of a grid size of 50 grid points. The standard deviations in the unrestricted case are 7-8 times bigger.

Table 8: 1-threshold TVECM - model class $\mathcal{H}_2$, threshold search for θ^2, β_1 variable

This table shows the parameter estimates for the threshold search for θ^2 as well as the theoretically expected values. All results are generated with a dynamic grid setup and this time β_1 is kept variable and a grid search is applied also for β_1 in the threshold search for θ^2.

	$\hat{\beta}_1$	$\hat{\beta}_0$	$\hat{\theta}_2$	fixed regressor α-error, CL			residual α-error, CL		
theor. value	1.10	10	n/a	90%	95%	99%	90%	95%	99%
restricted	1.02 (0.13)	8.74 (17.25)	-9.94 (17.22)	0.28	0.19	0.09	0.28	0.19	0.09

The results presented in Table 8 lead to two immediate observations: firstly, the estimator $\hat{\beta}_1$ is poor and the α-errors are larger than the ones presented in the first row and the last six columns of Table 7.

We can conclude that the performance of the second threshold search is better if β_1 is kept fixed to the value found in search for θ^1. Keeping β_1 fixed reduces also the computational burden of the method.

2.3.5 No-arbitrage model class $\mathcal{H}_3$

Finally, we present the results for a data generating process with two thresholds. Without loss of generality we have chosen the following parameters: $\beta_1 = 1.1$, $\beta_0 = 1$ as well as two thresholds $\theta^1 = -4$ and $\theta^2 = 6$. We use again the dynamic grid point method and a time series of length 1,000 as the baseline setup. Further, as part of our baseline case, we keep β_1 for the second threshold search fixed, because we have shown in Section 2.3.4 that this produces best results. However, we perform a variety of robustness checks, such as we extend the time series length to 2,000, analyse the unrestricted case, as well as test results where we keep β_1 variable in the second threshold search. We will start with the presentation of the results for the first threshold search in Table 9. Depending on the generated time series either of the two thresholds, ie θ^1 or θ^2, may be found in the first round.

Table 9: 2-threshold TVECM - model class $\mathcal{H}_3$, first threshold

This table shows the results of the first threshold search for our 2-threshold respective 3-regime TVECM. The point estimates of θ^i are imprecise, simply because we have a chance to find either of the two thresholds. The estimates of the two transaction costs are again very good.

		$\hat{\beta}_1$	$\hat{\beta}_0$	$\hat{\theta}$	$\hat{\beta}_0 + \hat{\theta}^1$	$\hat{\beta}_0 + \hat{\theta}^2$
	theor. value	1.10	1	-4 or 6	-3	7
type	periods	(stdev)	(stdev)	(stdev)	(stdev)	(stdev)
restricted	1,000	1.10	0.08	2.43	-2.98	6.94
		(0.02)	(6.15)	(8.69)	(0.65)	(1.59)
unrestricted	1,000	1.10	-1.02	4.00	-3.16	6.82
		(0.01)	(17.47)	(18.38)	(1.51)	(1.83)
restricted	2,000	1.10	-0.61	3.05	-2.97	6.88
		(0.01)	(8.04)	(10.18)	(0.62)	(1.26)

The estimation of β_1 is again very precise, whereas the estimate of β_0 is poor. The imprecise estimation of the threshold (the estimator is denoted as $\hat{\theta}$ in Table 9) is due to two reasons, firstly, we know from the one threshold case that the estimation of the threshold is imprecise and secondly, in the two threshold case the estimation procedure may find either the negative or the positive threshold. Again, the most important finding is, that the sum of $\hat{\beta}_1$ and the threshold is for the negative values very close to $\beta_0 + \theta^1 = -3$ and for the positive values very close to $\beta_0 + \theta^2 = 7$. The best performance is, as expected, achieved for longer time series.

The power of the two bootstrap tests is very high, as expected from the previous simulation results (see Table 10). The lowest β-errors are realized for the time-series with 2,000 periods.

Table 10: Bootstrap tests for the 2-threshold TVECM - model class $\mathcal{H}_3$, first threshold

This table shows the β-errors of the two bootstrap methods. The parameter estimates can be found in the previous Table 9.

		fixed regressor β-error, CL			residual β-error, CL		
	periods	90%	95%	99%	90%	95%	99%
restricted	1,000	0.000	0.001	0.001	0.000	0.000	0.002
unrestricted	1,000	0.000	0.001	0.002	0.001	0.001	0.004
restricted	2,000	0.000	0.000	0.000	0.000	0.000	0.000

The figures for the second threshold search are presented in Table 11 and are extremely promising, because the two transaction costs are well estimated. In addition to the three simulations presented Tables 9 and 10, we add for the second threshold search also the simulation, where β_1 is kept variable in the search for the second threshold. In this simulation we do not keep β_1 fixed to the value found in the first threshold search, but perform a grid search to find the "optimal" value based on our maximum likelihood method.

Table 11: 2-threshold TVECM - model class $\mathcal{H}_3$, second threshold

This table shows the results of the second threshold search for our 2-threshold TVECM, respective 3-regime TVECM. The estimator $\hat{\theta}$ is imprecise, simply because we have a chance to find either of the two thresholds. In addition to the first threshold search presented in Table 9 we also include a simulation where we kept β_1 variable (last two rows), ie where we have not fixed β_1 in the second threshold search to the estimate found in the first threshold search. This simulation is denoted by the appreciation "var" for variable. The precision of the estimates of the two transaction costs is again very high. β_1 is not reestimated, except for the simulations in the last two rows, hence we filled in n/a.

		$\hat{\beta}_1$	$\hat{\beta}_0$	$\hat{\theta}$	$\hat{\beta}_0 + \hat{\theta}^1$	$\hat{\beta}_0 + \hat{\theta}^2$
	theor. value	1.10	1	-4 or 6	-3	7
type	periods	(stdev)	(stdev)	(stdev)	(stdev)	(stdev)
restricted	1,000	n/a	0.84	0.43	-2.91	6.28
			(4.28)	(7.56)	(1.80)	(2.22)
unrestricted	1,000	n/a	1.65	-1.48	-3.44	5.94
			(18.51)	(19.97)	(4.15)	(2.82)
restricted	2,000	n/a	1.175	0.04	-2.97	6.28
			(6.39)	(9.27)	(2.43)	(2.30)
restricted	1,000/var	1.10	1.06	-0.21	-2.95	5.55
		(0.05)	(8.37)	(10.52)	(2.09)	(2.63)

We find a similar picture as for the first threshold search. The threshold estimate is poor. However, the sum of $\hat{\beta}_0$ and the threshold is for the negative values very close to $\beta_0 + \theta^1 = -3$ and for the positive values very close to $\beta_0 + \theta^2 = 7$. The β_1 estimation in the variable case (last two lines in Table 11) is as expected from previous discussions poorer compared to the cases where we have fixed the β_1 value to the estimator $\hat{\beta}_1$, found in the first threshold search. In general, the precision of the parameter estimates in the second search is not as good as compared to the first search.

In line with the previous simulation results, the power of the two bootstrap tests, presented in Table 12, is very high, even though not as high as for the first threshold search (see Table 10).

Table 12: Bootstrap tests for 2-threshold TVECM - model class $\mathcal{H}_3$, second threshold

This table shows the β-errors of the two bootstrap methods. The parameter estimates can be found in the previous Table 11.

		fixed regressor β-error, CL			residual β-error, CL		
	periods	90%	95%	99%	90%	95%	99%
restricted	1,000	0.007	0.008	0.009	0.006	0.007	0.009
unrestricted	1,000	0.015	0.016	0.018	0.014	0.019	0.021
restricted	2,000	0.008	0.008	0.010	0.007	0.007	0.013
restricted	1,000/var	0.011	0.011	0.013	0.009	0.010	0.014

The negative log-likelihood function for three different sample simulations is shown in Figure 7. The negative log-likelihood function for the threshold estimation is now more complex, reflecting the fact that we have two thresholds. The functional form shows immediately the non-differentiable structure. Further, it is again evident, why the estimate of the threshold and β_0 is imprecise. The estimation of β_1 is as in the previous cases very accurate.

Figure 7: Negative log-likelihood function - 2-threshold TVECM, first threshold

The figures show the negative log-likelihood functions for three sample simulations of the first threshold search. The functions are non-differentiable and a high precision can only be expected for β_1. We find evidence for a low individual precision of β_0 and θ^i. The two vertical lines in the left-hand graph represent the two thresholds θ^1 and θ^2.

threshold estimation | β_1 estimation | β_0 estimation

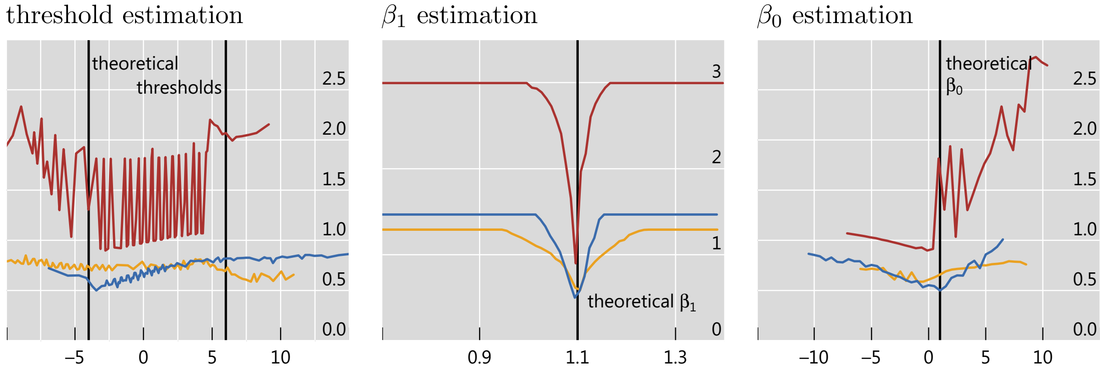

The distribution of the estimators $\hat{\beta}_0$, $\hat{\beta}_1$, $\hat{\theta}^i$ and the transaction costs for the first and second threshold search are presented in Figure 8. The distribution of $\hat{\beta}_0$ and the thresholds is wide in line with the form of the negative likelihood function in Figure 7. In each search we can either find θ^1 or θ^2, which cannot be distinguished due to the imprecise estimation. However, looking at the transaction costs (right-hand panel), which are estimated at a good precision, we can clearly unravel the two different transaction costs.

Figure 8: Distribution of 2-threshold TVECM parameter estimates, β_1 fixed

The distribution of the parameter estimates are based on 1,000 MC simulations. The first three graphs in the upper panel show the distribution of the three estimates for our model in Equation (20). The lower panel does not show the distribution of $\hat{\beta}_1$, as it was fixed to the value found in the first search. The graph on the very right shows the estimation of the transaction costs $\hat{\beta}_0 + \hat{\theta}^i$. The distribution of $\hat{\beta}_1$ is very narrow, whereas we find large outliers for $\hat{\beta}_0$ and $\hat{\theta}^i$. The distribution of $\hat{\beta}_0 + \hat{\theta}^i$ is also very narrow.

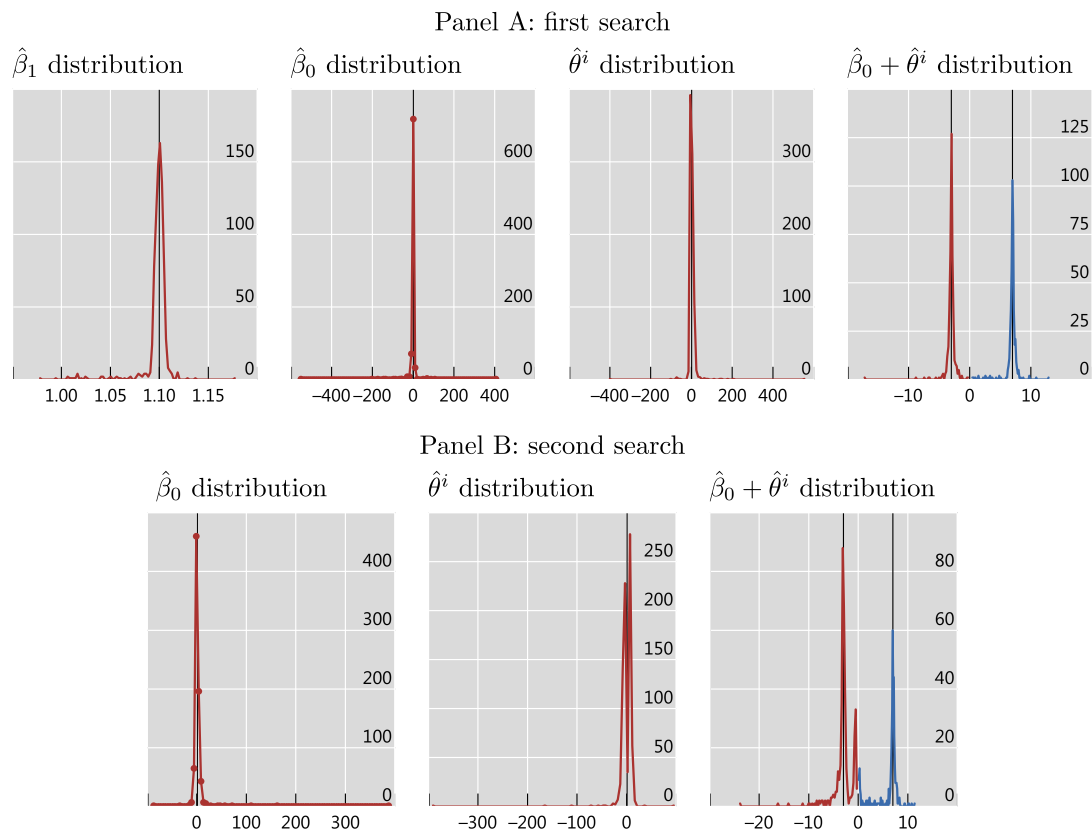

2.3.6 Conclusion from the simulation study

The simulations have revealed the following outcomes: The method discussed in Sections 2.1 and 2.2 leads to stable and robust results in the VECM case as well as for the 1- and 2-threshold TVECM. The transaction costs $\theta^i + \beta_0$ as well as β_1 can be estimated with high accuracy, whereas the individual distribution of $\hat{\theta}^i$ and $\hat{\beta}_0$ is very wide. The dynamic grid search strategy, where the grid size is determined with an accuracy parameter (distance between two grid points of the equidistant grid), has turned out to be superior. The value of β_1 can be fixed in the second threshold search to the value found in the first search. This yields best results and reduces the computational burden dramatically.

3 Application

We apply the proposed TVECM methodology to different markets that exhibit a non-zero basis (either in contango or backwardation) for two similar financial market instruments traded in the spot and the derivative market. We are discussing two types of basis trades, for commodities, such as gold and platinum and for stock indices, such as DAX and S&P 500.

It is unlikely that transaction costs remain stable over a long period of time, due to changing economic environments (index basis trading), weather and climate changes (demand for agricultural products) or technological improvements (demand for raw materials). Therefore, it is necessary to either have strong arguments for structural breaks along the time axis or to test for such breaks. This section is for illustration of the proposed methods only and not a comprehensive analysis of commodity trading. Nevertheless, we have made a search for a threshold over different time periods. In most case we have daily data from 2000 until December 2016. During this period, the following possible candidates for structural breaks exists: the Lehman Brothers' bankruptcy (15.09.2008), the beginning of the Greek sovereign debt crisis (20.10.2009), the start of the European sovereign debt crisis as defined by van Rixtel and Gasperini (2013) (01.04.2010) and the speech of ECB president Mario Draghi (26.07.2012).

In our application we present results for our dynamic grid setup for different periods and we run our model for 1 to 5 lags in the VAR term. We use the Schwarz information criterion (SIC) to determine the model with the best fit.

Further as an additional illustration of the usage of (T)VECM models we introduce briefly two measures of price finding and the terminology of a half-life of shock absorption. From the speed of adjustments $\lambda = (\lambda_1\,\lambda_2)^\mathsf{T}$ in Equation (1)[14] we can compute the Has-

[14] We have skipped the superscript j at λ indicating the regimes in order to keep the notation simple. However, measures of price finding and the half-life depend on the regime.

brouck (HAS) and Gonzalo-Granger (GG) measures of price discovery. The two relevant and independent HAS measures (Man and Wu; 2013) are defined as follows:

$$\mathrm{HAS}_1 = \frac{\lambda_2^2\left(\sigma_1^2 - \frac{\sigma_{12}^2}{\sigma_2^2}\right)}{\lambda_2^2\sigma_1^2 - 2\lambda_1\lambda_2\sigma_{12} + \lambda_1^2\sigma_2^2} \qquad \text{and} \qquad \mathrm{HAS}_2 = \frac{\left(\lambda_2\sigma_1 - \lambda_1\frac{\sigma_{12}}{\sigma_1}\right)^2}{\lambda_2^2\sigma_1^2 - 2\lambda_1\lambda_2\sigma_{12} + \lambda_1^2\sigma_2^2}\,. \tag{26}$$

To estimate the Equations (26) we rely on the estimated covariance matrix from the VECM to capture the terms σ_1^2, σ_{12} and σ_2^2. In the following we define HAS as the average of HAS_1 and HAS_2.

The second indicator for price discovery, the GG measure (Gonzalo and Granger; 1995) decomposes the common factor itself, but ignores the correlation of the innovations in the two markets. The following two measures exist

$$\mathrm{GG}^{\mathrm{spot}} = \frac{-\lambda_2}{\lambda_1 - \lambda_2} \qquad \text{and} \qquad \mathrm{GG}^{\mathrm{futures}} = \frac{\lambda_1}{\lambda_1 - \lambda_2}\,, \tag{27}$$

whereby it is obvious that $\mathrm{GG}^{\mathrm{spot}} + \mathrm{GG}^{\mathrm{futures}} = 1$.

HAS and GG measures greater than 0.5 imply that more than 50% of the price discovery occurs in the spot market. When the measures are close to 0.5 both markets contribute to price discovery without evidence on which market is dominant. GG and HAS below 0.5 suggest price leadership of the futures market.[15]

Finally, the vector error correction mechanism directly links the speed of adjustments to the cointegration residual u_t which follows an implied AR(1) process:

$$u_t = (1 + \lambda_1 - \beta_1\lambda_2)u_{t-1} + \varepsilon_t^{\mathrm{spot}} - \beta_1\varepsilon_t^{\mathrm{futures}} \equiv \phi u_{t-1} + \varepsilon_t^{\mathrm{spot}} - \beta_1\varepsilon_t^{\mathrm{futures}}\,. \tag{28}$$

The half-life of a shock, hl, can now be calculated from the AR(1) coefficient ϕ as:

$$hl = \frac{ln(0.5)}{ln(\phi)}\,. \tag{29}$$

All our time series used in the subsequent sections are as required I(1) (see Appendix A) and cointegrated (see Appendix B).

3.1 Commodities - raw materials

In the next section we empirically estimate unknown transaction costs in the markets for gold and platinum by analysing the basis defined as the spot price minus the price of

[15] Unlike the HAS measure, the GG measure is mathematically not confined to the interval [0,1] which seems to make an interpretation similar to the Hasbrouck measure difficult. GG measures below 0 and above 1 should be interpreted as 0 and 1, respectively (see Gyntelberg et al. (2013) for more discussion).

the future on the same underlying. The basis highly depends on the spot price that is directly impacted by supply and demand for the specific commodity, storage costs and profit margins.

3.1.1 Gold

Gold is a precious metal and as such most popular as an investment. It is subject to speculation and volatility and gold has the most effective safe haven and hedging properties across a number of countries, in a sense that it performs relatively strong in extreme market conditions (Low et al.; 2015). Therefore, we expect several structural breaks along the time dimension, representing different economic conditions. The threshold search suggests for the period starting in 2000 and the period starting with the Draghi speech (26. July 2012) that the VECM is the most appropriate model. As an illustration we present the results for the period starting with the European sovereign debt crisis, as defined by van Rixtel and Gasperini (2013), ie 1st April 2010 until end December 2016, for which we find a 1-threshold TVECM. The gold spot and future prices as well as the basis are presented in Figure 9. The basis (right-hand side in Figure 9) fluctuates between positive and negative values. In principle, from the basis plot, we could have expected three regimes or two thresholds.

Figure 9: Gold spot, future prices and spot-future basis

The figure on the left-hand side shows the gold spot and the future prices, for the period starting 1st April 2010, the start of the euro area sovereign debt crisis as defined by van Rixtel and Gasperini (2013). The prices are given in USD for a contract size of 100troy oz. The figure on the right-hand side shows the basis for the same period and the transaction costs for a positive basis trade. Source: Bloomberg, authors' calculations.

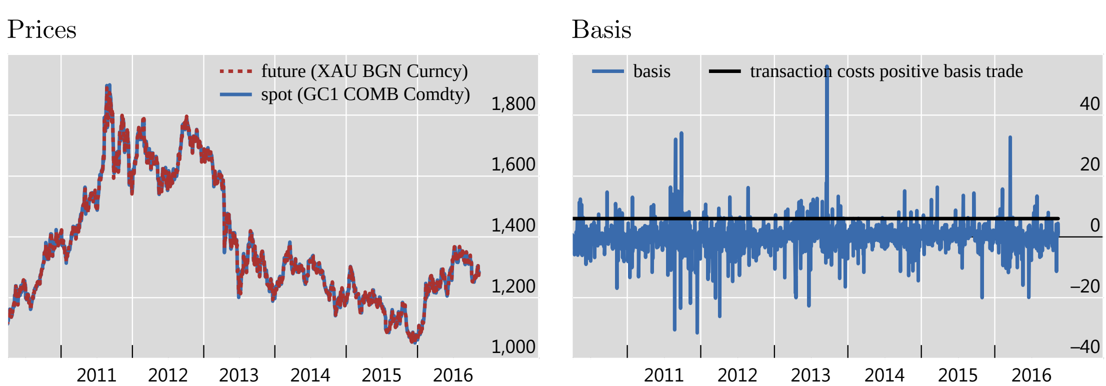

The detailed results for the threshold search are presented in Table 13. The estimated thresholds for different VAR lags are very stable and in the range of 4USD to 7USD. For all estimations we find the best model fit (lowest SIC) for a VAR lag of 1, however, the thresholds appear to be significant at lag 4 and 5 only.

Table 13: Threshold search for θ^1 - Gold

The table shows the results for the threshold search for θ^1 for the period April 2010 until end December 2016. The best model fit is achieved at lag=1, based on the SIC. The estimated transaction costs $\beta_0 + \theta^1$ are stable across all analysed lag lengths, but only significant beyond and including lag=4. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^1$	λ^1_{1L}	λ^1_{2L}	λ^1_{1U}	λ^1_{2U}	% upper regime
1	**8.55**	0.30	0.39	0.997	4.58	0.45*	1.49***	0.20*	1.12***	45
2	8.61	0.41	0.50	0.996	4.57	0.35	1.39***	0.36*	1.00***	56
3	8.63	0.75	0.82	0.997	6.19	0.39	1.23***	0.80***	1.42***	25
4	8.68	**0.07**	**0.10**	0.997	6.27	0.34	1.09***	0.92***	1.51***	25
5	8.74	**0.09**	0.13	0.997	6.41	0.38	1.09***	0.88***	1.46***	25

The figures in Table 13 suggest transaction costs $\beta_0 + \theta^1$ to be around 6USD based on the two lags (4 and 5) with significant thresholds as well as the findings for the model with the lowest SIC (lag=1). We find no evidence for a second threshold.

We present in Table 14 measures of price discovery for lag=4 where a significant threshold has been found. We find strong indication for the spot market to lead the price discovery process. The half-lives suggest a fast adjustment process of below 30mins in the upper regime. We do not report a half-life in the neutral regime, as it turns out to be negative. This implied unstable dynamics is due to speed of adjustments that are either insignificant or have a negative sign.

Table 14: Measures of price discovery - Gold

The table shows price discovery measures for lag=4. We report only two regimes, whereby in the reported case the lower regime is identical to the neutral regime. The GG measures with the superscript + should be interpreted as 1. The half-life in the lower regime is not reported as one speed of adjustment is insignificant and the other has a wrong sign.

regime	HAS	GG	leading market	half-life (in 30min)	comments
upper	0.65	2.56^{+}	spot	0.79	
lower (neutral)	0.65	1.45^{+}	spot	-	half-life negative

3.1.2 Platinum

Platinum is an extremely scarce metal. It is one of the least reactive metals and has a high resistance to corrosion. It is considered a precious metal. It has a wide range of usage, such as in catalytic converters, as an investment product or in laboratory equipment (due to its chemical inertness).

We analyse the period from 2000 until December 2016 where we find evidence for two thresholds. The time series of the spot and future prices, the basis as well as the two transaction costs are displayed in Figure 10. The basis fluctuates mainly between -20USD and 20USD.

Figure 10: Platinum spot, futures price and spot-future basis

The figure on the left-hand side shows the platinum spot and the futures price. The prices are displayed in USD for a contract size of 50troy oz. The figure on the right-hand side shows the basis and the estimated transaction costs for a positive and a negative basis trade. Source: Bloomberg, authors' calculations.

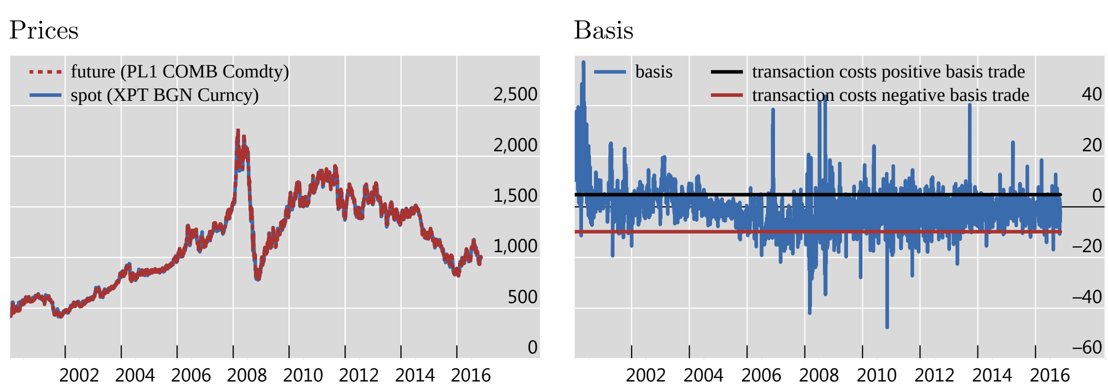

The threshold search for θ^1 results in the best model fit based on the SIC and a significant threshold at lag=3 (see Table 15). Our estimated transaction costs for a positive basis trade are 4.8USD for a contract size of 50troy oz.

Table 15: Threshold search for θ^1 - Platinum

The table shows the results for the threshold search for θ^1 for the period April 2000 until end December 2016. The best model fit and a significant threshold is found at lag=3. The estimated transaction costs $\beta_0 + \theta^1$ for a positive basis trade are 4.8USD. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^1$	λ^1_{1L}	λ^1_{2L}	λ^1_{1U}	λ^1_{2U}	% upper regime
1	8.96	**0.02**	**0.04**	0.992	8.66	0.06	0.42***	-0.05	0.28***	42
2	8.93	**0.03**	**0.04**	0.993	13.48	0.07	0.37***	0.05	0.26	11
3	**8.93**	**0.09**	0.12	1.002	4.77	0.12	0.34***	0.05	0.12***	10
4	8.95	**0.09**	0.13	1.003	3.42	0.15	0.34**	0.04	0.09***	11
5	8.93	**0.04**	**0.07**	0.995	11.16	0.10	0.35	-0.02	0.07***	10

The threshold search for θ^2 does not provide any evidence for threshold for lag=1 and 2 of the VAR term. However, for lag=3 and higher we find a statistically significant second threshold and a $\beta_0 + \theta^2$ in the order of -9.8, as shown in Table 16. This means that our estimated transaction costs for a negative basis trade are 9.8USD for a contract size of 50troy oz.

Table 16: Threshold search for θ^2 - Platinum

The table shows the results for the threshold search for θ^2 for the period April 2000 until end December 2016. The best model fit and a significant threshold is found at lag=3. The estimator $\hat{\beta}_0 + \hat{\theta}^2$ for a negative basis trade is -9.8 which means that the transaction costs are estimated to be 9.8USD. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^2$	λ^2_{1L}	λ^2_{2L}	λ^2_{1U}	λ^2_{2U}	% lower regime
3	**8.95**	**0.02**	**0.04**	1.002	-9.79	0.08	0.21	0.02	0.14^{***}	14
4	8.97	0.14	0.17	1.003	-11.13	0.08	0.20	0.01	0.10^{***}	13
5	8.95	**0.03**	**0.05**	0.995	-2.37	0.02	0.18	-0.01	0.15^{***}	11

Both transaction costs are also presented in Figure 10. The upper and the lower regime are clearly arbitrage regimes in the sense that only 10% and 14% of the total observations are part of these regimes, respectively (see right-hand column in Tables 15 and 16).

In Table 17 we present measures of price discovery for the relevant lag=3. We find relative strong evidence for a price leadership of the spot market.

Table 17: Measures of price discovery - Platinum

The table shows price discovery measures for lag=3. The GG measures with the superscript $^+$ should be interpreted as 1. The price discovery measures of the neutral regime are means of the estimation from the first and second threshold search.

regime	HAS	GG	leading market	half-life (in 30min)	
upper	0.72	1.63^+	spot	9.52	
neutral	0.58	1.34^+	spot	4.13	longer half-life expected
lower	0.56	1.62^+	spot	5.02	λ^2_{iL} are insignificant

The results in Table 17 suggest that shocks are absorbed within half a trading day, surprisingly faster in the neutral regime.

3.2 Index trading

Finally, we also estimate unknown transaction costs on arbitrage trades between a spot index, such as the DAX and the S&P 500 and their futures indices. The basis is again defined as the difference between the spot and the future prices.

3.2.1 DAX

The DAX is the German blue chip stock index formed out of the 30 major German companies trading at the Frankfurt Stock Exchange. We have analysed daily data of the DAX (spot) and FDAX (futures) for several periods starting from January 2000. In Figure 11 we present the futures and spot prices as well as the basis for the period starting from the speech of ECB president M. Draghi (26.07.2012). The basis moves in a range mainly between -100 and 100, suggesting the existence of three regimes. Indeed, the threshold search results in one significant positive threshold and one significant negative threshold, leading to transaction costs for a positive basis trade in the range of 20 index points and transaction costs for negative basis trading of around 48 index points. The transaction costs are shown as horizontal lines in the right-hand panel of Figure 11.

Figure 11: DAX spot, future prices and spot-future basis

The figure on the left-hand side shows the DAX and FDAX for the period starting with the speech of ECB president M. Draghi (26.07.2012). The prices are given in index points. The figure on the right-hand side shows the basis as well as the two transaction costs for a positive and a negative basis trade. Source: Bloomberg, authors' calculation.

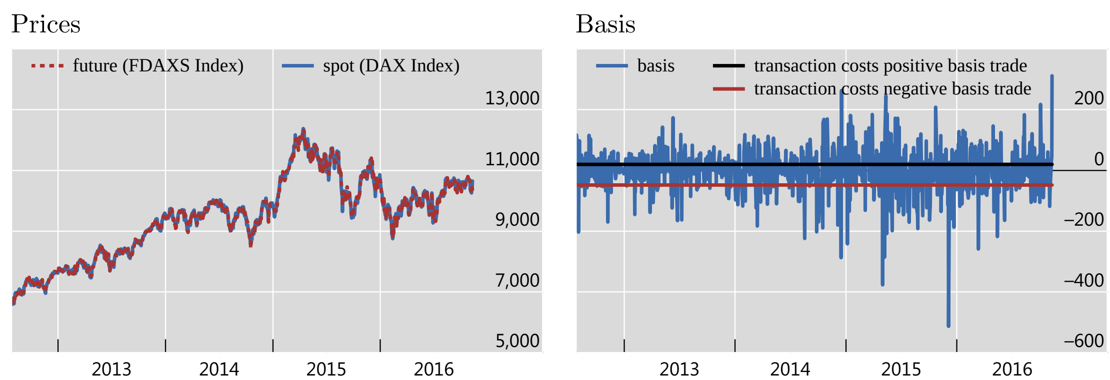

The results of the threshold search for θ^1 are presented in Table 18. The optimal lag length (based on the SIC) is two, however, the threshold is significant for four VAR lags. No results are shown for lag=1, because a VECM model is found to be the optimal model. The estimated transaction costs are robust over various lags (see Table 18).

Table 18: Threshold search for θ^1 - DAX

The table shows the results for the threshold search for θ^1 for the period starting with the speech of ECB president M. Draghi (26.07.2012). We do not show the result for one lag, as a VECM is the best model. The estimated transaction costs are robust for the tested lags and in the range of around 20 index points. The minimum value of the SIC is found for two lags and the most parsimonious model showing a significant threshold exists at four lags. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^1$	λ^1_{1L}	λ^1_{2L}	λ^1_{1U}	λ^1_{2U}	% upper regime
2	**17.59**	0.23	0.31	0.997	18.00	0.19	1.17***	-0.02	0.75***	57
3	17.64	0.19	0.27	0.997	20.03	0.30	1.22***	-0.13	0.61***	58
4	17.66	**0.08**	0.15	0.997	19.69	0.51*	1.43***	-0.14	0.62***	58
5	17.71	0.18	0.28	0.997	23.16	0.56*	1.36***	0.2	0.74***	58

The threshold search for θ^2 (see Table 19) produces only a significant result for lag=4. For the other lag choices we find no significant threshold or the 1-threshold TVECM is the better model choice compared to any 2-threshold TVECM. The evidence for a second threshold is weak based on the bootstrap results.

Table 19: Threshold search for θ^2 - DAX

The table shows the results for the threshold search for θ^2 for the period starting with the speech of ECB president M. Draghi (26.07.2012). We present only the fourth lag, because the other lags do not show a significant second threshold.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^2$	λ^2_{1L}	λ^2_{2L}	λ^2_{1U}	λ^2_{2U}	% lower regime
4	17.66	0.10	0.17	0.997	-47.84	0.96**	1.81***	0.02	0.94***	10

We present measures of price discovery and the half-life for the relevant fourth lag in Table 20. The spot market, ie the DAX, is leading the price discovery in all regimes. The half-lives suggest a very fast shock absorption, which takes place within the first 30min interval. We find surprisingly a shorter half-life in the neutral regime.

Table 20: Measures of price discovery - DAX

The table shows price discovery measures for lag=4. The GG measure with the superscript $^+$ should be interpreted as 1. The price discovery measures of the neutral regime are means of the estimation from the threshold search for θ^1 and θ^2.

regime	HAS	GG	leading market	half-life (in 30min)	comments
upper	0.62	0.81	spot	0.49	
neutral	0.61	0.99	spot	0.28	longer half-life expected
lower	0.75	2.15^+	spot	0.37	

3.2.2 S&P 500

As an illustration we present the analysis of the Standard & Poor's 500, abbreviated as the S&P 500 index, and its futures index for the period starting in 2000. The S&P 500 is an American stock market index based on the market capitalizations of 500 large companies having common stock listed on the New York stock exchange or NASDAQ. The index and the future prices, as well as the basis are presented in Figure 12.

Figure 12: S&P 500 spot, futures price and spot-future basis

The figure on the left-hand side shows the S&P 500 index and its futures for the period starting in 2000. The prices are given in index points. The figure on the right-hand side shows the basis as well as the two transaction costs for a positive and a negative basis trade. Source: Bloomberg, authors' calculations.

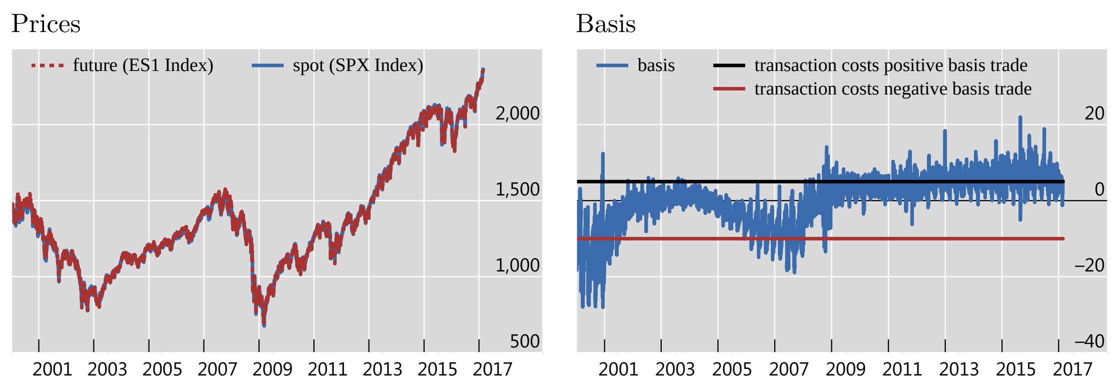

The threshold search suggests for the period starting in 2000 that we have average transaction costs on a negative basis trade of around 10 index points and on a positive basis trade of around 5 index points.

The detailed results for the threshold search for θ^1 are shown in Table 21. The transaction costs for a negative basis trade are fairly stable (exception is found for lag=2) and are statistically highly significant.

Table 21: Threshold search for θ^1 - S&P 500

The table shows the results for the threshold search for θ^1 for the period starting in 2000. The estimated transaction costs are significant at lag=3, for which the SIC also has a local minimum. The transaction costs are around -10 index points. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^1$	λ^1_{1L}	λ^1_{2L}	λ^1_{1U}	λ^1_{2U}	% lower regime
1	7.47	**0.00**	**0.00**	1.007	-7.68	0.01	0.06^{*}	0.43^{**}	0.74^{***}	48
2	7.43	**0.00**	**0.00**	0.998	-1.77	0.02	0.09	0.12^{*}	0.17^{***}	21
3	**7.38**	**0.03**	**0.06**	1.001	-10.13	-0.01	0.04	0.17^{***}	0.23^{***}	10
4	7.39	**0.07**	0.14	1.003	-12.41	-0.04	0.01	0.21^{***}	0.27^{***}	10
5	7.37	**0.04**	**0.07**	1.003	-12.23	-0.1	0.04	0.19^{***}	0.26^{***}	11

The SIC is minimized for lag=3, for which we also find a significant threshold and hence significant transaction costs for a negative basis trade of around 10 index points. The results for the threshold search for θ^2 are shown in Table 21. No results are presented for lag=1 as the 1-threshold TVECM is a better model fit compared to the 2-threshold TVECM model.

Table 22: Threshold search for θ^2 - S&P 500

The table shows the results for the threshold search for θ^2 for the period starting in 2000. We do not show the result for one lag, as a 1-threshold TVECM is the best model fit. The minimum value of the SIC is found at three VAR lags, for which the threshold is also significant. The transaction costs for a positive basis trade are around 5 index points. The smallest SIC value and the percentage of the bootstrap methods below 0.10 are boldfaced.

lag	SIC	Bootstrap fixed	Bootstrap residual	$\hat{\beta}_1$	$\hat{\beta}_0 + \hat{\theta}^2$	λ^2_{1L}	λ^2_{2L}	λ^2_{1U}	λ^2_{2U}	% upper regime
2	7.44	**0.09**	0.11	0.998	9.30	0.05	0.10*	0.06	0.15	14
3	**7.39**	**0.08**	**0.12**	1.001	5.46	0.04	0.08*	0.11	0.23	11
4	7.41	0.65	0.68	1.003	2.27	0.3	0.07	0.28	0.41	11
5	7.39	0.83	0.85	1.003	2.22	0.04	0.08*	0.21	0.29	11

Similar to the threshold search for θ^1, we find again a local minimum of the SIC at 3 lags, for which the threshold is significant. This leads to transaction costs of a positive basis trade of around 5 index points. Further, we show price discovery measures for the different regimes at three lags in Table 23.

Table 23: Measures of price discovery - period 2000 until December 2016 - S&P 500

The table shows price discovery measures for lag=3. The GG measures with the superscript $^+$ should be interpreted as 1. The price discovery measures of the neutral regime are means of the estimation from the threshold search for θ^1 and θ^2.

regime	HAS	GG	leading market	half-life (in 30min)	comments
upper	0.54	2.00^+	spot	5.41	
neutral	0.35	2.88^+	-	14.05	
lower	0.51	0.77	spot	13.50	

The spot market is leading in the two arbitrage regimes and as expected, the half-life in the neutral regime is the longest.

4 Conclusion

The identification of thresholds in nonlinear vector correction models is rather of complex origin with several unresolved problems. In this paper we present the solution for one unresolved issue which is the estimation of a 3-regime TVECM with an unknown cointegrating vector. Our proposed methodology extends the 2-regime TVECM model as proposed by Hansen and Seo (2002). In contrast to Hansen and Seo (2002) we also introduce an intercept β_0 in the cointegrating relation $(S - \beta_1 D - \beta_0)_t$ to account for a distorted parity relationship, eg the non-zero basis. Using a sequential grid search for the first and the second threshold, we estimate the cointegration relationship as well as the two thresholds (θ^1 and θ^2) by employing a maximum likelihood approach. As there are practically no empirical studies for TVECMs even for the 1-threshold case we present a comprehensive simulation study to understand the reliability of our proposed method. Hansen and Seo (2002) have suggested a grid point search to estimate the variables β_i and θ^i with a fixed number of grid points. We show that a dynamic grid point setting, where the distance between two grid points is set via a precision parameter, instead of defining a fixed number of grid points leads the best results, however at potentially high computational costs. The intercept β_1 in the cointegrating relationship is estimated at very high precision. We show that the estimator $\hat{\beta}_1$ found in the first threshold search can be fixed in the second search. This lowers the dimension of the grid space and hence reduces compute time in second grid search. The thresholds θ^i and the shift in the cointegration relation β_0 are estimated poorly, however the sum $\theta^i + \beta_0$ is estimated with a very good precision.

Our proposed methodology is particularly appealing for the analysis of distorted parity relationships in economics, such as no-arbitrage relationships with a non-zero basis. A persistent non-zero basis between two similar financial market instruments traded in the spot and in the derivative market points towards the presence of transaction costs on arbitrage trades that prevent a complete adjustment of market prices to the theoretical no-arbitrage condition of a zero basis.

We have presented four examples of basis trading: gold and platinum as well as DAX and S&P 500, to illustrate the method.

A Unit root tests

We present results of the Augmented Dickey-Fuller (ADF) test and the Phillips-Perron (PP) test for all our time series for the period beginning of 2000 until December 2016. The null hypothesis of both tests state: the series has a unit root. Table A.1 shows that our data is I(1) as required by the (T)VECM approach.

Table A.1: Unit root test

The table shows results of the ADF and PP test for period 2000 until December 2016 for our data in levels and in first differences.

	levels futures		levels spot		first differences futures		first differences spot	
	ADF	PP	ADF	PP	ADF	PP	ADF	PP
gold	0.63	0.63	0.75	0.75	0.00	0.00	0.00	0.00
platinum	0.20	0.19	0.33	0.30	0.00	0.00	0.00	0.00
DAX	0.78	0.83	0.82	0.85	0.00	0.00	0.00	0.00
S&P 500	0.90	0.95	0.90	0.95	0.00	0.00	0.00	0.00

B Cointegration tests

The VECM methodology requires that the spot and the futures time series must be cointegrated. We use the Johansen test with intercept but no deterministic trend in the cointegration equation. The results presented in Table B.1 indicate that we have cointegration for all our time series.

Table B.1: Johansen test

The table shows results of the Johansen test for period 2000 until December 2016.

	Trace test		Maximum eigenvalue test	
	None	at most 1	None	at most 1
gold	0.00	0.66	0.00	0.66
platinum	0.00	0.31	0.00	0.31
DAX	0.00	0.96	0.00	0.96
S&P 500	0.00	0.93	0.00	0.93

Part IIb: Arbitrage costs and the persistent non-zero CDS-bond basis: Evidence from intraday euro area sovereign debt markets[1]

1 Introduction

The theoretical no-arbitrage condition between credit default swaps (CDS) and credit-risky bonds based on Duffie (1999) is a cornerstone for empirical research on price discovery in credit risk markets. This condition requires that CDS spreads and (par floating rate) spreads on bonds issued by the entity referenced in the CDS contract must be equal, as any discrepancy would present investors with an arbitrage opportunity. For this no-arbitrage condition to hold, markets must be perfect and frictionless. In practice, however, frictions and imperfections often make such arbitrage trades difficult and costly to varying degree. These imperfections include limited and time-varying liquidity across market segments, unavailability of instruments with identical maturity and payout structures, and the fact that some arbitrage trades require tying up large amounts of capital for extended periods of time. As a result, the difference between the CDS premium and the bond spread, the so-called basis, is typically not zero. Moreover, the basis can become sizeable and persistent in times of market stress. This was particularly evident during the euro area sovereign debt crisis, when the basis widened significantly (see for example Fontana and Scheicher (2016) and Gyntelberg et al. (2013)). This paper adds to the existing literature by analysing the importance of arbitrage trading and arbitrage costs with respect to the size of the CDS-bond basis.

A persistent non-zero CDS-bond basis is likely to reflect the unwillingness of arbitrageurs to try to exploit it, unless the pricing mismatch is greater than the overall transaction costs of undertaking the arbitrage trade. Empirically, we would therefore expect to see such arbitrage forces intensifying as the magnitude of the basis exceeds some level that reflects the overall, average transaction costs for implementing the arbitrage trade. This suggests that the adjustment process towards the long-run equilibrium is nonlinear, in that it differs depending on the level of the basis. In order to capture such behaviour, we extend the vector error correction model (VECM) which has been the convention in existing studies (see for example Blanco et al. (2005) and Zhu (2004) for corporates, Ammer and Cai (2007) for emerging markets, and Fontana and Scheicher (2016), Gyntelberg et al. (2013), Mayordomo et al. (2011) and Palladini and Portes (2011) for euro area

[1] This chapter is joint work with Jacob Gyntelberg (Danske Bank, Copenhagen), Peter Hördahl (Bank for International Settlements, Hong Kong), Kristyna Ters (University of Basel, Basel) and has been published as: Gyntelberg, J., Hördahl, J., Ters, K. and Urban J., Arbitrage costs and the persistent non-zero CDS-bond basis: Evidence from intraday euro area sovereign debt markets. *BIS Working Papers No 631*, Bank for International Settlements.

sovereigns) to a nonlinear set-up using a threshold VECM (TVECM). This framework will help to answer the question if transaction costs on arbitrage trades were related to the widening of the CDS-bond basis during the sovereign debt crisis period. As it is impossible to disentangle the exact transaction costs for arbitrage trades in sovereign credit risk, our estimated transaction costs comprise overall costs that arbitrageurs face when implementing these trading strategies such as liquidity costs, funding cost, repo costs, risk compensation, search costs, cost associated with committing balance sheet space, etc..

One of the key contributions of our paper to the existing literature on price discovery in credit markets is that, in contrast to all studies mentioned above, we allow for a non-linear adjustment of prices in CDS and bond markets towards the long-run equilibrium. This allows us to determine whether a relationship exists between the overall costs that arbitrageurs face in the market for sovereign risk and the magnitude of the CDS-bond basis. Hence, with this model we can capture the possibility that arbitrageurs step into the market only when the trading opportunity is sufficiently profitable. Our TVECM approach can directly quantify the threshold beyond which such trading opportunities are seen by investors as 'sufficiently profitable'. Furthermore, our results show that even when markets in times of stress are liquid, the basis can widen as high market volatility makes arbitrage trades riskier, leading arbitrageurs to demand higher compensation (suggesting a higher threshold) before stepping into the market. This could explain why the basis reached very high levels during the euro area sovereign debt crisis as it was subject to considerable volatility in a stressed market environment.

Our analysis relies on intraday price data for both CDS and bonds, allowing us to estimate the spread dynamics and the price discovery implications substantially more accurately than existing studies that rely on lower frequency data. Our TVECM approach identifies thresholds in the CDS-bond basis, below which arbitrageurs are reluctant to step in. We also find that once the basis exceeds the estimated transaction costs (given by the threshold and the constant long-run mean of the basis), the adjustment speeds towards the long run equilibrium intensify. This supports our assumption that arbitrageurs only step into the market when the trade becomes profitable. We find that the estimated average transaction cost is around 80 basis points in the pre-crisis period. During the euro area sovereign debt crisis this average increased to around 190 basis points. This increase in the estimated threshold during the crisis period coincided with a higher CDS-bond basis volatility. As arbitrageurs face the risk that the arbitrage trade will go in the wrong direction in the short run, they will demand higher compensation for undertaking the arbitrage trade in volatile markets. Thus, our findings help to explain the persistent non-zero basis in markets for sovereign credit risk.

The remainder of the paper is structured as follows. Section 2 discusses in more detail the relationship between sovereign CDS and bonds. Section 3 explains our data, while

Section 4 discusses the set-up and estimation of our TVECM. Section 5 provides the empirical results and Section 6 concludes.

2 Relation between sovereign CDS and bonds

The importance of frictions in credit risk modelling is well-known. However, only few empirical studies analyse the effects of frictions on the price discovery process for credit risk. Several papers conclude that for example liquidity affects corporate bond spreads significantly (eg Chen et al. (2007), Ericsson and Renault (2006), Elton et al. (2001) and Mahanti et al. (2008)). By contrast, other papers argue that CDS spreads reflect pure credit risk, ie that they are not significantly affected by liquidity (eg Longstaff et al. (2005)). However, there are numerous papers reporting that CDS spreads are too high to represent pure credit risk (eg Berndt et al. (2005), Blanco et al. (2005), Pan and Singleton (2005)). Tang and Yan (2007) find that the level of liquidity and liquidity risk are important factors in determining CDS spreads. Hull and White (2000) address the effects of market frictions from a theoretical point of view and determine conditions under which CDS prices are affected. Longstaff et al. (2005) study price differences between CDS and bonds and attribute them to liquidity and counterparty risk. Also Zhu (2004) concludes that liquidity matters in CDS price discovery. Ammer and Cai (2007), Levy (2009) and Mayordomo et al. (2011) find evidence that liquidity (as measured by the bid-ask spread) is a key determinant for price discovery, but without explicitly modelling any market frictions. Tang and Yan (2007) focus on pricing effects in CDS and show that the liquidity effects on CDS premia are comparable to those on treasury and corporate bonds (Tang and Yan; 2007).

2.1 Frictionless markets

In a frictionless market, the CDS premium should equal the spread on a par fixed-rate bond (issued by the same entity as referenced by the CDS) over the riskfree interest rate (Duffie (1999)). Both the CDS premium and the risky bond's yield spread is compensation to investors for being exposed to default risk, and must therefore be priced equally in the two markets. However, for this no-arbitrage relationship to hold exactly, a number of specific conditions must be met, including that markets are perfect and frictionless, that bonds can be shorted without restrictions or cost, that there are no tax effects, etc. Any departures from this perfect environment will introduce potential wedges between the pricing of credit risk in CDS contracts and in bonds.

Moreover, given that floating rate notes are relatively uncommon, in particular for sovereigns, any comparison between CDS spreads and bond spreads based on fixed-rate bonds will introduce other distortions. Hence, the observed difference between the CDS premium and the bond spread, the basis, is typically not zero.

2.2 Markets with frictions

There are a number of recent papers that focus on the pricing of sovereign credit risk in the euro area, which all find that the theoretical no-arbitrage condition between CDS spreads and bond spreads does not hold (for example Fontana and Scheicher (2016), Gyntelberg et al. (2013), Arce et al. (2012), and Palladini and Portes (2011)). Gyntelberg et al. (2013) find that the basis across seven euro area sovereign entities[2] is almost always positive over the 2008-11 sample period for the 5 year and the 10 year tenor. Moreover, they find that the basis varies substantially across countries, with means ranging from 74 to 122 basis points for the 5-year tenor, and from 58 to 175 basis points for the 10-year tenor. Empirical research on corporate credit risk also points towards a non-zero basis as shown for example in Nashikkar et al. (2011), Blanco et al. (2005) and Zhu (2004), and for emerging markets sovereign credit risk according to Ammer and Cai (2007).

The CDS market is a search market as the contracts are traded over-the-counter (OTC) where parties have to search for each other in order to bargain and match a trade. Therefore, market trading is not continuous in the sense that it is not necessarily possible to buy or sell any amount immediately (Black; 1971). Moreover, other frictions and imperfections may make arbitrage trades difficult and costly. These imperfections include limited and time-varying liquidity in some or all market segments, unavailability of instruments with identical maturity and payout structures, and the fact that some arbitrage trades require tying up large amounts of capital for extended periods of time. As the costs associated with tying up space on banks' balance sheets have risen following the global financial crisis, this can represent a significant hurdle that traders face in the market. Furthermore, the no-arbitrage condition relies on the ability to short sell bonds, which is not always costless and sometimes even impossible due to illiquid markets. All of these imperfections contribute to explaining why the basis between CDS and bond spreads can deviate from zero, often substantially and persistently. However, we would expect to see arbitrage forces come into play if the basis becomes "too wide", thereby pushing it back towards zero. Clearly, we would also expect to see stronger adjustment forces in CDS and bond markets when the basis exceeds some critical threshold. The size of the threshold would reflect the various arbitrage costs traders face in markets, including costs for illiquidity as well as for tying up costly capital for possibly long periods of time.

[2] France, Germany, Greece, Ireland, Italy, Portugal, Spain; 5- and 10-year tenor from October 2008 to end-May 2011

3 Data

For our empirical analysis we use intraday price quotes for CDS contracts and government bonds for France, Germany, Greece, Ireland, Italy, Portugal and Spain. We choose this group of countries because they include those that were most affected by the euro sovereign debt crisis. Germany is included as a near-riskfree reference country, and France which we consider as a low-risk control country. We use 5- and 10-year USD-denominated CDS quotes for all countries in our sample. As documented in Gyntelberg et al. (2013), the 5-year segment is more liquid than the 10-year segment, particularly as the sovereign debt crisis intensified.

Our sovereign bond price data is provided by MTS (Mercato Telematico dei Titoli di Stato). The MTS data consists of both actual transaction prices and binding bid-offer quotes. The number of transactions of sovereign bonds on the MTS platform is however not sufficient to allow us to undertake any meaningful intraday analysis. Therefore, we use the trading book from the respective domestic MTS markets.[3]

The CDS data consists of price quotes provided by CMA (Credit Market Analysis Ltd.) Datavision. CMA continuously gathers information on executable and indicative CDS prices directly from the largest and most active credit investors. After cleaning and checking the individual quotes, CMA applies a time and liquidity weighted aggregation so that each reported bid and offer price is based on the most recent and liquid quotes.[4]

We construct our intraday data on a 30-minute sampling frequency for the available data sets that span from January 2008 to end-December 2011. The available number of indicative quotes for CDS does not allow higher data frequency than 30 minutes. The euro area sovereign CDS markets were very thin prior to 2008, which makes any type of intraday analysis before 2008 impossible (for a discussion please refer to Gyntelberg et al. (2013)).

When implementing our analysis we split the data into two sub-samples. The first sub-sample covers the period January 2008 to end-March 2010, and as such represents the period prior to the euro area sovereign debt crisis (van Rixtel and Gasperini; 2013). While this period includes the most severe phase of the financial crisis, including the default of Lehman Brothers, it is relatively unaffected by any major market concerns about the sustainability of public finances in euro area countries. The second sub-sample covers the

---

[3] We ignore quotes from the centralized European platform (market code: EBM), as quotes for government bonds on the centralised platform are duplicates of quotes on the domestic platforms. The MTS market is open from 8:15 to 17:30 local Milan time, preceded by a pre-market phase (7.30 to 8.00) and an offer-market phase (8:00 to 8:15). We use data from 8:30 to 17:30.

[4] The CDS market, which is an OTC market, is open 24 hours a day. However, most of the activity in the CMA database is concentrated between around 7:00 and 17:00 London time. As we want to match the CDS data with the bond market data, we restrict our attention to the period from 8:30 to 17:30 local Milan time.

euro area sovereign debt crisis period and runs from April 2010 to December 2011. We have tested other break downs in a pre-crisis and crisis period[5] and have found that our results remain robust.

In order to accurately match the maturities and the cash flow structures of the CDS and the cash components for the measurement of the CDS-bond basis, we calculate intraday asset swap (ASW) spreads based on estimated zero-coupon government bond prices according to Nelson and Siegel (1987). Appendix A provides details. The use of ASW spreads is also in line with the practice applied in commercial banks when trading the CDS-bond basis. By calculating ASW spreads we ensure that we are comparing like with like in our empirical analysis, and we avoid introducing distortions by using imperfect cash spread measures, such as simple "constant maturity" yield differences.

An asset swap is a financial instrument that exchanges the cash flows from a given security - eg a particular government bond - for a floating market rate[6]. This floating rate is typically a reference rate such as Euribor for a given maturity plus a fixed spread, the ASW spread. This spread is determined such that the net value of the transaction is zero at inception. The ASW allows the investor to maintain the original credit exposure to the fixed rate bond without being exposed to interest rate risk. Hence, the ASW is similar to the floating-rate spread that theoretically should be equivalent to a corresponding CDS spread on the same reference entity.

Finally, we note that using intraday data in our empirical analysis should enable us to obtain much sharper estimates and clearer results with respect to market mechanisms and price discovery compared to any analysis carried out with a lower data frequency (see Gyntelberg et al. (2013)).

Using the above methodology, we derive the intraday asset swap spreads for each country for the 5- and 10-year maturities (displayed in Appendix B). The corresponding CDS series are also shown in Appendix B while the CDS-bond basis is displayed in Figures 2 and 3.

In Appendix C we present information on liquidity such as number of ticks, number of trades and bid-ask spreads for CDS and bonds. Interestingly, we find that for example the number of data ticks for our sovereign bonds remained quite stable over the whole sample period and that the 5-year tenor is typically more liquid than the 10-year tenor. The number of indicative CDS prices (see Figure C.1) remained stable for the 5-year tenor (Greece is an exception) and decreased for the 10-year tenor. The number of trades

[5] We have for example tested the 20 October 2009 as the beginning of the crisis period. At that date the new Greek government announced that official statistics on Greek debt had previously been fabricated. Instead of a public deficit estimated at 6% of GDP for 2009, the government now expected a figure at least twice as high.

[6] See Appendix A. Gyntelberg et al. (2013) and O'Kane (2000) further discuss the mechanics and pricing of asset swaps.

reported in the EuroMTS platform decreased slightly for most GIIPS countries since the onset of the euro area sovereign debt crisis (see Figure C.3). On the other hand, the sovereign CDS data shows that the number of ticks more than doubled in 2010, as the crisis spread. The bid-ask spreads for our sovereign CDS and bonds tighten over our sample period in France and Germany. While CDS bid-ask spreads in GIIPS countries are typically very tight, the spread size is quite volatile for bonds. While we can see that the bid-ask spreads for the Irish, Italian, Portuguese and Spanish 5-year bonds widen during the sovereign debt crisis period, we can not see the same behaviour in the 10-year bond segment (Figure C.6).

Thus, the dramatic increase of the CDS-bond basis during the euro area sovereign debt crisis can not be exclusively explained by market liquidity, but seems to be linked to overall transaction cost in these markets.

4 Threshold vector error correction model (TVECM)

We begin our empirical analysis by examining the statistical properties of our spread time series. This analysis shows that the series are I(1) and that the CDS and ASW series are cointegrated (see Appendix D and E). As a result, we can employ a vector error correction model (VECM) to study the joint price formation process in both markets. From the estimated error correction model we calculate measures that indicate which of the two markets is leading the price discovery process as well as examine the speed of adjustment towards the long-term equilibrium.

The linear VECM concept implies that any deviation from the long-run equilibrium of CDS and ASW spreads will give rise to dynamics that will bring the basis back to the equilibrium due to an error correction mechanism as illustrated in the left panel of Figure 1. Thus, in a market with no frictions (such as transaction costs) every deviation from the non-zero basis will initiate arbitrage trades on the pricing differential between the spot and the derivatives market (Figure 1). Hence, in a frictionless market, the basis will typically fluctuate around zero.

Given that the CDS and bond markets are subject to market frictions and arbitrageurs face various trading costs, it is useful to extend the linear VECM approach[7] to a threshold vector error correction model (TVECM). Threshold cointegration was introduced by Balke and Fomby (1997) as a feasible mean to combine regime switches and cointegration. The TVECM model allows for nonlinear adjustments to the long-term equilibrium in CDS and bond markets. In our case, such nonlinear adjustment dynamics should be able to capture arbitrageurs' decisions to only step into the market when the basis exceeds some critical threshold, such that the expected profit exceeds the transaction costs. As a

[7] As in eg Fontana and Scheicher (2016), Gyntelberg et al. (2013) and Blanco et al. (2005).

result, adjustments to the long-term equilibrium would then be regime-dependent, with a relatively weak adjustment mechanism below the threshold (a 'neutral' regime) and a stronger adjustment mechanism above it. This is illustrated in the right panel of Figure 1. The example in this figure displays a predominantly positive basis as this is also the case in our underlying data (see Figures 2 and 3).

Figure 1: Linear versus Threshold Vector Error Correction Model

The linear VECM model in the left panel represents markets where the theoretical no-arbitrage condition holds approximately as the basis does not deviate too much from zero. Otherwise arbitrageurs step in immediately to trade on pricing differentials between the spot and the derivatives market which reverts the basis back towards zero. The right-hand panel shows the case for markets that are subject to non-negligible transaction costs. Arbitrageurs will only step in once the expected gain from the trade is above the transaction costs, in the "arbitrage regime". A predominantly positive basis is shown in this example as this reflects the typical conditions in euro sovereign debt markets.

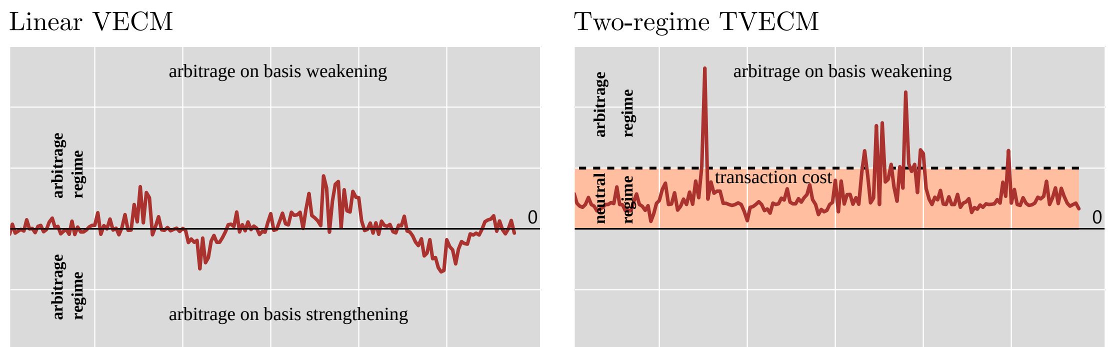

4.1 Model specification

Let $y_t = (CDS_t \quad ASW_t)^\mathsf{T}$ represent the vector of CDS and ASW spreads at time t for a specific sovereign entity. The TVECM approach allows the behaviour of y_t to depend on the state of the system. In our data, the basis for all reference entities is almost always positive. Hence, we expect to find at most two regimes with one threshold θ, above which arbitrageurs can be expected to step in to trade on the pricing difference in the two markets, but below which they will have little or no incentive to do so. One can formulate a two-regime TVECM as follows[8]:

$$\Delta CDS_t = \left[\lambda_1^L \mathrm{ec_{t-1}} + \Gamma_1^\mathrm{L}(\ell)\Delta \mathrm{y_t}\right] d_{Lt}(\beta,\theta) + \left[\lambda_1^U \mathrm{ec_{t-1}} + \Gamma_1^\mathrm{U}(\ell)\Delta \mathrm{y_t}\right] d_{Ut}(\beta,\theta) + \varepsilon_t^{CDS},$$
$$\Delta ASW_t = \left[\lambda_2^L \mathrm{ec_{t-1}} + \Gamma_2^\mathrm{L}(\ell)\Delta \mathrm{y_t}\right] d_{Lt}(\beta,\theta) + \left[\lambda_2^U \mathrm{ec_{t-1}} + \Gamma_2^\mathrm{U}(\ell)\Delta \mathrm{y_t}\right] d_{Ut}(\beta,\theta) + \varepsilon_t^{ASW}$$

[8] for a derivation of the TVECM see for example Balke and Fomby (1997)

or in vector form,

$$\Delta y_t = \left[\lambda^L \mathrm{ec_{t-1}} + \Gamma^{\mathrm{L}}(\ell)\Delta \mathrm{y_t}\right] d_{Lt}(\beta,\theta) + \left[\lambda^U \mathrm{ec_{t-1}} + \Gamma^{\mathrm{U}}(\ell)\Delta \mathrm{y_t}\right] d_{Ut}(\beta,\theta) + \varepsilon_t \qquad (1)$$

where $\mathrm{ec}_{t-1} = (CDS_{t-1} - \beta_0 - \beta_1 ASW_{t-1})$ is the error correction term, $\Gamma^j(\ell)\Delta y_t$, $j \in \{L, U\}$ represents the VAR term of some order, expressed in lag operator (ℓ) representation, and $\varepsilon_t = (\varepsilon_t^{CDS} \quad \varepsilon_t^{ASW})^{\mathsf{T}}$ is a vector of i.i.d. shocks. The lower regime (specified by the index L) is defined as $\mathrm{ec}_{t-1} \leq \theta$, and the upper regime (specified by the index U) as $\mathrm{ec}_{t-1} > \theta$. Hence d_{Lt} and d_{Ut} are defined using the indicator functions $I(\cdot)$ as follows:

$$d_{Lt}(\beta,\theta) = I(\mathrm{ec}_{t-1} \leq \theta),$$
$$d_{Ut}(\beta,\theta) = I(\mathrm{ec}_{t-1} > \theta).$$

The error correction term ec_{t-1} represents the long-term equilibrium of the two time series which has to be stationary by construction (Johansen; 1988). The number of lags in the VAR terms are determined using the Schwarz information criterion. We constrain β_1 to 1 which is motivated by our no-arbitrage discussion in Section 2. A non-zero estimated β_0 represents a persistent non-zero basis. The average transaction costs that arbitrageurs need to overcome, as implied by the model, can now be identified as $\theta + \beta_0$.

The speed of adjustment parameters λ^U and λ^L characterize to what extent the price changes in $\Delta y_t = (\Delta CDS_t \quad \Delta ASW_t)^{\mathsf{T}}$ react to deviations from the long-term equilibrium. In case price discovery takes place only in the bond market we would find a negative and statistically significant λ_1^j and a statistically insignificant λ_2^j, as the CDS market would adjust to correct the pricing differentials from the long-term relationship. In other words, in this case the bond market would move ahead of the CDS market as relevant information reaches investors. Conversely, if λ_1^j is not statistically significant but λ_2^j is positive and statistically significant, the price discovery process takes place in the CDS market only - that is, the CDS market moves ahead of the bond market. In cases where both λ's are significant, with λ_1^j negative and λ_2^j positive, price discovery takes place in both markets.

We expect to find the speed of adjustment parameters to indicate that arbitrageurs are engaging in CDS-ASW basis trades if the basis exceeds the average transaction costs of $(\theta + \beta_0)$. In a market with a positive basis (CDS > ASW), arbitrageurs will bet on a declining basis and will therefore short credit risk in the bond market and go long credit risk in the CDS market, ie sell the bond and sell the CDS (Gyntelberg et al.; 2013).[9] The

[9] In case of a negative basis (ASW > CDS), arbitrageurs bet on an increasing basis while carrying out the reverse trade. In markets where the basis regularly would fluctuate between being positive and negative, we would expect to find a 3-regime TVECM. With a lower regime $\mathrm{ec}_{t-1} \leq \theta^1$, a middle regime (neutral regime) $\theta^1 < \mathrm{ec}_{t-1} \leq \theta^2$, and a upper regime $\theta^2 < \mathrm{ec}_{t-1}$.

predominantly positive basis throughout our sample suggests the presence of at most one threshold.

Moreover, we expect to find higher transaction costs ($\theta + \beta_0$) in times of market stress. This can be explained by the fact that when the basis is subject to increased volatility, the risk increases that any arbitrage trade moves in the wrong direction in the short or medium term. Therefore, arbitrageurs will demand higher compensation for taking such positions in times when the basis volatility is high, resulting in higher estimated thresholds.

4.2 Estimating the threshold

As discussed above, the positive basis in our sample suggests the presence of at most one threshold. In order to test for the presence of a threshold effect, we follow the method proposed by Hansen and Seo (2002) who extend the literature by examining the case of an unknown cointegrating vector.[10] They implement maximum likelihood estimation (MLE) of a bivariate TVECM with two regimes. Their algorithm involves a joint grid search over the threshold and the cointegrating vector while using the error-correction term as the threshold variable (see Equation (1)). All coefficients are allowed to switch between these two regimes. Only the cointegrating vector β remains fixed across all regimes, by construction. We follow this grid search estimation approach, subject to the constraint $\beta_1 = 1$, motivated by our no-arbitrage discussion in Section 2.

As in Hansen and Seo (2002) we estimate the model while imposing the following additional constraint:

$$\pi_0 \leq P(\mathrm{ec}_{t-1} \leq \theta) \leq 1 - \pi_0 \tag{2}$$

where $\pi_0 > 0$ is a trimming parameter and P is the share of observations in each regime. This constraint allows us to identify a threshold effect only if the share of observations in each regime is greater than π_0. If this condition is not met, the model reduces to a linear VECM. Andrews (1993) argues that setting π_0 between 0.05 and 0.15 are typically good choices. As we use intraday data of the order of 10,000 observations, we set the trimming parameter to $\pi_0 = 0.10$, which will still ensure an adequate number of observations in both regimes.

4.3 Statistical testing for a threshold

Once a threshold has been identified, the next step is to determine whether the estimated threshold θ is statistically significant. Under the null hypothesis $\mathscr{H}_0$ there is no threshold, so the model reduces to a conventional linear VECM where $\lambda^L = \lambda^U$. The two regime

[10] Balke and Fomby (1997) and Tsay (1989) transform the TVECM specification into a univariate regression while the cointegrating vector is known a priori.

TVECM is the alternative hypothesis $\mathscr{H}_1$ with $\lambda^L \neq \lambda^U$ under the constraint in Equation (2). The linear VECM under $\mathscr{H}_0$ is nested in Equation (1), hence, a regular LM test with an asymptotic $\chi^2(N)$-distribution can be calculated based on Equation (1). However, the LM test can only be applied if the cointegrating vector β and the threshold variable θ are known a priori (Hansen and Seo; 2002). While the point estimate of β under $\mathscr{H}_0$ is $\hat{\beta}$ from the linear model, there is no estimate of θ under $\mathscr{H}_0$. This implies that there is no distribution theory for the parameter estimates and no conventionally defined LM statistic.

As there is no formal distribution theory under the $\mathscr{H}_0$ we follow Hansen and Seo (2002) and perform two different bootstrap analyses in order to estimate the distribution for our model specification in Equation (1). First, we implement a non-parametric bootstrap on the residuals, called the "fixed regressor bootstrap", which resamples (Monte-Carlo) the residuals from the estimated linear VECM. The second bootstrap methodology is parametric, called "residual bootstrap". It is assumed that the residuals are i.i.d. Gaussian from an unknown distribution with fixed initial conditions. The parametric bootstrap then calculates the sampling distribution of the supremum LM test in Equation (3) below using the parameter estimates obtained under the $\mathscr{H}_0$. The distribution is bootstrapped using Monte-Carlo simulations from the residual vector under the $\mathscr{H}_0$ while the vector series y_t are created by recursion given the linear VECM model.

For the critical value, we employ a supremum LM statistic based on the union-intersection principle, proposed by Davies (1987):

$$\text{SupLM} = \sup_{\theta_L \leq \theta \leq \theta_U} LM(\hat{\beta}, \theta). \tag{3}$$

According to the constraint in Equation (2) we set the search region $[\theta_L, \theta_U]$ such that θ_L is the π_0 percentile of $\hat{\text{ec}}_{t-1}$, and θ_U is the $(1-\pi_0)$ percentile. This grid evaluation over $[\theta_L, \theta_U]$ is necessary to implement the maximisation defined in Equation (3) because the function $\text{LM}(\hat{\beta}, \theta)$ is non-differentiable in θ.

We consider our model as threshold cointegrated if we can reject the null hypothesis of a linear VECM by either the "residual bootstrap" or the "fixed regressor bootstrap" methodology. We verify that our results are robust with respect to the choice of the trimming parameter.

4.4 Measure of price discovery

We calculate the Hasbrouck (1995) measure to investigate in which market segment – the CDS market or the bond market – price discovery takes place. The Hasbrouck measure is calculated based on the estimated speed of adjustment parameters λ^U and λ^L as well as the estimated covariance matrix of the error terms, and is by construction confined to the closed interval [0,1]. This makes interpretation straightforward. We specify our Hasbrouck measures such that HAS > 0.5 can be interpreted as the CDS market contributing more to price discovery than the cash market. Similarly, HAS < 0.5 means that the bond (ASW) market contributes more to price discovery.[11]

Finally, we are interested in examining the speed of adjustment towards the long-term equilibrium in each regime. As the CDS and ASW spreads in the bivariate VECM share a common stochastic trend, the speed of adjustments of the cointegrating residual to the long-run equilibrium can be used to determine the impulse response function (Zivot and Wang; 2006). The vector error correction mechanism directly links the speed of adjustment of CDS and ASW spreads to the regime dependent cointegrating error u_t^j which follows an implied AR(1) process:

$$\begin{aligned} u_t^j &= (1+\lambda_1^j-\beta_1\lambda_2^j)u_{t-1}^j+\varepsilon_t^{CDS}-\beta_1\varepsilon_t^{ASW} \\ &= (1+\lambda_1^j-\lambda_2^j)u_{t-1}^j+\varepsilon_t^{CDS}-\varepsilon_t^{ASW} \equiv \phi^j u_{t-1}^j+\varepsilon_t^{CDS}-\varepsilon_t^{ASW}\,, \end{aligned} \tag{4}$$

where we have set β_1 to 1 in the second line of the equation.[12] The superscript j stands for L and U. The half-life of a shock for each regime, hl^j, can now be calculated from the AR(1) coefficient ϕ^j as:

$$hl^j = \frac{ln(0.5)}{ln(\phi^j)}\,. \tag{5}$$

5 Results

In this section we first present results for the period before the euro area sovereign debt crisis (January 2008 to end-March 2010). These are followed by our findings using data for the sovereign debt crisis period (April 2010 to December 2011).

As a general result, we find a functioning relationship between the CDS market and the bond market during both samples. In cases where we find threshold cointegration, the adjustment process towards the long-term equilibrium is faster in the upper regime compared to the lower regime, in line with our reasoning on the behaviour of arbitrageurs. The estimated transaction costs in the pre-debt-crisis period average around 80 basis points. For the second sub-period (sovereign debt crisis) we find much higher thresholds of around 190 basis points. These estimated transaction costs, which are not directly observable, represent the overall costs that arbitrageurs face, such as liquidity costs, repo

[11] Specifically, we calculate the independent set of values HAS_1 and HAS_2 based on the CDS market for each regime, and we then define HAS as the average of HAS_1 and HAS_2.

[12] We include the intercept β_0 in our error correction term and set $\beta_1 = 1$, motivated by our no-arbitrage discussion in Section 2.

costs, search costs, cost associated with committing balance sheet space, as well as risk compensation, etc. The two to three times higher transaction costs during the crisis period are in line with our expectations, as markets were subject to stress in peripheral sovereign credit markets. The significant increase of the basis level during the sovereign debt crisis period can not be uniquely explained by illiquidity as already discussed in Section 3. For example, we also find an increased basis for sovereigns such as France where liquidity increased during the crisis period (as number of ticks, bid-ask spread, number of trades see Appendix C).

Instead, much of the increase in the thresholds during the crisis is likely related to arbitrageurs demanding higher compensation for undertaking arbitrage trades, as the risk of the trade moving in the wrong direction is elevated. In the short run, this risk is directly proportional to the basis volatility. By calculating a daily basis trade gain we show below that arbitrageurs demanded a higher compensation for elevated basis volatility while on a risk-adjusted level the overall compensation remained comparable to the pre-debt crisis period.

The estimated transaction costs ($\theta+\beta_0$) are displayed as red horizontal lines in Figures 2 and 3 in comparison to the overall basis level that is shown as blue curves. We find that the estimated overall transaction costs increased during the crisis period. This finding holds for the 5-year and the 10-year tenor.[13] Empirically we find moderate or no adjustment dynamics below the estimated transaction costs. Thus, we can say that in the lower regime (below the transaction costs $\theta + \beta_0$), the price dynamics are consistent with the notion that arbitrageurs have no incentive to carry out arbitrage trades. However, once the transaction costs ($\theta + \beta_0$) are exceeded (upper regime) and arbitrage trades become profitable, we find rapid adjustment dynamics. Thus, the increase in the basis and in the thresholds during the crisis period is consistent with an increase in overall transaction costs that arbitrageurs face in the market for sovereign risk.

[13] Except for Germany where we either find no significant threshold (10-year tenor) or no threshold at all (5-year tenor) for the sovereign debt crisis period.

Figure 2: CDS-ASW basis, 5 year tenor

The basis is the difference between the CDS spread and the ASW spread expressed in basis points for the period from January 2008 until December 2011. The figure shows data with 30-minute sampling frequency. Due to the Greek debt restructuring the data for Greece ends in September 2011. The red horizontal line represents the overall transaction costs ($\theta + \beta_0$) for the average arbitrageur. During the crisis period the linear VECM model for Germany (superscript $^+$) is a better model fit than any threshold model based on maximum likelihood estimation. Therefore, we do not plot the red horizontal line representing the overall transaction costs for the crisis period in Germany.

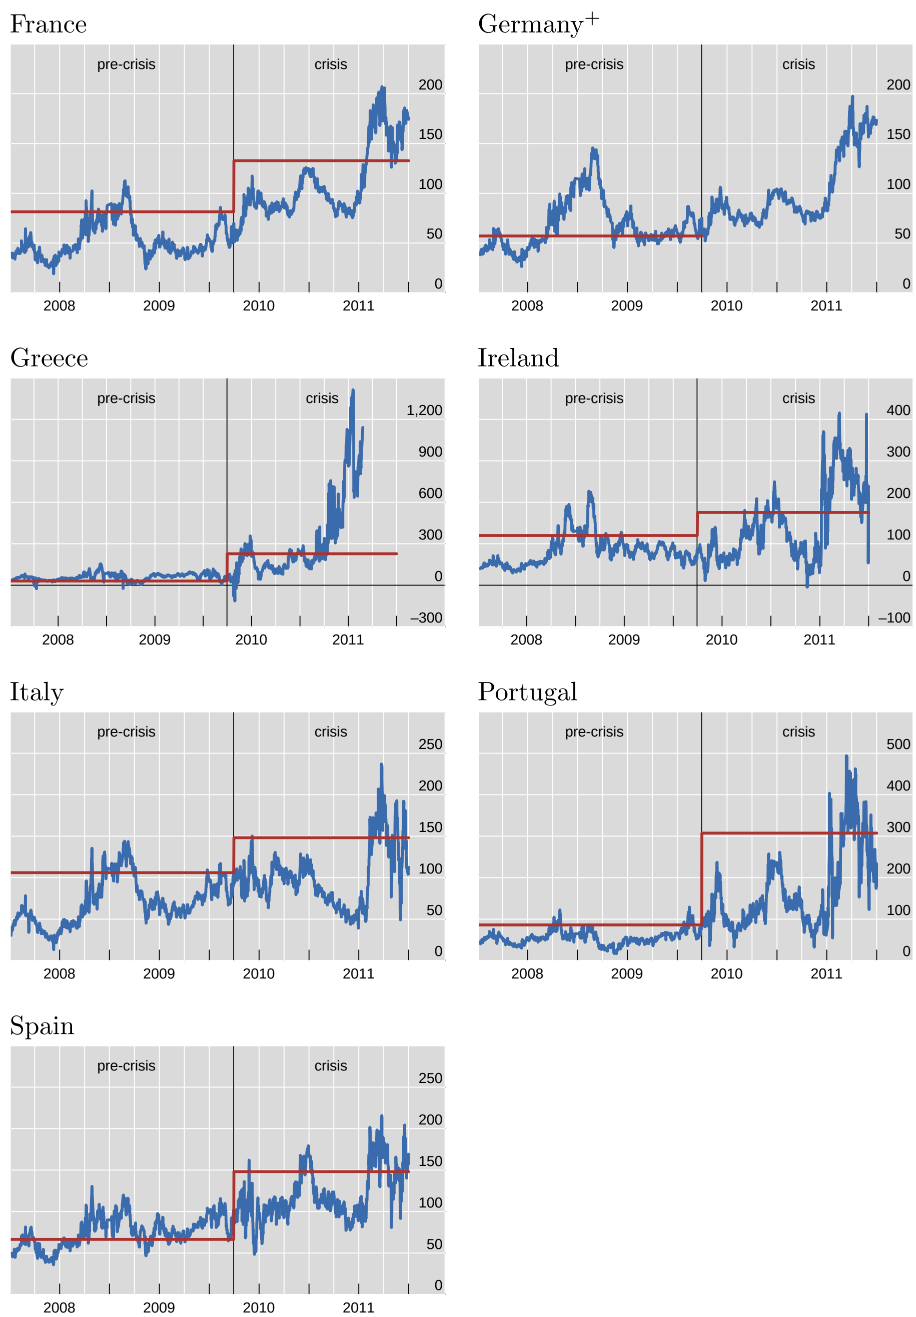

Figure 3: CDS-ASW basis, 10 year tenor

The basis is the difference between the CDS spread and the ASW spread expressed in basis points for the period from January 2008 until December 2011. The figure shows data with 30-minute sampling frequency. Due to the Greek debt restructuring the data for Greece ends in September 2011. The red horizontal line represents the overall transaction costs $(\theta + \beta_0)$ for the average arbitrageur.

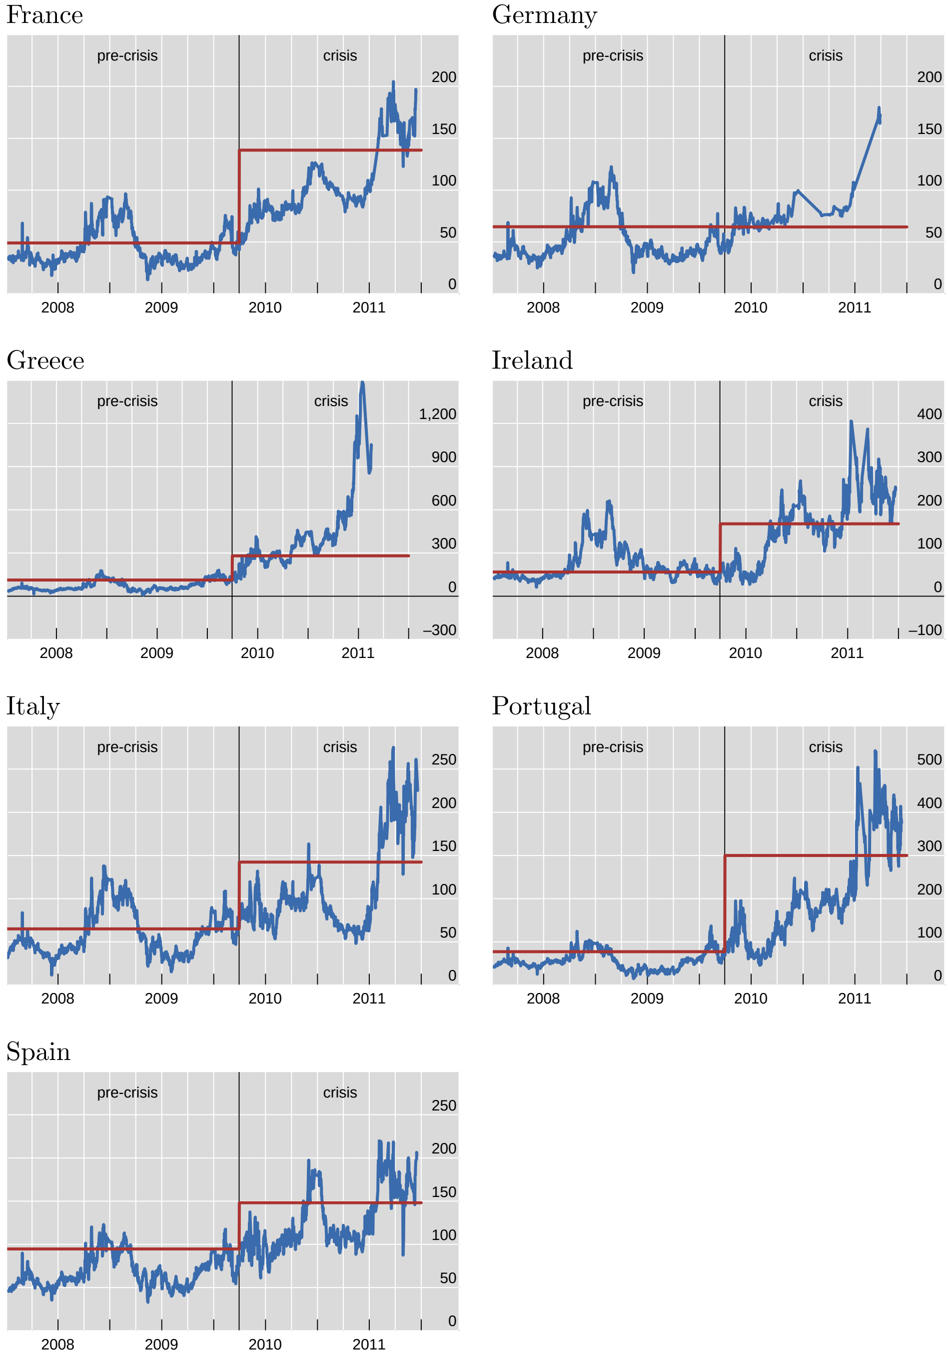

5.1 Results for the pre-debt-crisis period

The results for the first sub-sample from January 2008 to end-March 2010, ie prior to the euro area sovereign debt crisis, show that arbitrage trading intensifies in CDS and bond markets once some basis threshold is exceeded. In the lower (neutral) regime we find as expected either no adjustment dynamics, or speed of adjustments that are much smaller in magnitude than in the upper regime. The price discovery results for the 5-year and 10-year tenor are presented in Table 1. Countries in bold have a statistically significant threshold according to either the "fixed regressor bootstrap" or the "residual bootstrap" methodology, as well as speed of adjustments as expected by arbitrage theory. The sum $\theta + \beta_0$ represents the estimated transaction costs while the significance levels of the threshold significance test are represented by the superscript *, **, *** (90%, 95% and 99% CL). The column observations (obs.) denotes the share of observations in the lower regime as a percentage of the total number of observations.

For the 5-year tenor, we fail to find threshold effects for most countries. As expected we find more thresholds for the less liquid 10-year tenor in the pre-crisis period, because less liquid market segments have more frictions and higher arbitrage costs and are thus more likely to exhibit multi-regime behaviour.

The results are supportive of our hypothesis regarding arbitrageurs behaviour in markets with frictions. We find either faster adjustment dynamics towards the long-term equilibrium in the upper regime compared to the lower regime, or no adjustments in the lower regime (ie simple VAR dynamics). Table 2 shows that the half-lives of any basis widening are also either significantly shorter in the upper regime compared to the lower regime or undefined in the lower regime (the only exception is France). This suggests that arbitrage trading activity is much higher in the upper regime and therefore pricing differences due to credit risk shocks are reabsorbed much faster once the threshold is exceeded. Typically, the upper regime can be viewed as an extreme regime as the bulk of observations is in most cases concentrated in the lower (neutral) regime. This is due to the fact that if the basis moves into the upper regime, the actions of arbitrageurs will quickly move the basis back into the lower regime.

We also show the Hasbrouck (HAS) price discovery measure, which gives information on the relative price leadership of the respective markets (CDS versus bond). Here, the superscripts U and L denote the upper and lower regime, respectively. Overall, the price discovery results are mixed. Focusing on the countries in bold, in the upper regime, there appears to be a tendency for the CDS market to lead the bond market in the 10-year segment. In the 5-year segment, on the other hand, there is a (weak) tendency for the bond market to lead in the upper regime. The results for the lower regime are inconclusive across both maturities.

Table 1: Price discovery TVECM - pre-crisis period

This table reports the price discovery analysis for the period from January 2008 to end-March 2010. The values of the VECM coefficients λ are expressed in units of 10^{-4}. HAS is defined as the average of HAS_1 and HAS_2 (Hasbrouck; 1995). The transaction costs $\theta + \beta_0$ are presented in basis points. The superscripts U and L denote the upper and lower regime, respectively. The upper regime is above the overall transaction costs $\theta + \beta_0$ for the arbitrage trade and the lower regime is equal and below the transaction costs. The average of the transaction costs in the last line of each table takes only the significant thresholds into account. Boldfaced country names represent entities for which we have found a significant threshold and where at least one speed of adjustment in the upper regime is significant and has the correct sign to move the basis back to the long-run equilibrium.

Panel A - 5-year tenor

Sovereign	$\theta + \beta_0$	HAS^U	λ_1^U	λ_2^U	HAS^L	λ_1^L	λ_2^L	obs.
France	81.4	0.63	-14.93*	-21.60**	0.17	-1.81*	1.40	87.2%
Germany	57.0**	0.81	0.19	0.60	0.93	2.22	-7.46*	16.9%
Greece	30.1	0.94	-10.69	105.04***	0.67	-67.32*	254.66	12.6%
Ireland	120.1*	0.71	5.23	6.24	0.82	2.28	-5.66	87.6%
Italy	106.1**	0.06	-7.13	-1.48	0.01	6.29	1.01	83.17%
Portugal	86.0*	0.07	-54.32*	-1.36	0.73	13.25	30.95**	89.9 %
Spain	66.4*	0.24	-25.96*	13.77	0.19	-14.90	7.46	18.3%
average	87.1							

Panel B - 10-year tenor

Sovereign	$\theta + \beta_0$	HAS^U	λ_1^U	λ_2^U	HAS^L	λ_1^L	λ_2^L	obs.
France	49.1	0.90	13.45	79.29**	0.00	-7.22**	0.00	66.2%
Germany	64.7	0.92	0.54	3.87	0.55	1.92*	-2.01**	78.5%
Greece	113.0	0.72	-16.40	23.58**	0.05	20.97***	5.70	81.7%
Ireland	56.0*	0.93	-2.23	4.93**	0.66	4.18	-6.70	39.1%
Italy	65.1**	0.84	-2.92	6.79*	0.21	6.07*	-4.55	55.5%
Portugal	77.2**	0.75	-15.33	24.14**	0.06	14.39**	-4.15	81.1%
Spain	94.7**	0.01	-23.98**	-2.05	0.51	14.66**	7.74	90.0%
average	73.3							

Table 2: Half-life of shocks in days - pre-crisis period

This table reports the half-life of shocks of 5-year and 10-year CDS and ASW for the period from January 2008 to end-March 2010. The half-lives of shocks are expressed in days, and are calculated using the impulse response function to a one unit shock on the cointegrating error, using Equations (4) and (5). In case the speed of adjustment is of the wrong sign we do not report any half-life. "Lower" denotes results for the region below the threshold, and "upper" above it.

	5-year tenor		10-year tenor	
Sovereign	lower	upper	lower	upper
France	119.9	-	53.3	5.8
Germany	-	939.2	-	115.6
Greece	1.2	3.3	-	9.6
Ireland	-	381.2	-	53.8
Italy	-	68.1	-	39.6
Portugal	21.7	7.3	-	9.7
Spain	17.2	9.7	-	17.5

5.2 Results for the euro area sovereign debt crisis period

The results for the euro area sovereign debt crisis period that spans from April 2010 to end-December 2011 show that arbitrage forces continue to function despite the turbulent market conditions. Arbitrageurs step into the market once the basis exceeds the overall transaction costs ($\theta + \beta_0$), at which point the adjustment process towards equilibrium speeds up. During the crisis period we find either no, or much slower adjustment speeds in the lower regime, where significant thresholds are identified (Table 3). These results are in line with our findings for the pre-crisis period. However, we find that the estimated transaction costs are around two to three times higher than in the pre-crisis period with an average of 190 basis points.

The sharply higher estimated transaction costs can be explained by decreased liquidity in peripheral sovereign credit markets, in combination with a markedly higher volatility of the basis (see Appendix C and F). As arbitrageurs face the risk that the arbitrage trade will go against them in the short- to medium-run, they will demand a higher compensation for undertaking the trade in volatile markets. The crisis period is characterised by much higher basis volatility across the countries in our sample compared to the pre-crisis period.

For the crisis period we cannot draw any general conclusion with respect to which market typically leads in the price discovery for credit risk as we find mixed results (based on the HAS measure). For the 5-year tenor, we find CDS leadership for Portugal and

Greece (upper regime, above the transaction costs). Results for the French and Irish cases suggest bond leadership. For the 10-year tenor we find CDS leadership in the upper regime for France and Greece, whereas bonds dominate for Germany. In the lower regime we find either bond leadership or no error correction at all.

Table 3: Price discovery TVECM - crisis period

This table reports the price discovery analysis for intraday data on a 30-minute sampling frequency from the TVECM for the period from April 2010 to end-December 2011 for the 5- and 10-year tenor. In the case of Germany, 5 year tenor (superscript $^+$), the VECM is a better fit compared to any threshold model based on maximum likelihood estimation. For further details see Table 1.

Panel A - 5-year tenor

Sovereign	$\theta + \beta_0$	HAS^U	λ_1^U	λ_2^U	HAS^L	λ_1^L	λ_2^L	obs.
France	132.8**	0.01	-65.84***	-8.51	0.16	-1.58	1.30	77.1%
Germany$^+$	-	-	-	-	-	-	-	-
Greece	227.7**	0.55	123.78	498.80*	0.27	-10.45*	11.56	89.5%
Ireland	175.5*	0.14	-53.32***	33.74	0.01	-10.84***	-0.31	70.6%
Italy	148.2***	0.02	-15.16	-0.85	0.76	-6.95	16.78	87.5%
Portugal	307.3***	0.78	-10.45	75.54**	0.03	-16.52*	3.94	87.9%
Spain	148.27	0.12	-17.14	2.48	0.75	-31.68	70.51***	80.7%
average	198.3							

Panel B - 10-year tenor

Sovereign	$\theta + \beta_0$	HAS^U	λ_1^U	λ_2^U	HAS^L	λ_1^L	λ_2^L	obs.
France	138.6*	0.99	4.44	25.98***	0.02	-18.46**	-3.43	86.0%
Germany	64.5*	0.13	-13.12**	5.16	0.06	37.13*	-9.58	36.9%
Greece	280.0***	0.57	10.00	15.16*	0.93	-1.64	4.42	44.6%
Ireland	167.7	0.26	-12.66	4.26	0.00	19.02	0.31	62.2%
Italy	142.3*	0.13	-22.94	-4.92	0.91	9.50	17.30	89.0 %
Portugal	300.1*	0.88	-8.43	-19.72	0.83	4.24	7.99	89.9%
Spain	95.4	0.17	-13.95	-5.29	0.18	-293.89	76.58	16.1%
average	185.1							

All half-lives are displayed in Table 4. The few cases where the speed of adjustments have a wrong sign (either CDS or ASW move away from the long-term equilibrium) the half-lives are not reported as the implied dynamics are unstable. As in the case of the pre-crisis period, the half-lives of any basis widening tend to be shorter in the upper regime compared to the lower regime. Again, this is in line with the notion that arbitrage trading activity is higher in the upper regime, leading to quicker readjustment of the basis. There are, however, a few cases where the estimated speed of adjustment is somewhat lower in the upper regime, possibly due to market disruptions among some of the worst affected sovereigns during the sovereign debt crisis.

Table 4: Half-life of shocks in days - crisis period

This table reports the half-life of shocks of 5-year and 10-year CDS and ASW for the period from April 2010 to end-December 2011. The half-lives of shocks are expressed in days, and are calculated using the impulse response function to a one unit shock on the cointegrating error, using Equations (4) and (5). In case the speed of adjustment is of the wrong sign we do not report any half-life. "Lower" denotes results for the region below the threshold, and "upper" above it.

	5-year tenor		10-year tenor	
Sovereign	lower	upper	lower	upper
France	133.7	6.7	25.6	17.9
Germany	-	-	-	21.0
Greece	17.5	1.0	63.5	74.6
Ireland	36.6	4.4	-	22.7
Italy	16.2	26.9	49.4	21.4
Portugal	18.8	4.5	102.7	-
Spain	3.7	19.6	1.0	44.4

5.3 Adjusted basis trade gain

To get a sense of the risk-return trade-offs arbitrageurs face in the market once the basis exceeds the estimated trading cost $(\theta+\beta_0)$, we calculate the so-called adjusted basis trade gain ($\mathrm{BTG}_{\mathrm{adj}}$). This measure represents the daily risk and cost adjusted potential basis trade gain, expressed in basis points, that an arbitrageur can typically expect in the upper regime, as implied by the model estimates. In the upper regime, the arbitrageur will bet on a declining basis while going short credit risk in the bond market and going long credit risk in the CDS market, ie by selling the bond and selling the CDS (Gyntelberg et al.; 2013). We assume that the typical basis, arbitrageurs encounter in the upper regime, is the mean value of the basis in the upper regime. After deducting the overall estimated transaction cost $(\theta + \beta_0)$ from this typical basis, we get the expected basis trade gain denoted as E(BTG) in Equation (6).

In order to get a time dimension associated with the trade gain (along the lines of an expected return per period of time), we scale E(BTG) by the half-life in the upper regime, hl_{d}^{U} (when defined). Here, the subscript d denotes that the half-lives are measured in days. In the short run, the arbitrageur faces the risk of the trade moving in the wrong direction which is directly proportional to the basis volatility. To generate a risk-adjusted measure, we adjust the daily potential trading gain by the daily basis volatility ($\mathrm{vola}_{\mathrm{d}}$).[14] Given this, the daily adjusted basis trade gain ratio $\mathrm{BTG}_{\mathrm{adj}}$ is then given by

$$\mathrm{BTG}_{\mathrm{adj}} = \frac{\mathrm{E(BTG)}}{hl_{\mathrm{d}}^{U}} \cdot \frac{1}{\mathrm{vola}_{\mathrm{d}}} \tag{6}$$

Table 5 shows that the expected basis trade gain is typically larger, and sometimes substantially so, in the crisis period compared to the pre-crisis period (Germany is an exception). Hence, despite higher trading costs facing arbitrageurs in the crisis period, the typical gains they may expect net of costs also tend to be higher. Once we adjust for the expected speed of adjustment and the risk (as measured by the basis volatility) associated with implementing arbitrage trades, the differences between the two periods are less stark. This suggests that part of the rise in the expected trading gains in the crisis period reflects higher compensation for risk. However, the fact that our $\mathrm{BTG}_{\mathrm{adj}}$ estimates do not fully equalize is likely due to a combination of imprecise parameter estimates and the imperfect nature of the basis volatility as a measure of arbitrage trade risk.

14 Daily volatilities of the basis are displayed in Appendix F.

Table 5: Daily risk and cost adjusted basis trade gain

This table reports the daily risk and cost adjusted basis trade gain (BTG_{adj}) from Equation (6) on a typical basis widening trade (arithmetic mean of the basis in the upper regime), in basis points. The few cases where the speed of adjustments have wrong signs are left empty.

Panel A - 5-year tenor

	pre-crisis		crisis	
Sovereign	E(BTG)	BTG_{adj}	E(BTG)	BTG_{adj}
France	12.53		34.63	4.02
Germany	24.70	0.02		
Greece	35.58	2.24	190.88	19.55
Ireland	42.22	0.05	82.59	5.27
Italy	15.39	0.13	26.27	0.49
Portugal	12.53	0.71	60.97	2.73
Spain	18.56	1.15	21.30	0.66

Panel B - 10-year tenor

	pre-crisis		crisis	
Sovereign	E(BTG)	BTG_{adj}	E(BTG)	BTG_{adj}
France	19.12	1.58	24.30	1.02
Germany	24.70	0.16	11.84	0.62
Greece	23.76	0.77	143.94	0.73
Ireland	41.89	0.30	55.95	1.11
Italy	24.44	0.29	60.13	1.27
Portugal	13.64	0.59	94.09	
Spain	8.40	0.32	33.87	0.37

6 Conclusions

The persistence of a positive basis between sovereign CDS and sovereign bond spreads in the euro area points to the presence of arbitrage costs that prevent a complete adjustment of market prices to the theoretical no-arbitrage condition of a zero basis. These include transaction costs and costs associated with committing balance sheet space for implementing arbitrage trades. Using a TVECM modelling approach, we are able to quantify these unobservable costs and study their properties.

We find that the adjustment process towards the long-run equilibrium intensifies once the CDS-bond basis exceeds a certain level/threshold. Above this estimated threshold, arbitrage trades become profitable for arbitrageurs while below the threshold, arbitrageurs have no incentive for trading as the costs they face are higher than the expected gain from the trade. As a result, we typically find faster adjustment dynamics towards the long-term equilibrium once the estimated threshold is exceeded (upper regime) compared to the lower regime, and the half-life of any basis widening therefore tends to be shorter in the upper regime compared to the lower regime. This supports our assumption that arbitrageurs step in and carry out basis trades only when the expected gain from the arbitrage trade is greater than the trading costs.

During the euro sovereign credit crisis in 2010-11, we find very high estimated transaction costs of around 190 basis points on average, compared to around 80 basis points before the crisis. This increase was likely due to higher costs facing arbitrageurs in the market, as well as higher risk that the trade would go against them due to substantially more volatile market conditions. In response, arbitrageurs demanded higher compensation for undertaking such trades during the crisis, resulting in higher thresholds. In line with this, we find that the expected trading gains facing arbitrageurs when the basis exceeds the threshold are higher in the crisis period than pre-crisis. Risk-adjusted trading gain ratios, which adjust for the expected speed of adjustment of the basis and for the volatility of the basis, displayed less stark differences, suggesting that part of the rise in expected trading gains during the crisis reflects compensation for higher trading risk.

Finally, we note that the divergence of CDS and ASW spreads during the crisis period can not be fully explained by decreased liquidity in peripheral sovereign credit markets. In fact, we find a significant increase of the CDS-bond basis in countries where measures of market liquidity increased during crisis period, as for example in France and Spain.

A Asset Swap Spreads

The asset swap spread, ASW, is the fixed value A required for the following equation to hold[15] (O'Kane (2000))

$$\underbrace{100 - P}_{\substack{\text{Upfront payment for bond} \\ \text{asset in return for par}}} \overbrace{+ \underbrace{C \sum_{i=1}^{N_{\text{fixed}}} d(t_i)}_{\text{Fixed payments}} = \underbrace{\sum_{i=1}^{N_{\text{float}}} \big(L_i + A\big) d(t_i)}_{\text{Floating payments}}}^{\text{Interest rate swap}}, \tag{7}$$

where P is the full (dirty) price of the bond, C is the bond coupon, L_i is the floating reference rate (eg Euribor) at time t_i, and $d(t_i)$ is the discount factor applicable to the corresponding cash flow at time t_i.

In order to compute the spread A several observations and simplifications have to be made. First, in practice it is almost impossible to find bonds outstanding with maturities that exactly match those of the CDS contracts and second, the cash-flows of the bonds and the CDS will not coincide. To overcome these issues, in what follows we use synthetic asset swap spreads based on estimated intraday zero-coupon sovereign bond prices. Specifically, for each interval and each country, we estimate a zero-coupon curve based on all available bond price quotes during that time interval using the Nelson and Siegel (1987) method. With this procedure we are able to price synthetic bonds with maturities that exactly match those of the CDS contracts, and we can use these bond prices to back out the corresponding ASW. As this results in zero coupon bond prices, we can set C in Equation (7) to zero.

A CDS contract with a maturity of m years for country j at time interval k of day t, denoted as $S_j(t_k, m)$, has a corresponding ASW $A_j(t_k, m)$:

$$100 - P_j(t_k, m) = \sum_{i=1}^{N_m} \big(L_i(t_k) + A_j(t_k, m)\big) \cdot d(t_k, t_i), \tag{8}$$

where $P_j(t_k, m)$ is our synthetic zero coupon bond price.

For the reference rate L_i in Equation (8), we use the 3-month Euribor forward curve to match as accurately as possible the quarterly cash flows of sovereign CDS contracts. We construct the forward curve using forward rate agreements (FRAs) and Euro interest rate swaps. We collect the FRA and swap data from Bloomberg, which provides daily (end-of-day) data. 3-month FRAs are available with quarterly settlement dates up to 21 months ahead, ie up to 21×24. From two years onwards, we bootstrap zero-coupon swap

[15] This assumes that there is no accrued coupon payment due at the time of the trade; otherwise, an adjustment factor would need to be added to the floating payment component.

rates from swap interest rates available on Bloomberg and back out the corresponding implied forward rates. Because the swaps have annual maturities, we use a cubic spline to generate the full implied forward curve, thereby enabling us to obtain the quarterly forward rates needed in Equation (8).

Given our interest in intraday dynamics, we follow Gyntelberg et al. (2013) and generate estimated intraday Euribor forward rates by assuming that the intraday movements of the Euribor forward curve are proportional to the intraday movements of the German government forward curve.[16] To be precise, for each day, we calculate the difference between our Euribor forward curve and the forward curve implied by the end-of-day Nelson-Siegel curve for Germany.[17] We then keep this difference across the entire curve fixed throughout that same day and add it to the estimated intraday forward curves for Germany earlier on that day to generate the approximate intraday Euribor forward curves. This approach makes the, in our view, reasonable assumption that the intraday variability in Euribor forward rates will largely mirror movements in corresponding German forward rates.

Finally, we need to specify the discount rates $d(t_k, t_i)$ in Equation (8). The market has increasingly moved to essentially risk-free discounting using the overnight index swap (OIS) curve. We therefore take $d(t_k, t_i)$ to be the euro OIS discount curve, which is constructed in a way similar to the Euribor forward curve. For OIS contracts with maturities longer than one year, we bootstrap out zero-coupon OIS rates from interest rates on long-term OIS contracts. Thereafter, we construct the entire OIS curve using a cubic spline. We use the same technique as described above to generate approximate intraday OIS discount curves based on the intraday movements of the German government curve.

[16] Euribor rates are daily fixing rates, so we are actually approximating the intraday movements of the interbank interest rates for which Euribor serves as a daily benchmark.

[17] Here we use the second to last 30-minute interval, because the last trading interval is occasionally overly volatile.

B CDS and ASW spreads

Figure B.1: CDS and ASW spreads in basis points

The figures are based on data with a 30-minute sampling frequency. Source: CMA Datavision, MTS

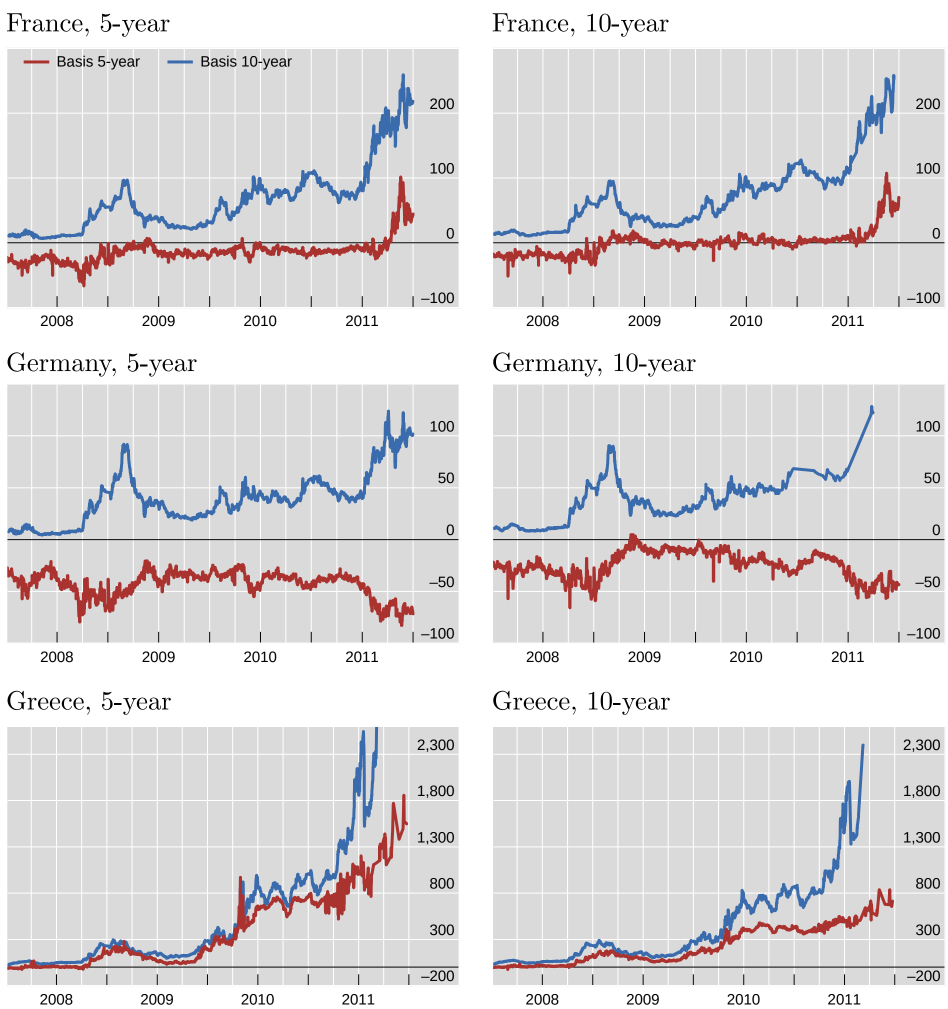

Figure B.1: (Cont.) CDS and asset swap spreads

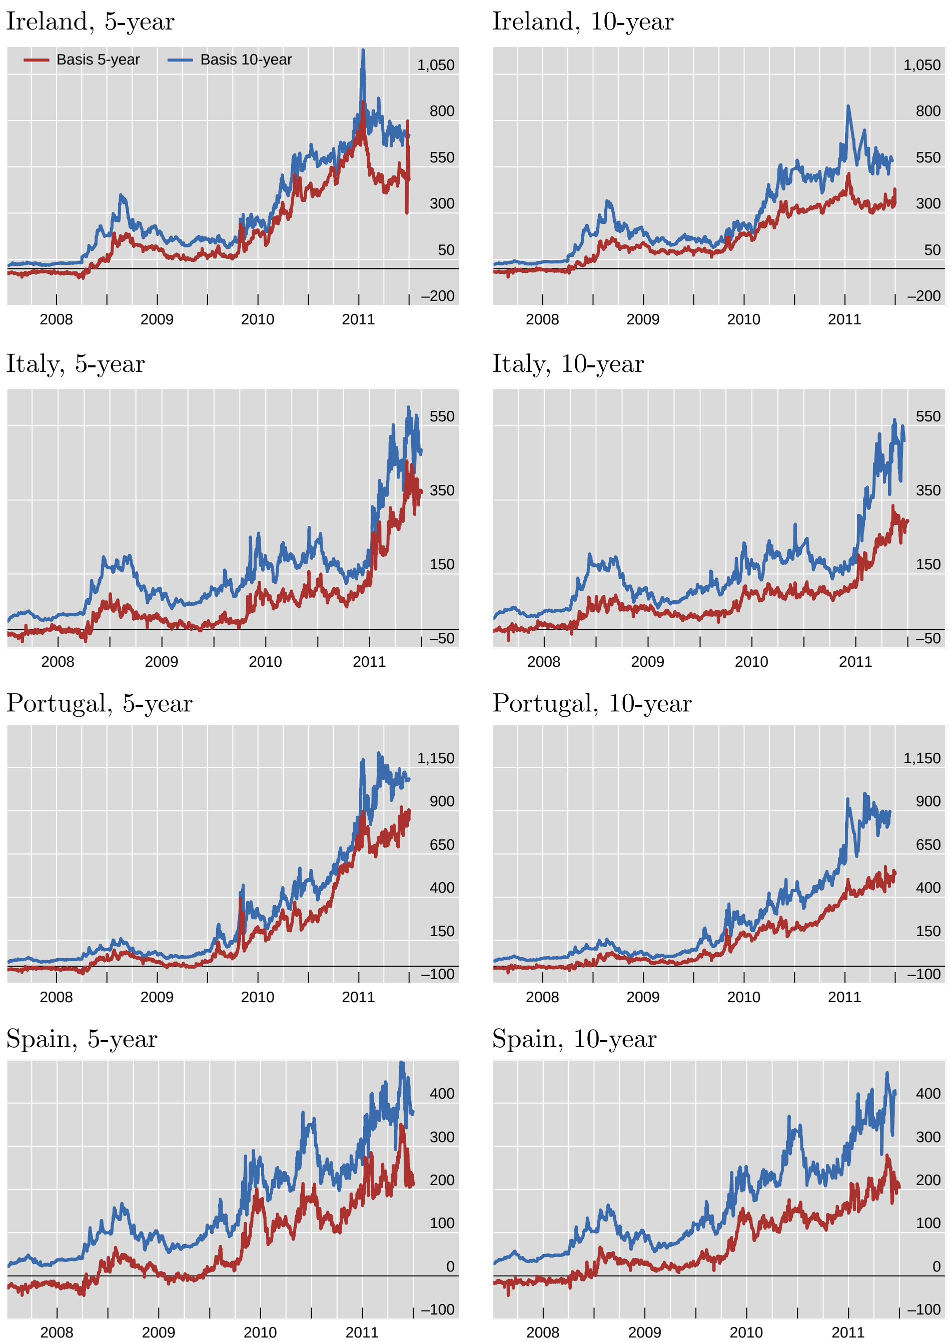

C CDS and Bond data and liquidity

Figure C.1: CDS data from CMA Datavision – tick-by-tick data

The right-hand scale shows the number (in thousands) of data ticks per year.

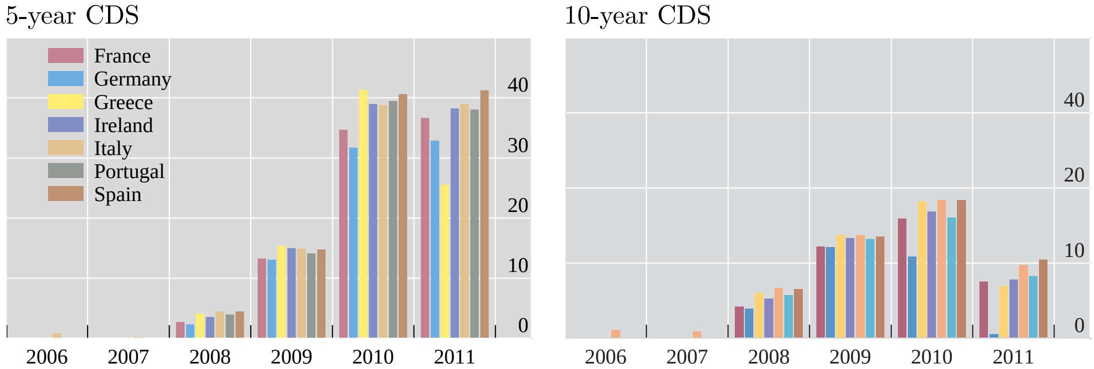

Figure C.2: CDS data from CMA Datavision – 30 min aggregates

The right-hand scale shows the number (in thousands) of non-empty half hour intervals per year. We consider 18 half hour slots per trading day, from 8:30 to 17:30 CET/CEST. The left-hand side scale shows the percentage of 30 min. intervals which contain at least one data tick during the 18 daily half-hour intervals we consider.

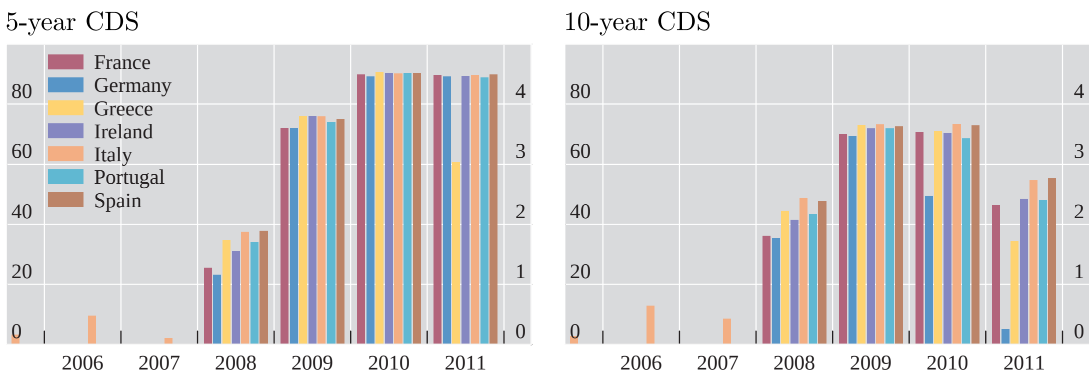

Figure C.3: EuroMTS number of trades

The right-hand side scale shows the number (in thousands) of trades per year. Italy is shown separately because the number of trades are more than an order of magnitude higher than for the other countries.

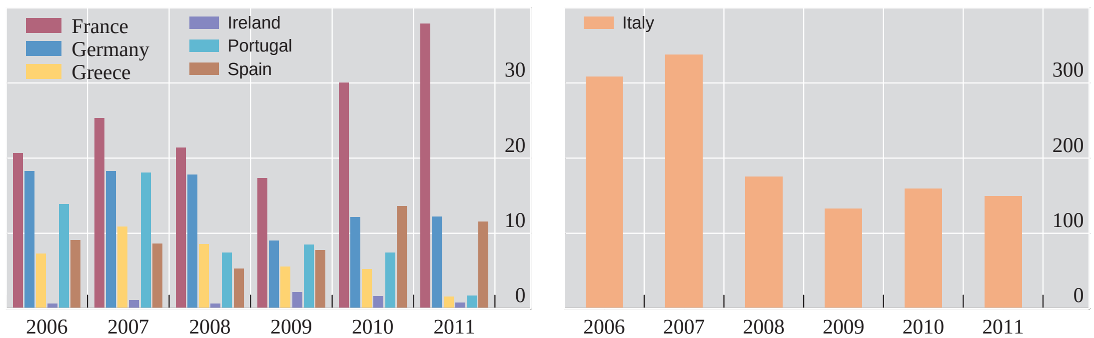

Figure C.4: EuroMTS bond price data from the trading book – tick-by-tick data

The right-hand side scale shows the number (in millions) of data ticks in the trading book. This includes all bonds with a maturity between 4 and 6 years and 9 and 11 years in the 5-year and 10-year segment, respectively.

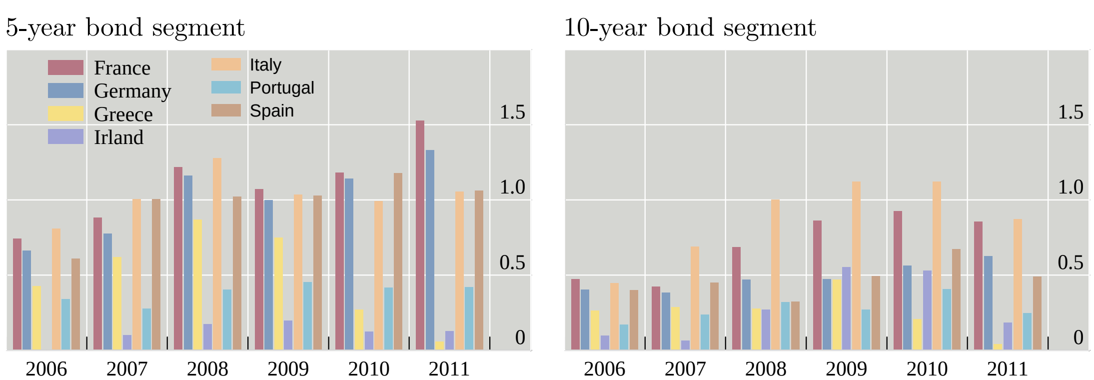

Figure C.5: EuroMTS bond price data from the trading book – 30 min aggregates

The left-hand side scale shows the percentage of 30 min. intervals during the trading period, which contain at least one data tick in the trading book. The right-hand scale shows the number (in thousands) of non-empty half hour intervals per year. We consider 18 half hour slots per trading day, from 8:30 to 17:30 CET/CEST.

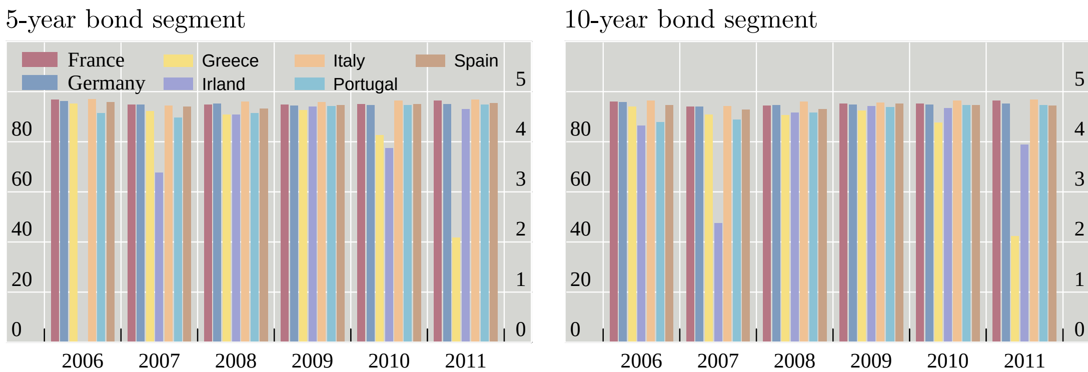

Figure C.6: Bid-Ask spreads for CDS and ASW in basis points

The figures are based on data with 30 minute sampling frequency. Source: CMA Datavision, EuroMTS

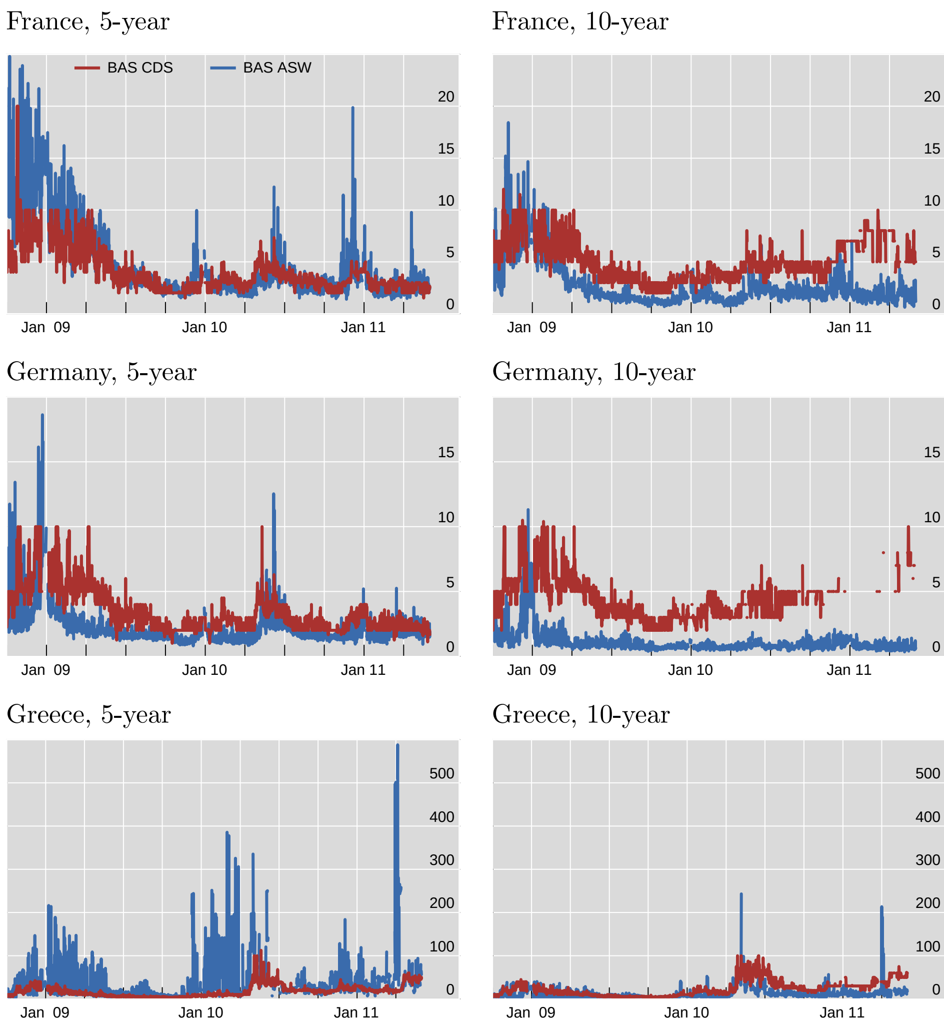

Figure C.6: (Cont.) Bid-Ask spreads for CDS and ASW in basis points

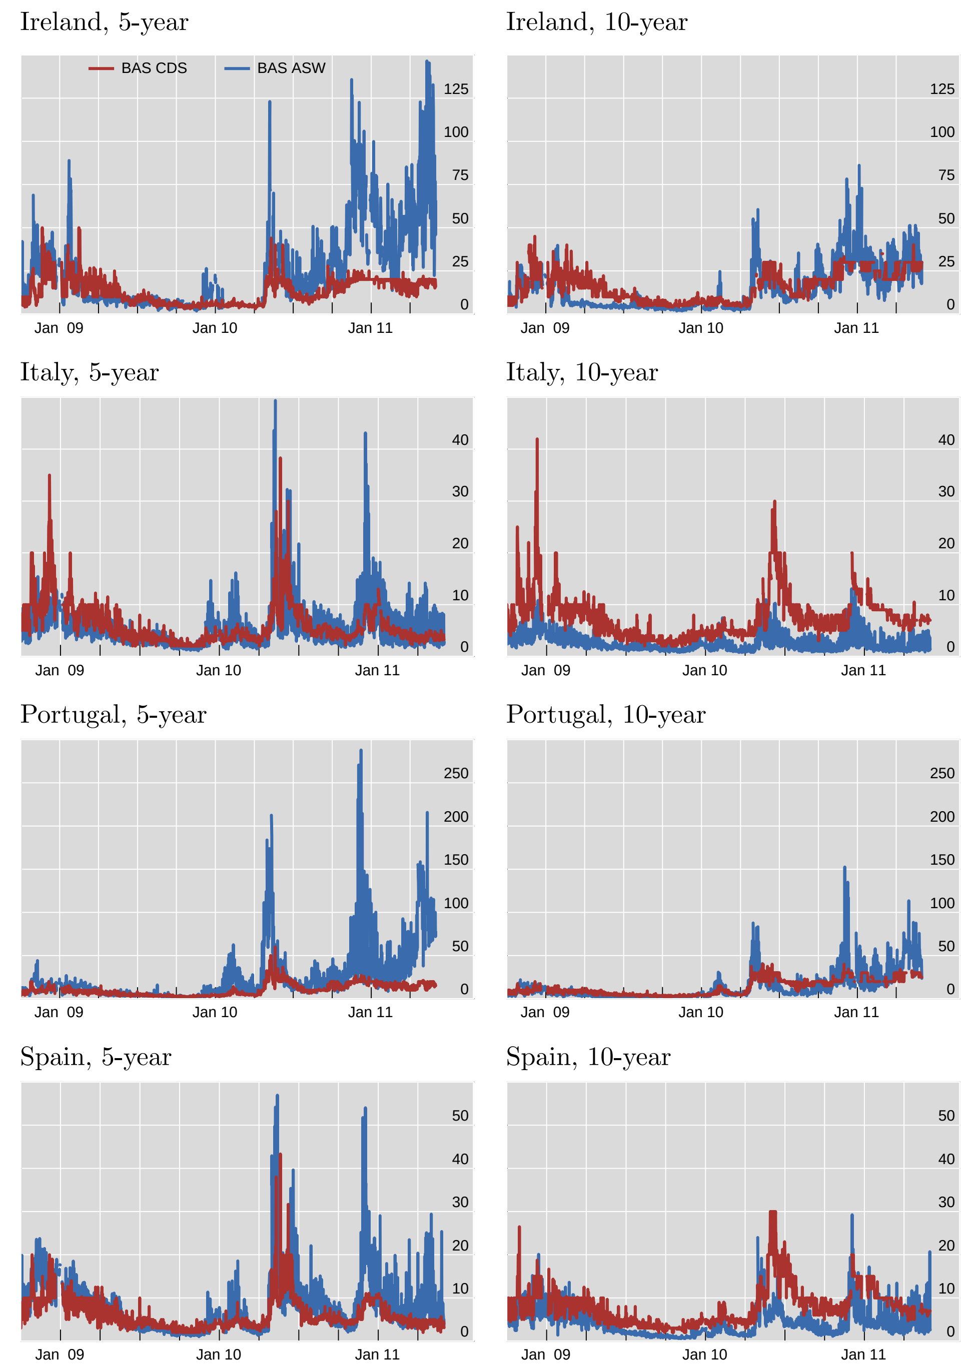

D Unit root and stationarity tests

We test for unit roots and stationarity in the CDS and ASW time-series using the following three methods:

1. the Augmented Dickey-Fuller (ADF) test,
2. the Phillips-Perron (PP) test and
3. the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test.

The null hypothesis of the ADF and PP test states: *the series has a unit root.* The null hypothesis of the KPSS test is: *the series is stationary.* Therefore, if our CDS and ASW data are $I(1)$ time series, we should be unable to reject the null hypothesis in levels for the ADF and PP test and reject H0 under the KPSS test, and vice versa for first differences.

Based on these three different tests we conclude that both the CDS and the asset swap spreads have a unit root for both tenors and periods (pre-crisis and crisis).

Our findings in Tables D.1 and D.2 show that for none of the CDS series in levels we are able to reject the null hypothesis of a unit root using either the ADF or the PP test. For the asset swap spread series the null is rejected for a few countries and tenors in levels using both the ADF and PP test. The KPSS rejects stationarity for all countries and both maturities. Test results for the first differenced spread data show that for all test methods we reject the unit root hypothesis across the board, indicating that all series are integrated of order one. To conserve space, we do not show these test results, but they are available from the authors on request.

Table D.1: Unit root and stationarity tests in levels - pre-crisis

The table reports the statistics of unit root and stationarity tests for the period from January 2008 to end-March 2010. The ADF and PP test for a unit root under the null hypothesis. For the KPSS test, the null is stationarity, and the 0.01, 0.05 and 0.10 critical values for the test statistics are 0.739, 0.463 and 0.347, respectively.

Panel A: 5-year spreads

	Credit default swap			Asset swap		
Sovereign	p_{ADF}	p_{PP}	KPSS stat.	p_{ADF}	p_{PP}	KPSS stat.
France	0.88	0.91	1.35	0.01	0.00	5.57
Germany	0.27	0.28	1.52	0.02	0.00	2.94
Greece	0.91	0.87	6.21	0.98	0.76	6.67
Ireland	0.48	0.48	2.87	0.13	0.12	7.04
Italy	0.45	0.62	2.24	0.01	0.00	3.73
Portugal	0.80	0.78	3.46	0.07	0.02	5.22
Spain	0.50	0.42	4.45	0.22	0.00	5.60

Panel B: 10-year spreads

	Credit default swap			Asset swap		
Sovereign	p_{ADF}	p_{PP}	KPSS stat.	p_{ADF}	p_{PP}	KPSS stat.
France	0.66	0.76	2.31	0.00	0.00	8.00
Germany	0.92	0.75	2.18	0.14	0.00	7.54
Greece	0.92	0.93	7.23	0.74	0.83	8.49
Ireland	0.47	0.28	4.68	0.20	0.62	9.58
Italy	0.31	0.30	3.36	0.06	0.23	6.77
Portugal	0.72	0.66	4.26	0.05	0.07	7.92
Spain	0.68	0.37	5.83	0.01	0.02	9.21

Table D.2: Unit root and stationarity tests in levels - crisis

The table reports the statistics of unit root and stationarity tests for the period from April 2010 to end 2011. Further details are presented in Table D.1.

Panel A: 5-year spreads

	Credit default swap			Asset swap		
Sovereign	p_{ADF}	p_{PP}	KPSS stat.	p_{ADF}	p_{PP}	KPSS stat.
France	0.80	0.74	7.57	0.15	0.18	4.43
Germany	0.80	0.62	7.35	0.60	0.31	7.16
Greece	1.00	1.00	7.79	0.00	0.00	9.67
Ireland	0.30	0.29	10.12	0.06	0.19	9.01
Italy	0.77	0.71	7.17	0.67	0.35	8.45
Portugal	0.79	0.69	10.80	0.26	0.14	11.29
Spain	0.11	0.08	7.83	0.03	0.03	7.51

Panel B: 10-year spreads

	Credit default swap			Asset swap		
Sovereign	p_{ADF}	p_{PP}	KPSS stat.	p_{ADF}	p_{PP}	KPSS stat.
France	0.99	0.98	7.94	0.49	0.81	5.21
Germany	0.48	0.49	4.10	0.59	0.09	6.58
Greece	0.94	0.97	8.68	0.17	0.00	5.36
Ireland	0.10	0.26	10.44	0.01	0.02	9.30
Italy	0.92	0.90	6.65	0.82	0.46	8.46
Portugal	0.77	0.79	11.26	0.01	0.09	11.26
Spain	0.86	0.73	8.02	0.19	0.15	8.51

E Cointegration analysis

We test for a long-run relationship in the form of cointegration between the bond and CDS market using the tests of Phillips and Ouliaris (1990) and Johansen (1988).

We view two series as cointegrated if either the null hypothesis of no cointegration is rejected using the Johansen or the Phillips-Ouliaris methodology. We use the Johansen test with intercept but no deterministic trend in the co-integrating equation. We use the Schwarz information criterion to estimate the optimal lag length for the Johansen test.

The test results indicate that in all cases, the CDS and the ASW spread series are cointegrated.

Table E.1: Cointegration - p-values, pre-crisis

This table reports the probabilities in decimals obtained from the Johansen cointegration and the Phillips-Ouliaris cointegration tests for the period from January 2008 to end-March 2010. For the Johansen test a constant is included in the co-integrating equation and the number of lags in the vector autoregression is optimized using the Schwarz information criterion. The Phillips-Ouliaris tests for no cointegration under the null hypothesis by estimating the long-term equilibrium relationship from a regression of CDS_t on ASW_t or from a regression of ASW_t on CDS_t among the levels of the time series. The column header ASW and CDS indicates which variable is used as dependent variable in the test.

Panel A: Johansen test

	Trace test				Maximum eigenvalue test			
	5-year		10-year		5-year		10-year	
Sovereign	None	at most 1	None	at most 1	None	at most 1	None	at most 1
France	0.000	0.435	0.003	0.612	0.000	0.435	0.001	0.612
Germany	0.143	1.000	0.159	0.664	0.039	1.000	0.104	0.664
Greece	0.001	0.786	0.000	0.441	0.000	0.786	0.000	0.441
Ireland	0.022	0.949	0.015	0.557	0.005	0.949	0.008	0.557
Italy	0.004	0.517	0.001	0.944	0.002	0.517	0.000	0.944
Portugal	0.001	0.354	0.000	0.728	0.001	0.354	0.000	0.728
Spain	0.024	0.783	0.000	0.618	0.009	0.783	0.000	0.618

Panel B: Phillip-Ouliaris test

	τ-statistics				z-statistics			
	5-year		10-year		5-year		10-year	
Sovereign	CDS	ASW	CDS	ASW	CDS	ASW	CDS	ASW
France	0.182	0.002	0.935	0.000	0.379	0.026	0.935	0.000
Germany	0.023	0.001	0.393	0.053	0.147	0.043	0.511	0.138
Greece	0.006	0.024	0.001	0.001	0.002	0.004	0.002	0.002
Ireland	0.002	0.001	0.017	0.016	0.028	0.022	0.080	0.076
Italy	0.000	0.000	0.000	0.000	0.005	0.000	0.001	0.000
Portugal	0.001	0.000	0.039	0.004	0.022	0.008	0.035	0.008
Spain	0.523	0.000	0.010	0.000	0.585	0.021	0.028	0.002

Table E.2: Cointegration - p-values, crisis

This table reports the probabilities in decimals obtained from the Johansen cointegration and the Phillips-Ouliaris cointegration tests for the period from April 2010 to end 2011. Further details are presented in Table E.1.

Panel A: Johansen test

	Trace test				Maximum eigenvalue test			
	5-year		10-year		5-year		10-year	
Sovereign	None	at most 1	None	at most 1	None	at most 1	None	at most 1
France	0.149	0.732	0.022	0.057	0.086	0.732	0.097	0.057
Germany	0.001	0.682	0.975	0.955	0.000	0.682	0.940	0.955
Greece	0.984	0.978	0.016	0.990	0.951	0.978	0.003	0.990
Ireland	0.011	0.104	0.050	0.224	0.030	0.104	0.077	0.224
Italy	0.209	0.721	0.168	0.516	0.134	0.721	0.145	0.516
Portugal	0.000	0.312	0.360	0.374	0.000	0.312	0.458	0.374
Spain	0.023	0.130	0.326	0.441	0.054	0.130	0.364	0.441

Panel B: Phillip-Ouliaris test

	τ-statistics				z-statistics			
	5-year		10-year		5-year		10-year	
Sovereign	CDS	ASW	CDS	ASW	CDS	ASW	CDS	ASW
France	0.951	0.106	0.008	0.401	0.945	0.103	0.163	0.319
Germany	0.000	0.000	0.028	0.015	0.001	0.001	0.102	0.074
Greece	0.000	0.000	0.019	0.034	0.000	0.000	0.012	0.019
Ireland	0.034	0.024	0.000	0.000	0.008	0.006	0.002	0.003
Italy	0.002	0.001	0.000	0.000	0.014	0.011	0.000	0.000
Portugal	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Spain	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000

F Basis volatility

Table F.1: Basis volatility - 5-year tenor

The table shows the volatility based on log changes and in bps. The pre-crisis period starts in January 2008 and ends in March 2010. The crisis period begins in April 2010 and our data ends in December 2011. The upper regime is above the estimated overall transaction cost $\theta + \beta_0$ and the lower regime is equal or below this level.

	pre-crisis		crisis	
Sovereign	lower regime	upper regime	lower regime	upper regime
France	2.14	1.25	1.08	1.29
Germany	1.01	1.08		
Greece	26.71	4.81	11.69	9.29
Ireland	2.40	2.31	7.27	3.56
Italy	2.18	1.72	2.26	1.99
Portugal	2.86	2.43	5.12	5.01
Spain	1.73	1.67	2.44	1.65

Table F.2: Basis volatility - 10-year tenor

The table shows the volatility based on log changes and in bps. The pre-crisis period starts in January 2008 and ends in March 2010. The crisis period begins in April 2010 and our data ends in December 2011. The upper regime is above the estimated overall transaction cost $\theta + \beta_0$ and the lower regime is equal or below this level.

	pre-crisis		crisis	
Sovereign	lower regime	upper regime	lower regime	upper regime
France	3.03	2.07	1.21	1.33
Germany	1.43	1.30	1.23	0.91
Greece	3.85	3.22	4.05	2.62
Ireland	2.84	2.62	4.24	2.21
Italy	3.05	2.15	2.31	2.23
Portugal	3.05	2.36	3.69	3.84
Spain	2.21	1.49	2.48	2.05

Concluding Remarks

The key focus of this thesis is the behaviour of credit risk markets before and during the euro area sovereign debt crisis. We have used intraday data sets for sovereign CDS from CMA Datavision and for government bonds from the MTS trading platform for the period from 2008 until 2011. With these intraday datasets it was possible to achieve a high statistical inference as well as investigate intraday behaviour of sovereign CDS and bonds. We focused on the European peripheral countries, the GIIPS countries (Greece, Ireland, Italy, Portugal and Spain), as well as France and Germany as riskfree/control countries. In order to broaden the scope we also investigated the impact of credit risk shocks from the GIIPS countries onto Central European countries, more specifically the Visegrad group (Czech Republic, Hungary, Poland and Slovakia) plus Austria as a control country.

The thesis had two main strands, firstly, we analysed contagion affects from and amongst the GIIPS countries before and during the euro area debt crisis, using CDS and bond data as credit risk channels. Secondly, we analysed the microstructure of the credit risk markets itself. Here, we focused on arbitrage costs mainly in CDS-bond basis trading as well as price discovery in credit risk markets.

In part I of the thesis we employed a panel VAR model. The results indicate that the CDS market was the main venue for the transmission of sovereign credit risk during the euro area sovereign debt crisis. In contrast, we find that, prior to the crisis, the two markets (CDS and bond) were similarly important in the transmission of financial spillovers. There is clear evidence for sovereign credit risk contagion during the euro area sovereign debt crisis period, as our results show more drastic reactions to shocks in terms of magnitude and absorption compared to the pre-crisis period. We find comovement effects, ie responses to credit risk shocks remain moderate in magnitude, rather than contagion during the pre-crisis period, as markets reacted rationally to economic fundamentals.

Our results using an unexpected exogenous macroeconomic news shock suggest that, during the pre-crisis period, markets for sovereign credit risk were driven by macroeconomic news. Positive news led to a decrease in credit spreads and negative news to an increase. For the euro area sovereign debt crisis period, our results show that movements in sovereign credit spreads did not respond to macroeconomic news but were rather driven by either monetary policy or exaggerations in financial markets due to lack of belief (a self-fulfilling crisis).

In 2010 markets become increasingly worried about the sustainability of sovereign debt especially in Greece, Ireland and Portugal and so concerns intensified on the side of policy makers about the fiscal outlook for the entire euro zone. These concerns led to the implementation of a series of financial support measures such as the European

Financial Stability Facility (EFSF) and the European Stability Mechanism (ESM). The Troika, a decision group formed by the European Commission, the ECB and the IMF, negotiated and supervised several bailout programmes during the euro area sovereign debt crisis. Our dataset covers the first four bailouts, two for Greece, one for Ireland and one for Portugal. By estimating contagion dynamics before and after the announcement of bailouts/economic adjustment programmes (EAP), we found that the magnitude of contagion, the interlinkages across countries, and the timeline of a shock absorption is strongly reduced when the country which received the bailout is also the origin of the shock. Our findings imply that the introduction of the EAPs had a positive effect on financial market's risk perception for the individual country under bailout. However, for the first three EAPs we do not find an effect on the joint systemic risk contagion among GIIPS countries after the bailout events. The exception is the fourth bailout which was the announcement of the preliminary draft of the second bailout for Greece. This draft addressed previous shortfalls and for the first time the EFSF was extensively enlarged. As a result, we find a stabilising effect across all GIIPS countries as the joint credit risk contagion and interlinkages are dramatically reduced.

Our findings show evidence that the Visegrad group countries have been immune to strong sovereign credit risk contagion from the GIIPS countries as we find comovement and no contagion. This result points to the fact, that investors must have focused on the countries that were most affected by the euro area sovereign debt crisis. Additionally, even though Hungary has experienced a significant increase in its debt-to-GDP ratio in 2010 with respective ratings downgrades, its impulse responses to a shock in the GIIPS countries do not differ from the other Visegrad group member countries. Also interesting in this perspective is the case of Slovakia which is the only euro area member amongst the Visegrad group countries. Contrary to the expectation that Slovakia should have a higher sensitivity to shocks from GIIPS countries due to its dependence on the euro, we do not find differences in the impulse responses for Slovakia compared to the rest of the Visegrad group.

In part II of the thesis we employed bi-variate threshold vector error correction models (TVECM), in order to understand the joint behaviour of two similar financial market instruments traded in the spot and in the derivative market (eg CDS versus ASW, gold versus gold futures or DAX (spot) versus FDAX (futures)). The basis is defined differently in different markets. We used the definition of the basis as the spot price minus the derivative prices, except for credit risk markets where we defined the basis as CDS minus ASW. The basis should fluctuated around zero in perfect markets, with no persistent deviations from zero. A persistent non-zero basis points towards the presence of transaction costs on arbitrage trades that prevent a complete adjustment of market prices to the theoretical no-arbitrage condition of a zero basis.

For economic reasons we can expect up to two thresholds, one threshold for basis weakening (positive transaction costs) and one threshold for basis strengthening (negative transaction costs). We have extended existing research on 2-regime TVECMs to 3-regime models (2-threshold TVECM). As there is practically no empirical research we carefully tested the statistical reliability of the proposed method, using large numbers of MC simulations with a VECM, a 1-threshold and a 2-threshold TVECM as a data generating process. We analysed a large variety of different setups, for example we used different grid sizes (for the parameter searches) and different time series length. We found that the slope β_1 of the error correction term $S_t - \beta_1 D_t - \beta_0$ is very precisely estimated. The threshold θ and the intercept β_0 of the error correction term are poorly estimated. However, the sum $\theta + \beta_0$ is found with high precision. This is an important finding, because this sum, represents the transaction cost.[18]

We perform an in-depth analysis for the example of sovereign CDS and bond data. The persistence of a positive basis between sovereign CDS and sovereign bond spreads in the euro area points to the presence of arbitrage costs that prevent a complete adjustment of market prices to the theoretical no-arbitrage condition of a zero basis. These include transaction costs and costs associated with committing balance sheet space for implementing arbitrage trades. Using a TVECM modelling approach, we were able to quantify these unobservable costs and study their properties. The CDS-bond basis for the considered period (2008-2011) suggests a 1-threshold TVECM.

We found that the adjustment process towards the long-run equilibrium intensifies once the CDS-bond basis exceeds a certain level/threshold. Above this estimated threshold, arbitrage trades become profitable for arbitrageurs while below the threshold, arbitrageurs have no incentive for trading as the costs they face are higher than the expected gain from the trade. As a result, we typically found faster adjustment dynamics towards the long-term equilibrium once the estimated threshold is exceeded (upper regime) compared to the lower regime, and the half-life of any basis widening therefore tends to be shorter in the upper regime compared to the lower regime. This supports our assumption that arbitrageurs step in and carry out basis trades only when the expected gain from the arbitrage trade is greater than the trading costs.

During the euro sovereign credit crisis in 2010-11, the estimated transaction costs are around 190 basis points on average, much higher as compared to around 80 basis points before the crisis. This increase is likely due to higher costs facing arbitrageurs in the market, as well as a higher risk that the trade would go against them due to substantially more volatile market conditions. In response, arbitrageurs demanded higher compensation for undertaking such trades during the crisis, resulting in higher thresholds. In line with

[18] The transaction cost is measured relative to the observed basis and not relative to the estimated basis, which is shifted by β_0.

this, we found that the expected trading gains facing arbitrageurs when the basis exceeds the threshold are higher in the crisis period than pre-crisis. Risk-adjusted trading gain ratios, which adjust for the expected speed of adjustment of the basis and for the volatility of the basis, displayed less stark differences, suggesting that part of the rise in expected trading gains during the crisis reflects compensation for higher trading risk.

Finally, the divergence of CDS and ASW spreads during the crisis period can not be fully explained by decreased liquidity in peripheral sovereign credit markets. In fact, we saw a significant increase of the CDS-bond basis in countries where measures of market liquidity increased during crisis period, as for example in France and Spain.

It is crucial for policy makers and regulators to understand the dynamics in the market for sovereign credit risk, especially in the derivative market, where contagion effects are more severe during our analysed crisis sample. Even though policy makers may not be interested in intraday movements in credit risk, our results show that the level impacts from the short-term dynamics are persistent. Hence, our results are important with regard to financial stability.

References

Abeysinghe, T. and Forbes, K. (2005). Trade linkages and output-multiplier effects: a structural VAR approach with a focus on Asia, *Review of International Economics* **13**: 356–375.

Adams, K. and van Deventer, D. (1993). Comment on intra-day arbitrage opportunities and price behavior of the hang seng index futures, *Review of Futures Markets* .

Adrian, T. and Brunnermeier, M. K. (2011). COVAR, *Working Paper 17454*, National Bureau of Economic Research.

Ammer, J. and Cai, F. (2007). Sovereign CDS and bond pricing dynamics in emerging markets: Does the cheapest-to-deliver option matter?, *International Finance Discussion Papers 912*, Board of Governors of the Federal Reserve System.

Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point, *Econometrica* **61**: 821–1414.

Ang, A. and Longstaff, F. A. (2011). Systemic sovereign credit risk: Lessons from the U.S. and Europe, *Working Paper 16982*, National Bureau of Economic Research.

Arce, O., Mayordomo, S. and Pena, J. I. (2012). Credit-risk valuation in the sovereign CDS and bonds markets: Evidence from the euro area crisis, *Documentos de Trabajo CNMV* **53**: 1–37.

Arghyrou, M. and Kontonikas, A. (2012). The EMU sovereign-debt crisis: Fundamentals, expectations and contagion, *The Journal of International Financial Markets, Institutions and Money* **25**: 658–677.

Bai, J. (1997). Estimating multiple breaks one at a time, *Econometric Theory* **13**(3): 315–352.

Bai, J., Julliard, C. and Yuan, K. (2015). Eurozone sovereign bond crisis: Liquidity or fundamental contagion, *Working Paper mimeo*, Federal Reserve of New York.

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes, *Econometrica* **66**(1): 47–78.

Baillie, R. T., Booth, G. G., Tse, Y. and Zabotina, T. (2002). Price discovery and common factor models, *Journal of Financial Markets* **5**: 309–321.

Balke, N. S. and Fomby, T. B. (1997). Threshold cointegration, *International Economic Review* **38**(3): 627–645.

Bekaert, G., Harvey, C. R. and Ng, A. (2005). Market integration and contagion, *Journal of Business* **78**: 39–69.

Berndt, A., Douglas, R., Duffie, D., Ferguson, M. and Schranz, D. (2005). Measuring default risk premia from default swap rates and edfs, *Working paper*, Cornell University.

Binder, M., Hsiao, C. and Pesaran, M. H. (2005). Estimation and inference in short panel vector autoregressions with unit roots and cointegration, *Econometric Theory* **21**: 795–837.

Black, F. (1971). Toward a fully automated stock exchange, *Financial Analysts Journal* **27**: 28–35.

Blanco, R., Brennan, S. and Marsh, I. W. (2005). An empirical analysis of the dynamic relation between investment-grade bonds and credit default swaps, *The Journal of Finance* **LX**(5): 2255–2281.

Brutti, F. and Saure, P. (2015). Transmission of sovereign risk in the euro crisis, *Journal of International Economics* **97**: 231–248.

Canova, F. and Ciccarelli, M. (2009). Estimating multicountry VAR models, *International Economic Review* **50**: 929–959.

Canova, F. and Ciccarelli, M. (2013). Panel vector autoregressive models: A survey, *Working paper series 1507*, European Central Bank.

Caporin, M., Pelizzon, L., Ravazzolo, F. and Rigobon, R. (2012). Measuring sovereign contagion in Europe, *Working Paper 05*, Norges Bank.

Chen, L., Lesmond, D. A. and Wei, J. (2007). Corporate yield spreads and bond liquidity, *Journal of Finance* **62**: 119–149.

Chinn, M. and Forbes, K. (2004). A decomposition of gloabl linkages in financial markets over time, *Review of Economics and Statistics* **86**: 705–722.

Claeys, P. and Vasicek, B. (2014). Measuring bilateral spillover and testing contagion on sovereign bond markets in europe, *Journal of Banking and Finance* **46(C)**: 151–165.

Constancio, V. (2011). *Contagion and the European debt crisis, Keynote lecture.*
URL: *https://www.ecb.europa.eu/press/key/date/2011/html/sp111010.en.html*

Corsetti, G., Pericoli, M. and Sbracia, M. (2011). *Correlation analysis of financial contagion, in: Financial Contagion: The Viral Threat to the Wealth of Nations*, Wiley.

Davies, R. B. (1987). Hypothesis testing when a nuisance parameter is present only under the alternative, *Biometrika* **74(1)**: 33–43.

De Santis, R. (2012). The euro area sovereign debt crisis: Safe haven, credit rating agencies and the spread of fever from Greece, Ireland and Portugal, *Working Paper 1419*, European Central Bank.

Duffie, D. (1999). Credit swap valuation, *Financial Analysts Journal* **55**: 73–87.

Ehrmann, M. and Fratzscher, M. (2017). Euro area government bonds - Fragmentation and contagion during the sovereign debt crisis, *Journal of International Money and Finance* **70**: 26–44.

Eichengreen, B. and Mody, A. (2000). What explains changing spreads on emerging-market debt?, *in* S. Edwards (ed.), *Capital Flows and the Emerging Economies: Theory, Evidence, and Controversies*, University of Chicago Press, pp. 107–136.

Eichengreen, B., Rose, A. and Wyplosz, C. (1996). Contagious currency crises, *Scandinavian Economic Review* **98**: 463–484.

Elton, E. J., Gruber, M. J., Agrawal, D. and Mann, C. (2001). Explaining the rate spread on corporate bonds, *Journal of Finance* **56**: 247–277.

Engle, R. F. and Granger, C. W. J. (1987). Co-integration and error correction: Representation, estimation and testing, *Econometrica* **35**: 251–276.

Ericsson, J. and Renault, O. (2006). Liquidity and credit risk, *Journal of Finance* **61**: 2219–2250.

Fama, E. F. and French, K. R. (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage, *The Journal of Business* **60**: 55–73.

Fomby, T. B., Lutz, K., Murphy, A. et al. (2013). *VAR Models in Macroeconomics - New Developments and Applications: Essays in Honor of Christopher A. Sims (series: Advances in Econometrics)*, Emerald.

Fontana, A. and Scheicher, M. (2016). An analysis of euro area sovereign CDS and their relation with government bonds, *Journal of Banking and Finance* **62**.

Forbes, C. S., Kalb, G. R. J. and Kofhian, P. (1999). Bayesian arbitrage threshold analysis, *Journal of Business & Economic Statistics* **17**: 364–372.

Forbes, K. (2012). The "Big C": Identifying contagion, *Working Paper 18465*, NBER.

Fulop, A. and Lescourret, L. (2007). How liquid is the CDS market?, *Technical report*, ESSEC and CREST.

Gale, G. (2006). Using and trading asset swaps, *Technical report*, Morgan Stanley.

Ghysels, E., Idier, J., Manganelli, S. and Vergote, O. (2017). A high frequency assessment of the ECB securities markets programme, *Journal of the European Economic Association* **15:1**.

Gibson, H., Hall, S. and Tavlas, G. (2012). The Greek financial crisis: Growing imbalances and sovereign spreads, *Journal of International Money and Finance* **31**: 498–516.

Goldberg, L. S. and Grisse, C. (2013). Time variation in asset price responses to macro announcements, *Working Paper No. w19523*, National Bureau of Economic Research.

Gonzalo, J. and Granger, C. W. J. (1995). Estimation of common long-memory components in cointegrated systems, *Journal of Business & Economic Statistics* **13**(1): 27–36.

Grammatikos, T. and Vermeulen, R. (2012). Transmission of the financial and sovereign debt crises to the EMU: Stock prices, CDS spreads and exchange rates, *Journal of International Money and Finance* **31**: 517–533.

Gyntelberg, J., Hoerdahl, P., Ters, K. and Urban, J. (2013). Intraday dynamics of euro area sovereign CDS and bonds, *BIS Working Papers 423*, Bank for International Settlements.

Gyntelberg, J., Hoerdahl, P., Ters, K. and Urban, J. (2017). Arbitrage costs and the persistent non-zero CDS-bond basis: Evidence from intraday euro area sovereign debt market, *BIS Working Papers 631*, Bank for International Settlements.

Hansen, B. E. (1996). Inference when an nuisance parameter is not identified under the null hypothesis, *Econometrica* **64**: 413–430.

Hansen, B. E. (2000). Testing for structural change in conditional models, *Journal of Econometrics* **97**: 93–115.

Hansen, B. E. and Seo, B. (2002). Testing for two-regime threshold cointegration in vector error-correction models, *Journal of Econometrics* **110**(2): 293–318.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery, *Journal of Finance* **50**(4): 1175–1199.

Hull, J. and White, A. (2000). Valuing credit default swaps i: No counterparty default risk, *Journal of Derivatives* **8**(1): 29–40.

Johansen, S. (1988). Statistical analysis of cointegrated vectors, *Journal of Economic Dynamics and Control* **12**: 231–254.

Kamin, S. and von Kleist, K. (1999). The evolution and determinants of emerging market credit spreads in the 1990s, *BIS Working Paper 68*, Bank for International Settlements.

Kaminsky, G., Reinhart, C. and Vegh, C. (2003). The unholy trinity of financial contagion, *Journal of Economic Perspectives* **17**: 51–74.

Kliber, A. (2014). The dynamics of sovereign credit default swaps and the evolution of the financial crisis in selected central european economies, *Czech Journal of Economics and Finance* **64(4)**: 330–350.

Komarek, L., Ters, K. and Urban, J. (2016). Intraday dynamics of euro area sovereign credit risk contagion, *BIS Working Papers 573*, Bank for International Settlements.

Koop, G. and Korobilis, D. (2016). Model uncertainty in panel vector autoregressive models, *European Economic Review* **81**: 115–131.

Levy, A. (2009). The cds bond basis spread in emerging markets: Liquidity and counterparty risk effects (draft).

Lien, D. and Yang, L. (2008). Asymmetric effect of basis on dynamic futures hedging: Empirical evidence from commodity markets, *Journal of Banking and Finance* **32**: 187–198.

Longstaff, F., Mithal, S. and Neis, E. (2005). Corporate yield spreads: Default risk or liquidity? new evidence from credit-default swap market, *Journal of Finance* **60**: 2213–2253.

Longstaff, F., Pan, J., Pedersen, L. H. and Singleton, K. J. (2011). How sovereign is sovereign credit risk?, *American Economic Journal: Macroeconomics* **3**: 75–103.

Low, R., Yao, Y. and Faff, R. (2015). Diamonds vs. precious metals: What shines brightest in your investment portfolio?, *International Review of Financial Analysis* **43**: 1–14.

Lütkepohl, H. (2006). *New Introduction to Multiple Time Series Analysis*, Springer.

Mahanti, S., Nashikkar, A., Subrahamanyam, M., Chacko, G. and Mallik, G. (2008). Latent liquidity: A new measure of liquidity, with an application to corporate bonds, *Journal of Financial Economics* **Vol. 88**: 272–298.

Man, K., W. J. and Wu, C. (2013). Price discovery in the u.s. treasury market: Automation versus intermediation, *Management Science* **59**: 695–714.

Mauro, P., Sussman, N. and Yafeh, Y. (2002). Emerging market spreads: Then versus now, *Quarterly Journal of Economics* **117**: 695–733.

Mayordomo, S., Pena, J. I. and Romo, J. (2011). The effect of liquidity on the price discovery process in credit derivatives markets in times of financial distress, *European Journal of Finance* **17**: 851–881.

McMillan, D. G. (2005). Threshold adjustment in spot-futures metal prices, *Applied Financial Economics Letters* **1**: 5–8.

Missio, S. and Watzka, S. (2011). Financial contagion and the European debt crisis, *Working Paper 3554*, CESifo.

Nashikkar, A., Subrahmanyam, M. G. and Mahanti, S. (2011). Liquidity and arbitrage in the market for credit risk, *Journal of Financial and Quantitative Analysis* **46**: 627–656.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves, *Journal of Business* **60**(4): 473–489.

Neri, S. and Ropele, T. (2013). The macroeconomic effects of the sovereign debt crisis in the euro area, *Workshops and conferences*, Bank of Italy.

O'Kane, D. (2000). Introduction to asset swap spreads, *Technical report*, Lehman Brothers.

Palladini, G. and Portes, R. (2011). Sovereign cds and bond pricing dynamics in the euro-area, *NBER Working Paper Series 17586*, The National Bureau of Economic Research.

Pan, J. and Singleton, K. J. (2005). Default and recovery implicit in the term structure of sovereign cds spreads, *Working paper*, MIT and Stanford.

Paulsen, J. W. (2014). Economic and market perspective, *Monthly report*, Wells Capital Management.

Phillips, P. C. B. and Ouliaris, S. (1990). Asymptotic properties of residual based tests for cointegration, *Econometrica* **58**(1): 165–193.

Rogers, J., Scotti, C. and Wright, J. H. (2014). Evaluating asset-market effects of unconventional monetary policy: a multi-country review, *Economic Policy* **29**: 749–799.

Sarkozy, N., Juncker, J.-C., Merkel, A. and Papandreaou, G. (2010). *Letter to J. M. Barroso and J. L. Zapatero concerning CDS trading in European government bonds.*
URL: *https://www.gouvernement.lu/732758/lettre-en.PDF*

Scotti, C. (2013). Surprise and uncertainty indexes: Real-time aggregation of real-activity macro surprises, *Working Paper International Finance Discussion Paper 1093*, FRB.

Stevens, J. (2015). Do transaction costs prevent arbitrage in the market for crude oil? Evidence from a threshold autoregression, *Applied Economic Letters* **22**: 169–172.

Stock, J. H. and Watson, M. W. (1988). Testing for common trends, *Journal of the American Statistical Association* **83**: 1097–1107.

Stock, J. H. and Watson, M. W. (2016). *Factor Models and Structural Vector Autoregressions in Macroeconomics*, Elsevier.

Sutcliffe, C. M. S. (2006). *Stock Index Futures*, Ashgate Publishing.

Tang, D. Y. and Yan, H. (2007). Liquidity and credit default swap spreads, *Research paper*, Kennesaw State University and University of South Carolina.

Ters, K. and Urban, J. (2017a). Credit risk contagion before and during the euro area sovereign debt crisis: evidence from central europe, *mimeo*, Bank for International Settlements.

Ters, K. and Urban, J. (2017b). The importance of the bond and the cds markets in the transmission of euro area sovereign credit risk contagion, *mimeo*, Bank for International Settlements.

Tsay, R. S. (1989). Testing an modeling threshold autoregressive processes, *Journal of the American Statistical Association* **84**(405): 231–240.

Urban, J. (2017). Limits to arbitrage: Estimating unknown arbitrage costs from a persistent non-zero basis, *mimeo*, Bank for International Settlements.

van Rixtel, A. and Gasperini, G. (2013). Financial crises and bank funding: recent experience in the euro area, *BIS Working Paper 406*, Bank for International Settlements.

Zhu, H. (2004). An empirical comparison of credit spreads between the bond market and the credit default swap market, *BIS Working Paper 160*, Bank for International Settlements.

Zivot, E. and Wang, J. (2006). *Modeling Financial Time Series with S-PLUS*, Springer-Verlag New York, Inc.

Erklärung

(gemäß §4, Abs. 4 der Promotionsordnung vom 15. August 2006)

Ich versichere wahrheitsgemäß, die Dissertation bis auf die in der Abhandlung angegebene Hilfe selbständig angefertigt, alle benutzten Hilfsmittel vollständig und genau angegeben und genau kenntlich gemacht zu haben, was aus Arbeiten anderer und aus eigenen Veröffentlichungen unverändert oder mit Abänderungen entnommen wurde.