Does it Matter where you Search? Twitter versus Traditional News Media

Costas Milas
Management School
University of Liverpool, United Kingdom

and

Rimini Centre for Economic Analysis, Italy
e-mail: costas.milas@liverpool.ac.uk

Theodore Panagiotidis
Department of Economics
University of Macedonia, Greece

and

Rimini Centre for Economic Analysis, Italy
e-mail: tpanag@uom.edu.gr

Theologos Dergiades
Department of International and European Studies
University of Macedonia, Greece
e-mail: dergiades@uom.edu.gr

Abstract

We compare news in Twitter with traditional news outlets and emphasize their differential impact on Eurozone's sovereign bond market. We reveal a two-way information flow between Twitter's "Grexit" tweets and "Grexit" mentions in traditional news which suggests not only that both types of news serve as important empirical predictors for the sovereign bond market but also that the 'old' (traditional news) and the 'new' (Twitter) media are connected; however, the influence of Twitter on traditional news is stronger. Grexit tweets raise the Greek spread more than Grexit mentions in traditional news. Weak contagion effects are recorded for Portugal and Ireland.

Keywords: Grexit, Twitter, Traditional news outlets, Sovereign spreads.

***JEL* Classification**: C10, G01, G12

We thank two anonymous reviewers for very helpful comments and suggestions. We thank Chaido Kotsiampasouli for collecting the data from the traditional news outlets as well as Lida Dimitriadou, Athansios Ziogas and Ioannis Arampatzidis for classifying (with respect to the language and the origin country) the text items extracted from the various traditional news outlets. We thank Ioannis A. Venetis for sharing the code to implement the Dufour *et al.* (2006) testing approach. We also thank Jason Douglas (The Wall Street Journal) for discussing with us issues related to the use of Twitter by financial journalists and Alex Kostakis, Jane Binner, Steven Vass and participants at the 2019 Money, Macro and Finance Conference (hosted by The London School of Economics) for comments. We have discussed the results of the paper in our opinion pieces for Barron's Investment News and Analysis (https://www.barrons.com/articles/twitter-influences-markets-more-than-traditional-media-51546875240) and The Conversation (https://theconversation.com/what-moves-markets-more-twitter-or-traditional-news-107854).

"The brains of humans and other animals contain a mechanism that is designed to give priority to bad news"

Daniel Kahneman (2011)

1. Introduction

Social media platforms allow news to travel much faster and wider compared to any time in the past. Social media platforms have become a popular open forum for analyzing economic/financial issues and, equally important, they reflect topics of public interest minute by minute. It has long since become essential for economic commentators, policymakers and their followers. For instance, New York Times Columnist and 2008 Nobel Laurate Paul Krugman runs a Twitter account with some 1 million followers during the Greek-related Eurozone crisis in early 2013; his followers have risen to approximately 4.6 million (in May 2020).[1] Also updating in real time are dedicated websites such as Real Time Economics of The Wall Street Journal (with its twitter account having approximately 0.8 million followers) that discuss hot economic topics. Not surprisingly, people pay attention. This is more so in the case of US President Donald Trump who currently has some 79.4 million Twitter followers (in May 2020) and what he writes makes a(n) (financial) impact. In fact, in an interview with The Financial Times on April 2, 2017, President Donald Trump noted "without the tweets, I wouldn't be here".[2]

There are good reasons to suggest that the information content on Twitter differentiates from the respective content of traditional news outlets. In terms of speed, for instance, reports published on social media sites can be accessed instantly whereas traditional media takes time to disseminate information (this is limited to once a day for newspapers; obviously television or radio can update their reports more frequently). In terms of creation and dissemination of content, traditional media work on the 'one-to-many' principle; an Editor decides what news is and the news consumers (readers and viewers) do not play a role in the creation or dissemination of content. Contrast this with the 'many-to-many' principle of social media, where any individual can create and share content. In terms of interactivity, all comments in social media occur in real time; traditional media instead is tightly patrolled. Further, social media connects billions of individuals across the globe, whereas traditional media limit their reach to the number of readers or viewers that individual newspapers or channels may have.[3]

[1] Krugman is the top economist in terms of followers (see: https://ideas.repec.org/top/top.person.twitter.html).
[2] See: https://www.ft.com/content/9ae777ea-17ac-11e7-a53d-df09f373be87.
[3] Worldwide, weekly social media use for news has enjoyed a steady increase over the 2013-2017 period (only to drop slightly in 2018); for 2018, the use was 45% in the US, 39% in the UK, 36% in France and 31% in Germany. See: http://www.digitalnewsreport.org/survey/2018/overview-key-findings-2018/.

The recent literature provides convincing evidence that information on social media influences financial markets. For instance, Azar and Lo (2016) show that publicly available tweets contain valuable information that can be used to forecast stock market returns over and above the impact of asset pricing factors. Agrawal et al. (2018) find that a sentiment index constructed based on social media, influences the supply and demand for liquidity in the stock market; they also find evidence supporting the capacity of social media to unveil insights above the respective insights of more traditional news feeds. By focusing on the global equity market, Beckers (2019) compares the predictive content of social media and traditional news to find that the information set delivered by the former is not different from the latter.

This paper is motivated by the agenda-setting theory of media in politics (see e.g. the discussion in Schroeder, 2018a and Neuman *et al.*, 2014) which suggests that the media have the power to determine important issues of the day and this applies to both Twitter and traditional media. This is not contradictory; as social media gain popularity over time, there is a shift from their *ability* to set the agenda to their *power* to set the agenda (Conway *et al.* 2015, refer to this as the crux of the current agenda-setting debate).

More specifically, the paper focuses on the information content of Twitter and traditional news media to raise the following question: does it matter where you get the information from, and if so, why? To answer this question, we proceed as follows. For a given topic related to financial markets, we compare within a bivariate testing framework, news appearing in both news sources and address the issue of which source of information "sets the agenda". We recognize, however, that verification of the agenda-setting hypothesis lacks sufficient practice in the absence of relevant validation. Therefore, we move on to validate the derived inference from the testing of the agenda-setting hypothesis within the context of financial markets. If, for instance, one source of news (say Twitter) is setting the agenda, then the underlying information set should provide insights that go beyond the respective insights of the competing source (say traditional media) in affecting financial markets (in terms of impact and duration). Therefore, if the hypothesis holds for one news source and is indeed validated, then it matters where one searches for information since a well-timed update of the information set will result in more informative investment decisions.

With this in mind, we examine the agenda-setting hypothesis by testing its empirical relevance for a topic related to the sovereign bond markets. We focus on the Greek debt crisis because it relates to a set of characteristics that allow an "apples-to-apples" comparison of the two news sources. These characteristics are as follows: first, the time persistence of the crisis that permits collection of data for a reasonable duration (our sample extends from 3/5/2012 to 6/24/2016, including 1,573 daily observations); second, the global interest for the crisis that creates adequate volume of resources and, third, the existence of the unique and untranslated acronym "Grexit" that directly refers to the Greek debt crisis.

Thus, we construct a unique time series dataset based on "Grexit" related news from both news sources. The comparison is conducted within a context that grants topic homogeneity, global geographic coverage, and inclusion of all potential languages. By selecting a topic that is described by the untranslatable term "Grexit" (adopted by financial reporters, commentators, and individuals), we establish, to a great extent, topic homogeneity. Twitter readily allows for global coverage. In order to establish the same geographic coverage for the traditional news outlets, we collect full-text documents from around the world based on more than 3,700 news sources (newspapers, magazines, broadcast transcripts from TV and radio as well as wire services). Finally, the untranslatable nature of the "Grexit" term allows us to identify tweets and text documents irrespective of their language. Methodologically, we test in a bivariate system the lead/lag relationship of the two news sources by conducting the Dufour *et al.* (2006) non-causality test that relies on the estimation of multiple-horizon Vector AutoRegressive (VAR) specifications. In addition, the response of one news source to a shock on the other source is evaluated by implementing the Jordà (2005, 2009) local projections approach, relying also on multiple-horizon VAR specifications.

The paper then examines the validity of the agenda-setting hypothesis by concentrating on Eurozone's sovereign bond market. More specifically, we focus on the borrowing costs of Eurozone's peripheral countries (namely Greece, Ireland, Italy, Portugal and Spain; hereafter the GIIPS). For comparison reasons, we also consider France as a core Eurozone country. In particular, we look at the differential impact of the "Grexit" mentions coming from Twitter, traditional news media and the orthogonal Twitter (we take out the effects of the traditional news in the latter) on the sovereign spreads, over and above the impact of economic/financial fundamentals, namely measures of default risk, liquidity risk and global financial risk. Sovereign bond spreads are defined as the difference between the 10-year government bond yield in each of the GIIPS and France relative to the German government bond yield. Methodologically, we act in a similar manner as in testing the agenda-setting hypothesis. The predictive capacity at different horizons of each news source over the sovereign spreads is evaluated by the Dufour *et al.* (2006) non-causality test, while each news source impact on the sovereign spreads is estimated by the Jordà (2005, 2009) local projections approach.

We have three main findings. First, there is a bidirectional information flow between Twitter and traditional news outlets, suggesting not only that both types of news serve as important empirical predictors for the respective market, but also that the 'old' (traditional news) and the 'new' (Twitter) media are connected. This connection brings into the analysis the concept of 'hybrid' media (Chadwick, 2013). Indeed, rather than making arbitrary distinctions according to simplified categories (such as 'old' versus 'new' media, or bloggers versus journalists), Chadwick (2013) postulates the side by-side existence of both media. Nevertheless, we find that the impact of Twitter on traditional news is more prolonged, stronger, and more robust (in terms of statistical significance) than the reverse. This finding gives prominence to Twitter as the agenda-setting news source in the context of the Grexit related discussion.

Second, the impact of Twitter's "Grexit" mentions on the Greek sovereign spread is positive and of higher magnitude than that of the traditional news outlets; in addition, the predictive power of Twitter persists even by taking out the effects of the traditional news (in terms of orthogonalizing the Twitter variable on the traditional news variable). Hence, the validation testing results robustify the role of Twitter as an agenda-setting news source in the sovereign bond market. This finding underlines the increasing power of information that appears and is shared on Twitter, bringing into the picture the importance of regulating social media; we return to this very issue in Section 6 where we discuss our main findings more in detail. Finally, our analysis shows weak contagion effects primarily for the case of Portugal and Ireland.

The paper proceeds as follows: Section 2 discusses the literature and Section 3 discusses the data used. Section 4 provides an outline of the methodology, whereas Section 5 reports the empirical results. Finally, Section 6 provides a discussion on our findings and concludes.

2. Literature Review

Rising social media popularity might explain why Twitter has recently decided to double the length of the classic 140-character tweet to 280.[4] Twitter is extremely popular with journalists.[5] Journalists tend to tweet immediately a breaking news story and often let their

[4] See: https://www.bloomberg.com/news/articles/2017-11-07/brevity-begone-twitter-doubles-tweet-limit-to-280-characters.

[5] Journalists were making up nearly a quarter (24.6 percent) of Twitter's authenticated users according to a 2015 report (see: https://www.poynter.org/news/report-journalists-are-largest-most-active-verified-group-twitter).

audience know that full coverage will appear on the newspaper's site soon. Since 2013, at least half of Twitter users in the US have reported getting news on the site; in 2017, however, that share went up to 74% and then fell slightly to 71% in 2018.[6] It is also interesting to note that Bloomberg has recognized the rapidly growing importance of Twitter in releasing financial information by integrating, since 2013, real-time Twitter feeds in its financial platform. In 2018, Bloomberg and Twitter expanded their relationship so that enterprise clients could incorporate Twitter-relevant news to their advanced trading strategies.[7] Despite the popularity of Twitter, a recent survey by Shearer and Gottfried (2017) notes that "many social media news consumers still get news from more traditional platforms". For instance, 55% of Twitter users often get news from news websites and 11% of Twitter users often get news from print newspapers.[8]

Twitter has become very popular in politics as it proxies attention paid to political issues (Barberá *et al.*, 2019). It has been shown to have some predictive power for election outcomes in the US (Heredia *et al.*, 2018), the UK (Boutet *et al.*, 2012), Germany (Tumasjan *et al.*, 2010) and around the world (Gayo-Avello, 2013). In examining the issue of whether Twitter or traditional media "set the agenda", existing studies find that the relationship between Twitter and traditional media is generally reciprocal. This simultaneity of 'old' and 'new' relates to the concept of 'hybrid' media (Chadwick, 2013) which seeks to integrate these seemingly disparate mediums (Twitter and traditional news).[9] Noting the lack of a model of how social media work differently from traditional media, Schroeder (2018b) views the new media as complementing traditional media rather than constituting a break with them.[10]

[6] See: http://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/ and http://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/.

[7] See: https://www.bloomberg.com/company/announcements/bloomberg-launches-twitter-feed-optimized-trading/.

[8] Although Facebook is the first among social media as news source, we focus on Twitter. This is justified by the nature of the topic. As the topic is about the sovereign risk of a country, the ideal user profile demands high education with capacity to comprehend economic topics. Twitter users fit better this profile as they are more educated relative to the users of other social media (Mitchell *et al.*, 2012). Dergiades *et al.* (2015) confirm, for the same topic, that the number of Grexit mentions mainly come from Twitter rather than Facebook.

[9] Chadwick (2013) notes that social media were central to Barack Obama's 2008 election campaign. At the same time, Obama's social media strategy was linked to old-fashioned mass-rallies of supporters and was coordinated with mass events, which, in turn, were carefully timed to generate maximum exposure on traditional media (television and newspapers). This suggests that older and newer media and political strategies have now become connected in a number of ways.

[10] Schroeder (2018b) notes that during the 2016 US presidential primaries, Donald Trump dominated the news headlines on the side of the Republican race to become the nominee largely because of Twitter, where he tweeted some rather controversial positions on a range of issues. These positions were subsequently featured in traditional media (TV and newspapers). In fact, traditional media appeared to give a lot of time to Trump's views because the American system is characterized by market competition for audience share and Trump's views boosted media ratings.

Neuman *et al.* (2014) use Granger causality tests for 29 issues (9 of which are economic ones, ranging from unemployment to corporate issues) that caught the attention of American politics in 2012 to find a two-way causality between traditional media and an aggregate index of social media (Twitter, blogs and discussion-forum data). Jungherr (2014) relies on Principal Component Analysis (PCA) to show that the temporal dynamics and content of Twitter messages are similar to those of traditional media only for some cases during the 2009 German federal election. Conway *et al.* (2015) focus on the 2012 US presidential election and use lead-lag cross-correlation coefficients to conclude that Twitter influences traditional news and vice versa with varying levels of intensity and differential time lags in the case of 7 issues (ranging from economic to healthcare) discussed during the election. Kruikemeier *et al.* (2018) study candidate visibility during the 2012 Dutch parliamentary election by comparing how often candidates are mentioned in the traditional media and social media (Twitter, Facebook). Using panel data regression analysis, they conclude that candidate visibility in the traditional media influence visibility in social media and vice versa.

Twitter is also a point of interest to other fields such as media studies or computer science. Araujo and van der Meer (2018) focus on 18 publicly listed companies in the Netherlands to show that Twitter activity about organizations has a positive influence on media coverage within a two-day window whereas media coverage influences Twitter activity positively in the same day and negatively in the following day. Meyer and Tang (2015) find that Twitter is widely used by US traditional news organizations as an additional channel for disseminating news, supporting the concept of 'hybrid' media, while Moon and Hadley (2014) demonstrate that journalists in the US utilize Twitter as a supplier for collecting information. An earlier study by Zhao *et al.* (2011) takes a different approach by comparing the content of Twitter with a typical traditional news source (The New York Times) in the context of unsupervised topic modelling techniques in statistics to show that Twitter is a good source of topics on celebrities and brands that have low coverage in traditional news media and that Twitter users actively help spread news of important world events.

Increasing focus of the recent finance literature is placed on the impact of social media rather than the relationship between social media news and traditional media news or their potential differential impact. From a theoretical point of view, media not only reflects, but also drives the expectations of managers and investors alike; expectations in turn feed into asset prices. To the extent that media content works as a proxy for investment sentiment, it carries predictive power for financial assets; at the same time, media visibility and content can increase an asset's investor base and also direct investment attention (see the discussion in Tetlock, 2015 and references therein and the empirical evidence, from a historical perspective, in Turner *et al.*, 2018). The empirical literature identifies significant social media effects on stock returns, stock volatility and earnings surprises (see, among others, Boudoukh *et al.*, 2019; Ben-Rephael *et al.*, 2017; Chen *et al.*, 2014 and Sprenger *et al.*, 2013). Media tone through a textual analysis of the content across newspaper articles has significant predictive power for future house prices (Soo, 2018). He (2017) finds that exposure to pessimistic news suppresses hiring and employment decisions. Dergiades *et al.* (2015) identify significant social media effects in sovereign bond markets.

Building on the above literature, we take the relationship between Twitter and traditional news outlets to sovereign bond markets, where rises in bond yields add to the risk of sovereign default and harm aggregate financial activity. Indeed, Gennaioli *et al.* (2014) use data for 46 countries to focus on the dire consequences of sovereign default on aggregate financial activity in the defaulting country; the impact is stronger in countries where domestic banks hold more public debt. Gennaioli *et al.* (2018) use a dataset of over 20,000 banks in 191 countries to quantify a significant negative relationship between a bank's holdings of government bonds and its lending during sovereign defaults. Altavilla *et al.* (2017) flag the amplification effect of sovereign stress on bank lending to domestic firms for a sample of Euro-area banks. Augustin *et al.* (2018), Wolski (2018) and Bedendo and Colla (2015) identify spillover effects from sovereign to corporate risk across Europe.

The next section of the paper proceeds with a detailed discussion of our dataset which will be used in the subsequent empirical analysis.

3. Data

Carrying out a comparison between Twitter and traditional news media is not a trivial task because of a set of emerging challenges that do not permit "apples-to-apples" evaluation. For instance, the discovery and classification of a topic necessitate the implementation of natural language processing algorithms, which have reduced efficiency when the text size is short, as is the case of Twitter. Hence, topic homogeneity is a concern. Other challenges, especially for topics of global interest, are the geographic coverage and the language coverage. Thus, a direct and meaningful comparison between the two news sources demands topic homogeneity and the same geographic and language coverage.

To examine whether the information content on Twitter feeds/leads the information content on the traditional news outlets (and vice versa), by pre-selecting a topic related to sovereign bonds markets, we construct a unique dataset based on the "Grexit" mentions coming from the two news sources. Among the existing community-based content sharing social media (e.g. Twitter, Facebook, Digg, Google+, Reddit), Twitter is used extensively as a platform for spreading news and principally economic/financial news. Hence, Twitter is used as one source of information with global coverage.[11] On the other hand, to establish global coverage for the traditional news outlets, we collect full-text documents from around the world based on more than 3,700 news sources (newspapers, magazines, broadcast transcripts from TV and radio and wire services).

The term "Grexit" was inaugurated in February 2012 by two Citigroup economists (Willem Buiter and Ebrahim Rahbari). As such, the sample starts from the first week of March 2012 (3/5/2012) and ends in June 2016 (6/24/2016), including 1,573 daily observations. For the above period, using as source of data the premium Twitter historical database of Followthehashtag,[12] we collect 936,837 unique tweets that contain the keyword "Grexit" or "#Grexit". By defining the number of followers of each tweet contributor as a measure of influence, then the average influence per tweet is 8,024 and the cumulative influence of the 936,837 tweets is 7.5 billion. Our sample covers geographically all countries around the globe (collected tweets come from 195 countries) and all languages (collected tweets are written in 14 different languages). In more detail, 76.1% of the sample tweets are written in English, Spanish and German (42.6%, 22.9% and 10.6%, respectively), while the countries ranked at the 95th percentile and above (see in Figure 1 countries in red color) contribute 79% to the total number of the collected tweets. Figure 1 groups the 195 countries in percentiles according to the density of contributed tweets. Countries in white color signify no contribution to the sample. "Grexit" seems to be an issue of discussion for Twitter users from North America and Europe.[13]

To construct the series that captures the intensity of the "Grexit" discussion in Twitter, we count the collected tweets, on a daily basis, the total number of mentions for

[11] Kümpel *et al.* (2015) mention that 69% of the studies dealing with news sharing use Twitter as a source platform. Twitter is in the lead of Facebook (17%), YouTube (12%) and Digg (8%).
[12] See: http://www.followthehashtag.com. The dataset is available from the website for a fee.
[13] The density of the sample tweets for the Eurozone countries is presented analytically in Figure A1.1 (see Appendix 1). Given the availability of coordinates for the Twitter data, a higher degree of disaggregation is also presented in Figure A1.2 for the European continent (see Appendix 1). Finally, the dynamic evolution of Figure A1.2 over time is available at: https://www.youtube.com/watch?v=WDo1FPH8kwU .

both terms “Grexit” and “#Grexit” and we assign each value to the respective day. In total, we identify 1,338,086 mentions. The created time series is presented in Figure 2.

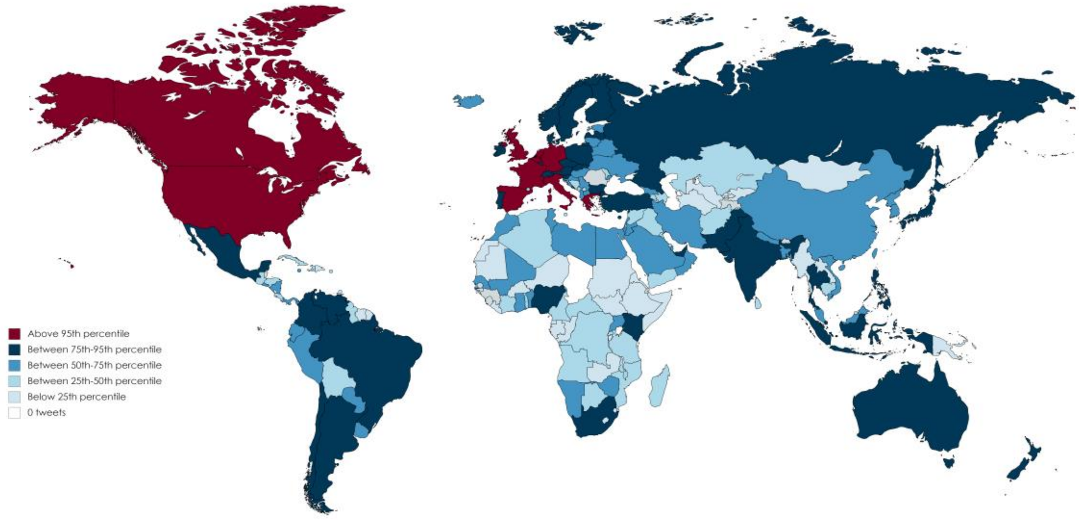

Figure 1. Tweets density per country.

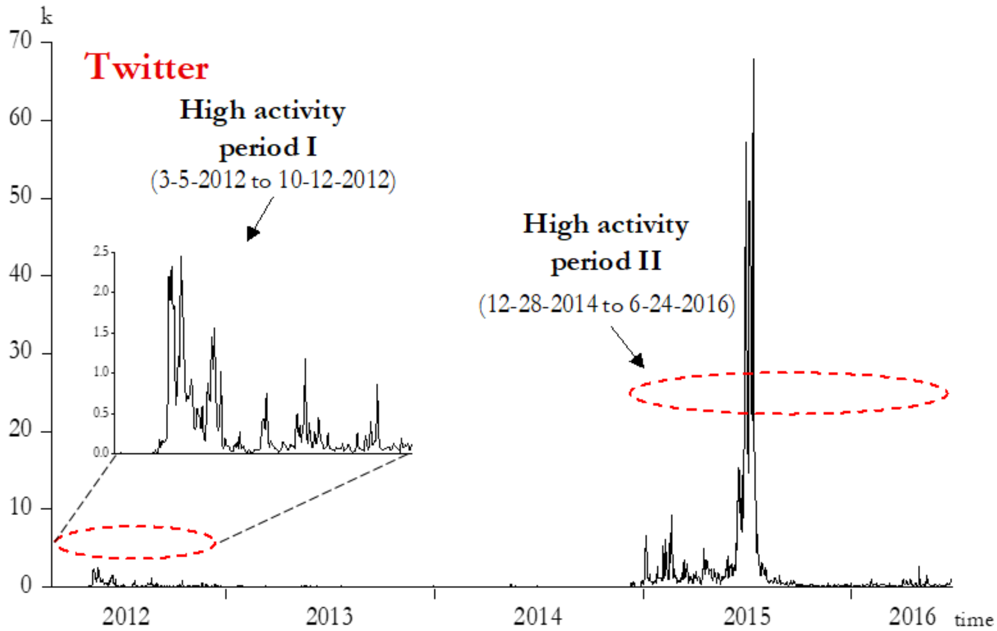

Figure 2. Grexit mentions in Twitter over time.

To construct the respective time-series with a daily frequency for the traditional news outlets, we use as source the LexisNexis Academic database. The coverage of the database is worldwide, offering access to multilingual text sources coming from newspapers, magazines, broadcast transcripts from TV and radio news as well as wire services. Hence, for the same sample period (3/5/2012 - 6/24/2016), we collect 40,341 unique text sources (e.g. newspaper articles) containing the keyword "Grexit" at least one time. Our sample covers geographically all countries around the globe (the collected text sources come from 83 countries) and all languages (collected text sources are written in 18 different languages). In more detail, 84.5% of the text sources from the traditional news are written in English, German, Dutch and French (42.5%, 22.7%, 10.8% and 8.5%, respectively), while the countries ranked at the 95th percentile and above (see red areas in Figure 3) contribute 63% to the total number of collected text items. Figure 3, groups the 83 countries in percentiles according to their text item contribution in the sample. Countries in white color signify no contribution to the sample. For the traditional news outlets, "Grexit" appears to be a topic of discussion mostly in Europe.[14]

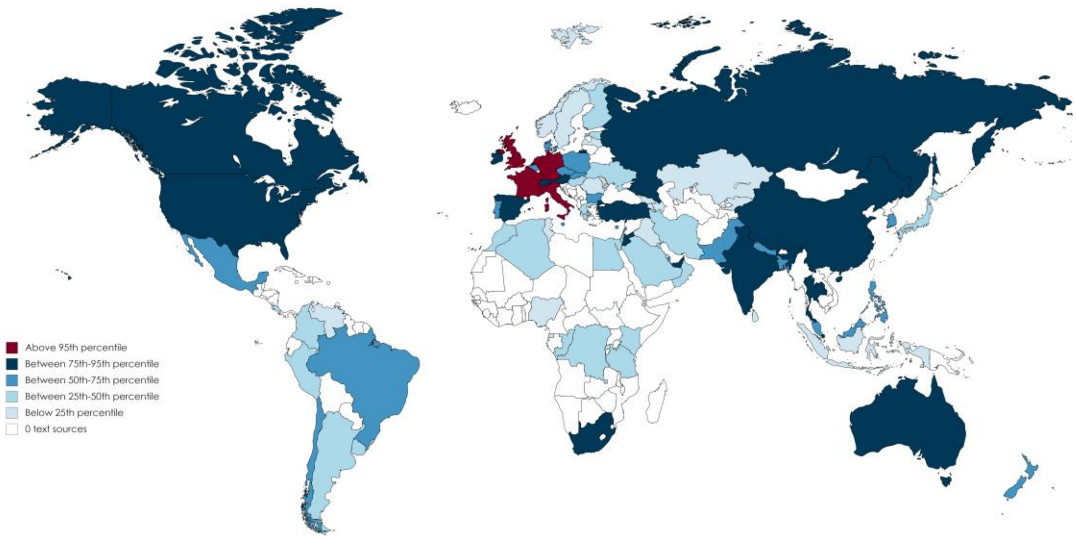

Figure 3. Traditional news density per country.

For the traditional news, we follow the same procedure as in the case of Twitter. To build the series that captures the intensity of the topic in the traditional news outlets, we count "Grexit" mentions for the collected text items on a daily basis and we assign each

[14] The density of the sample traditional news text items for the Eurozone countries is presented analytically in Figure A1.3 (see Appendix 1). For the traditional news text items no further disaggregation can be conducted (given data availability).

value to the respective day. In total, we identify 66,246 mentions. The constructed time-series is illustrated in Figure 4.

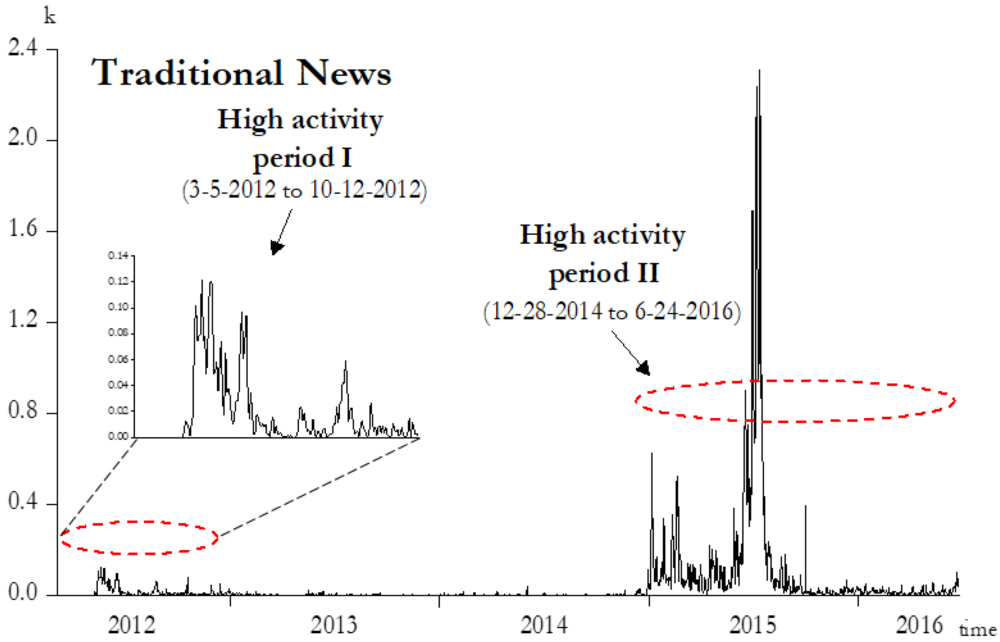

Figure 4. Grexit mentions in Traditional news outlets over time.

From Figures 2 and 4, we identify two high-activity periods for both series. The first high-activity period starts on 3/5/2012 (right after the introduction of the "Grexit" term) and ends at 10/12/2012. In particular, following the declaration of Mario Draghi to "do whatever it takes to preserve the euro" (7/26/2012),[15] the lower "Grexit" media activity after 10/12/2012 had to do with ECB's announcement for the unlimited bond buying plan on the secondary market (9/9/2012) and the Greek parliamentary vote on the 2013 budget (10/11/2012), which foresaw €13.5 billion budget cuts as a precondition to secure a new bailout loan from the European Union and the International Monetary Fund (discussed on the 10/12/2012 Eurogroup by the Eurozone finance ministers).

The second high-activity period begins on 12/28/2014 and terminates on 6/24/2016. The main reason for the reignited "Grexit" discussion was the snap general national election announcement, triggered by the failure of the Greek parliamentarians to elect a new head of state (12/29/2012). The prospect, revealed by opinion polls, that the election outcome might bring into power the radical left party of Syriza was sufficient to revive in a keenly manner the "Grexit" discussion.

[15] See https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html.

Descriptive statistics for both series referring to the full-sample and to the two high-activity samples, are reported in Table 1. The activity on Twitter is more intense than that in traditional news, as the Grexit mentions in Twitter reveal higher means and higher volatilities (in all samples). For instance, on 7/12/2015, when Grexit mentions in Twitter reached their maximum value of 67,948, the respective mentions in the traditional news media were 1,047 (this is almost a sixty fivefold difference). The considerably higher maximum values observed for Twitter mentions reflect Twitter's ultra-speed in disseminating news. In a failed Eurogroup meeting that took place on 7/11/2015 (that is, shortly after the 7/5/2015 Greek referendum), "Grexit" was closer than ever following Germany's proposal for Greece to take a 'time-out' of the common currency block for five years.[16] After the end of the meeting, in the early hours of 7/12/2015, it was Jeroen Dijsselbloem's (President of the Eurogroup) exit doorstep comment "it is still very difficult" that triggered the "Grexit" mentions in Twitter to reach their highest value (67,948 - 98% increase compared to the previous day). Traditional news reached their maximum value (2,312 - 120% increase compared to the previous day) a day later, on 7/13/2015.

Table 1. Descriptive statistics for the constructed time-series.

Statistic	Variable					
	Twitter mentions			Traditional news mentions		
	S^{f}	S^{I}	S^{II}	S^{f}	S^{I}	S^{II}
Mean	850.65	247.76	2310.00	42.11	14.90	112.55
Median	49.00	70.50	433.00	2.00	3.00	28.00
st. deviation	4263.59	440.18	7011.63	166.43	25.82	268.45
skewness	9.77	2.86	5.72	8.23	2.32	4.87
kurtosis	113.53	11.62	39.87	85.68	7.96	31.08

Notes: S^{f}, S^{I} and S^{II} denote the full sample, the first and the second high-activity periods, respectively.

Moreover, both constructed series appear to deviate vastly from normality, as inferred from the respective values of the skewness and the kurtosis. Hence, our empirical analysis is executed based on the logarithmic transformation of both constructed time-series, as a convenient way to (a) deal with the scaling issue of the constructed variables and (b) move from highly skewed variables to variables that are closer to normal.

Additionally, we report analytically, in Table 2, unit-root and stationarity tests for the log-levels of the two series (traditional news and the Twitter). The results of Table 2 support that the order of integration for the log-levels of both constructed series is I(0).

[16] See: https://www.theguardian.com/business/2015/jul/12/greek-crisis-surrender-fiscal-sovereignty-in-return-for-bailout-merkel-tells-tsipras.

This inference about the identified order of integration of the series is further supported by the Johansen (1991) cointegration test, which identifies two cointegrating vectors for a VAR system of two endogenous variables (see Table 3). The identification of two cointegrating vectors in this case, implies the lack of a unit root for both series favoring a VAR specification in the log-levels.

Table 2. Unit-root and Stationarity tests

Test	ADF	DF-GLS	PP	ERS	HOAC	BG	joint inference
Null	I(1)	I(1)	I(1)	I(1)	I(1)	I(0)	
N	**	*	***	***	**	ø	I(0)
T	*	*	***	***	**	ø	I(0)

Notes: ADF is the Augmented Dickey-Fuller test, DF-GLS is the Generalized Least Squares Dickey-Fuller test, PP is the Phillips-Perron test, ERS is Elliott-Rothenberg-Stock optimal point test, HOAC is the Bierens' Higher-Order Autocorrelation test and BG is the Bierens-Guo test. Across the second row, I(1) and I(0) reveal the null hypothesis of each test (unit root and stationarity, respectively). ***, ** and * signify the rejection of the null hypothesis at the 0.01, 0.05 and 0.1 significance level, respectively. The symbol ø signifies failure to reject the null hypothesis (for all four test statistics of BG test) at the conventional levels of significance. The joint inference is concluded by summarizing the individual inference of each test.

Table 3. Johansen (1991) cointegration test

H_0	**H_1**	**Trace Statistic**	**0.05 Critical Value**	**H_0**	**H_1**	**Max-Eigen Statistic**	**0.05 Critical Value**
$r = 0$	$r = 1$	60.426**	15.494	$r = 0$	$r = 1$	52.558**	14.264
$r \leq 1$	$r = 2$	7.867**	3.841	$r \leq 1$	$r = 2$	7.867**	3.841

Notes: ** signify the rejection of the null hypothesis at the 0.05 significance level and r refers to the number of cointegrating vectors. The reported *p*-values come from the MacKinnon *et al.* (1999).

The two independently constructed time-series (Twitter and traditional news; see Figures 2 and 4), present a high degree of positive linear correlation (the respective correlation coefficient is equal to 0.89). This high degree of linear association can be perceived as a signal of credibility towards the procedures used to build the series. Moreover, to provide further evidence on the robustness of the procedures used to construct the raw series, we first disaggregate the news items of both series by country of origin[17] and then calculate within (i.e. within the countries of each news source) and between (i.e. for each country between the two news sources) correlations. The within and between correlations for the disaggregated series are illustrated jointly in Figure 5.[18]

[17] Therefore, we have a separate time series for each country that participates in our sample.
[18] The countries presented are the first ten common countries to both news sources with the highest density in the sample. The exact correlation values are presented in Figure A1.4 in Appendix 1.

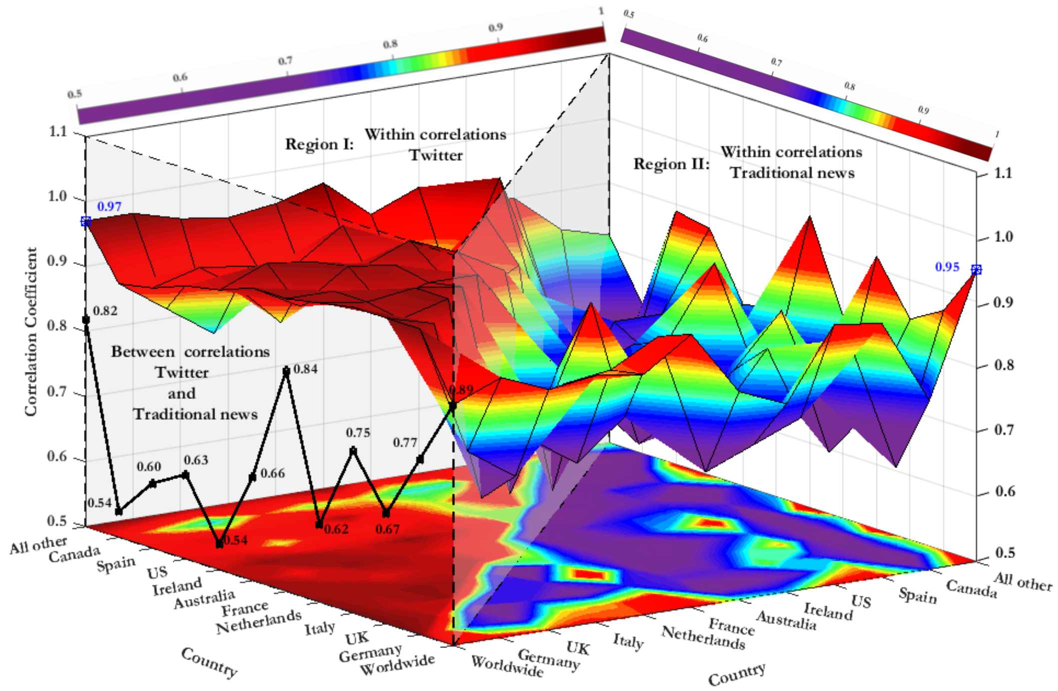

Figure 5. Surface of the within and between correlations by country.

Notes: "Worldwide" is the aggregate series. "All other" refers to all other countries of the sample and is constructed by summing each country's Grexit mentions per day. The left half surface (Region I) shows the within correlations for Twitter series (e.g. the left edge of the surface, signified by the **blue** symbol, suggests that the worldwide Twitter time series has a correlation of 0.97 with the Twitter time series of "All other" countries). The right half surface (Region II) shows the within correlations for the traditional series (e.g. the right edge of the surface, signified by the **blue** symbol, suggests that the worldwide traditional news time series has a correlation of 0.95 with the traditional news time series of "All other" countries). For convenience, the between correlations, signified by the **bold black** line, are presented in the transparent surface defined by the two left axes of the figure (e.g. for the case of Germany the correlation between the Twitter time series and the traditional news time series is 0.77).

Figure 5 shows that the within correlations for the Twitter series (Region I), on average, are higher than the within correlations for the traditional news series (Region II). This difference is mainly attributed to the variation in the density of observations per country. The density of the disaggregated series in Twitter is much higher than the respective disaggregated series in the traditional news (the unique tweets are 936,837, while the unique text sources are 40,341). Overall, we may argue that the disaggregated series are correlated consistently in a positive manner and primarily to a high degree, robustifying further the procedures used to build the Twitter and traditional news series.

Finally, to provide a visual sense of the disaggregated series' co-movement, we present, for both news sources, the normalized log-level series in Figures 6 and 7. We select to normalize the series by their respective standard deviation only for presentation purposes, as scaling issues in the raw data make the joint presentation non-informative. From both Figures 6 and 7, the series appear to reveal similar peaks and troughs; in addition, the Twitter series shows higher co-movement relative to the traditional news series. Overall, the disaggregation process shows that the country level series move in a quite uniform way.

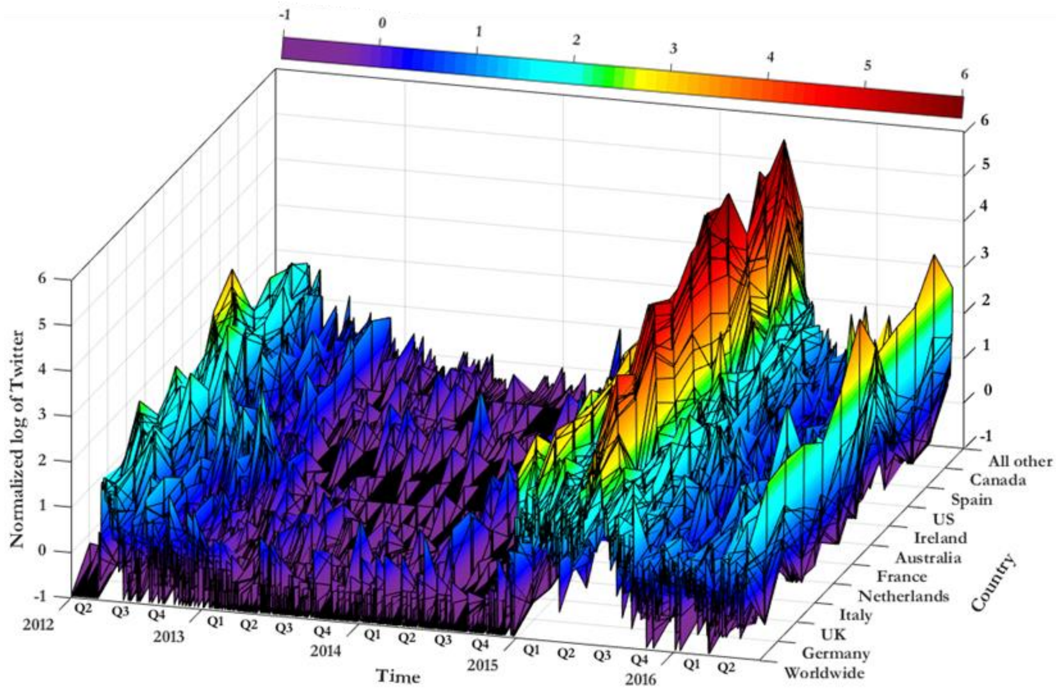

Figure 6. Log-normalized series of Twitter at country level.

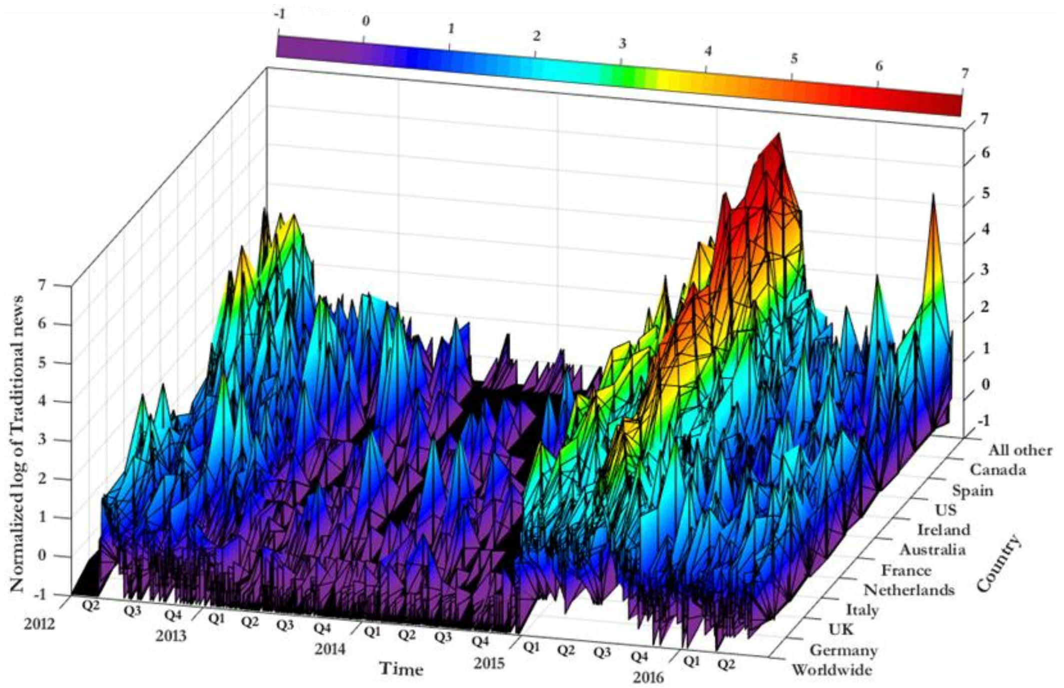

Figure 7. Log-normalized series of traditional news at country level.

4. Methodology

4.1. Non-causality at various horizons

To compare Twitter with traditional news media and examine their impact on the sovereign bond market, we rely on the Dufour *et al.* (2006) causal framework, which permits inferences on the linkages of a multivariate process, not only at a single horizon, but also at a multiple horizons framework. For finite order VAR processes, impediment in testing the non-causality hypothesis, at horizons different than one, is the non-linear nature of the imposed restrictions. Hence, the typical Wald-type statistics do not conform to the standard asymptotic theory. To alleviate these complications, Dufour *et al.* (2006) introduce a multiple-horizon VAR. After correcting for serial correlation in the error term, the validity of the restrictions is examined via a Wald-type test statistic (hereafter DPR).

Within this context, for the $W(t) = (w_{1t}, \ldots, w_{mt})'$ vector of random variables, the projection of a VAR process of order p at horizon h $\left(\mathrm{VAR}(p,h)\right)$ can be written as follows:

$$W(t+h) = \mu^{(h)} + \sum_{i=1}^{p} \pi_i^{(h)} W(t+1-i) + \sum_{j=0}^{h-1} \psi_j \alpha(t+h-j) \tag{1}$$

where $\mu^{(h)}$ is the constant term at horizon h $(h = 1,2,\ldots,H)$, $\pi_i^{(h)}$ are $m \times m$ coefficient matrices at horizon h and finally, ψ_j are $m \times m$ coefficient matrices that correspond to components of the MA$(h-1)$ process assumed for the error-term. The derivation of $\pi_i^{(h)}$ and ψ_j matrices is described in Dufour and Renault (1998). *Eq.* 1 is rewritten as follows:

$$W(t+h)' = \mu^{(h)\prime} + \sum_{i=1}^{p} W(t+1-i)' \pi_i^{(h)\prime} + u^{(h)}(t+h)' \tag{2}$$

with $u^{(h)}(t+h)' = \sum_{j=0}^{h-1} \alpha(t+h-j)' \psi_j'$. Using matrix notation *Eq.* 2 is represented as:

$$W(t+h) = \overline{W}_p(h) \Pi^{(h)} + U(t+h) \tag{3}$$

where $\Pi^{(h)}$ is a matrix of coefficients and $\overline{W}_p(h)\Pi^{(h)}$ the matrix of the variables.

The multiple-horizon VAR system in *Eq.* 3 is estimated by OLS. Once the estimate $\hat{\Pi}^{(h)}$ of $\Pi^{(h)}$ is attained, we may impose zero restrictions to test the non-causality hypothesis at horizon h. Suppose that it is of our interest to know whether the variable w_{st} causes at horizon h another variable, say w_{qt} $(1 \le s \le m,\ 1 \le q \le m$ and $s \ne q)$.

To test whether w_{st} does not cause w_{qt} given the available information set $\left(w_{st} \nrightarrow w_{qt} \left| \mathbf{I}_s \right.\right)$, we need to impose the following zero restrictions:

$$H^{(h)}_{w_{st} \nrightarrow w_{qt}} : \pi^{(h)}_{qsi} = 0, \qquad i = 1, \ldots, p \tag{4}$$

The non-causality hypothesis illustrated in *Eq.* 4 is tested through the Wald-type test statistic $\mathscr{W}[H_0(h)]$ (DPR statistic) that follows the χ^2 distribution with p degrees of freedom.

$$\mathscr{W}[H_0(h)] = T(R\hat{\beta}_q(h) - r)'\left[R\hat{V}_T(\hat{\beta}_q)R'\right]^{-1}(R\hat{\beta}_q(h) - r) \tag{5}$$

where, $R_{p\times(n+pm)}$ is selection matrix, $\hat{\beta}_q(h)$ is the $(n+pm)\times 1$ vector of OLS estimates for the q^{th} equation of the VAR system, $r_{p\times 1}$ is a vector of zeros and $\hat{V}_T(\hat{\beta}_q)$ is the Newey-West estimate of the $(n+pm)\times(n+pm)$ variance-covariance matrix. Finally, as noted in Dufour *et al.* (2006), failure to reject the null hypothesis consistently up to horizon $L = (m-2)p + 1$ is an adequate condition to verify absence of long-run causality.

Under the Dufour *et al.* (2006) framework, the testing procedure is adjusted to account for integrated processes up to order $d \geq 1$. The proposed adjustment follows the lines of the lag-extension practice introduced by Toda and Yamamoto (1995). Hence, if the involved process is integrated of order d, the optimal lag structure of the system illustrated in *Eq.* 2 is augmented by adding d extra lags. Once augmentation is done, the null hypothesis of no causality is examined by imposing restrictions on the optimal lag structure of the system (the extra lags are ignored).

Unfortunately, the asymptotic distribution of the $\mathscr{W}[H_0(h)]$ statistic proves to perform quite poorly in small samples. The performance of the test deteriorates further when the testing procedure is conducted in VAR systems with large order and with large number of variables. Furthermore, inference of non-causality at long horizons also disturbs the size and the power of the test as a consequence of the observed serial correlation. To control these concerns, Dufour *et al.* (2006) assess the validity of the null hypothesis by implementing a parametric bootstrap procedure. The bootstrap technique performs asymptotically considerably better in small samples, provided that the asymptotic distribution of $\mathscr{W}[H_0(h)]$ is nuisance-parameter-free.

4.2. Local projections

Starting from the VAR specification of order p at horizon h $\left(\mathrm{VAR}(p,h)\right)$, illustrated in *Eq.* 2, we further compute impulse responses based on local projections, as proposed by Jordà (2005). To circumvent the algebraic complexity involved in the estimation of the impulse responses within a standard VAR framework (introduced by the unique set of the VAR coefficients estimates), Jordà (2005) suggests obtaining a new set of coefficients estimates for each horizon h of *Eq.* 2. For instance, at horizon $t+h$, local projections constitute the response of the vector $W(t+h)$ to an experimental shock on the VAR reduced form residuals e at time t, given the available information set $\mathbf{I}_t$. Such a response is formally presented below:

$$\mathscr{L}_h^p = E\left(W(t+h)\middle|e_t = 1;\mathbf{I}_t\right) - E\left(W(t+h)\middle|e_t = 0;\mathbf{I}_t\right) \tag{6}$$

The structural impulse response is given by the following structural decomposition:

$$\Theta_h^p = \mathscr{L}_h^p A_0^{-1} \tag{7}$$

Hence, to construct $\hat{\Theta}_h^p$, we need an estimate of $\mathscr{L}_h^p$ which is attained by the coefficient matrix $\hat{\pi}_1^{(h)}$ of *Eq.* 2, while the impact matrix A_0^{-1} is recovered from the standard $\mathrm{VAR}(p)$ specification, after implementing an appropriate identification scheme.

To assess the shape of the response trajectories, given by (7), and the individual significance of each response, we construct the Scheffe´ and the conditional confidence bands respectively, as proposed by Jordà (2009). By letting $\hat{\theta}_{ij}$ to denote the estimated response of the variable i to a shock on variable j up to horizon h, then the Scheffe´s confidence interval is defined as:

$$\hat{\theta}_{ij} \pm \hat{A}_{ij}\hat{D}_{ij}\left(\sqrt{c_\alpha^2/H}\right)\mathbf{i}_H \tag{8}$$

where, $\hat{A}_{ij}$ is a lower triangular matrix and $\hat{D}_{ij}$ is diagonal matrix (both are estimated through Cholesky decomposition), c_α^2 is the critical value that corresponds to the χ^2 with H degrees of freedom and $\mathbf{i}_H$ is a vector of ones. Additionally, the conditional confidence bands are constructed as follows:

$$\hat{\theta}_{ij} \pm z_{\alpha/2} diag\left(\hat{D}_{ij}^{1/2}\right) \tag{9}$$

where, $z_{\alpha/2}$ is the critical value that corresponds to the standard normal distribution and $diag\left(\hat{D}_{ij}^{1/2}\right)$ is the vector with the diagonal elements of $\hat{D}_{ij}^{1/2}$. The benefits of using local projections over the standard VAR impulse analysis are as follows: first, the robustness over potential model misspecification; second, the joint inference for the impulse response coefficients and, third, the applicability of the approach to non-linear models. These advantages are discussed more analytically in Jordà (2005, 2009). The disadvantages of the local projections approach are summarized as follows: first, the impulse responses derived from small samples may be less precise compared to the standard VAR responses (although asymptotically local projections remain superior); second, the responses in the long run may be quite volatile, and third, the associated standard errors may be serially correlated. A comprehensive criticism on the use of local projections is provided by Kilian and Kim (2009).

5. Empirical findings

5.1. The information flow between Twitter and traditional news

To examine whether there is a two-way information flow between the Twitter (T_t) and traditional news (N_t) series, we implement the Dufour *et al.* (2006) non-causality testing approach. By identifying the causal dynamics, at horizons greater than one, we can assess not only the nature of the relationship between Twitter and traditional news outlets, but also the persistence/strength of the predictive content over time. To conduct the causal testing, we estimate *Eq.* 2 for the bivariate vector $W(t) = (T_t, N_t)'$ by specifying the optimal lag-length through the Schwarz Information Criterion assuming that the involved variables are not cointegrated.[19] *Eq.* 2 is estimated repetitively to obtain the $\mathscr{W}\left[H_0(h)\right]$ (or DRP) statistics up to twenty horizons (or days) ahead (h =1, 2,…, 20). The significance of the statistics is evaluated through bootstrapped *p*-values with 1000 replications. Moreover, starting from *Eq.* 2, we estimate the twenty periods response trajectory of T_t after a shock on N_t (and vice versa), under the Jordà (2005) local projections approach. The overall shape and the individual significance for each point on the impulse response path, are assessed by calculating respectively, the Scheffe´ and the conditional confidence bands. Furthermore, we test the significance of the impulse responses' twenty-period cumulative sum as proposed by Jordà (2009). This way, we assess the direction (positive or negative) and the significance of the overall impact of the shock on the target variable.

[19] As a robustness check we executed the same estimates under the assumption of cointegration. The derived causal inference is qualitatively the same as in the case of no cointegration and is available on request.

The p-values, for both hypotheses, are presented jointly in Figure 8.[20] The p-values for the hypothesis of no-causality running from Twitter to the traditional news outlets $(T_t \nrightarrow N_t)$ and vice versa $(N_t \nrightarrow T_t)$ are depicted by the red and the black line, respectively. The dark grey area and the light grey area imply significance at the 0.05 and 0.01 levels, respectively. From Figure 8, we infer that the first hypothesis $(T_t \nrightarrow N_t)$ is rejected at the 0.01 significance level for all horizons, except the last horizon where the rejection takes place at the 0.05 significance level. For the reversed case $(N_t \nrightarrow T_t)$, the rejection is the regular decision up to the ninth horizon, mainly at the 0.05 significance level, while for longer horizons, the non-rejection is the dominant inference. The results reveal a two-way predictive capacity between T_t and N_t with the effect of T_t on N_t being more prolonged and more robust in terms of significance.[21]

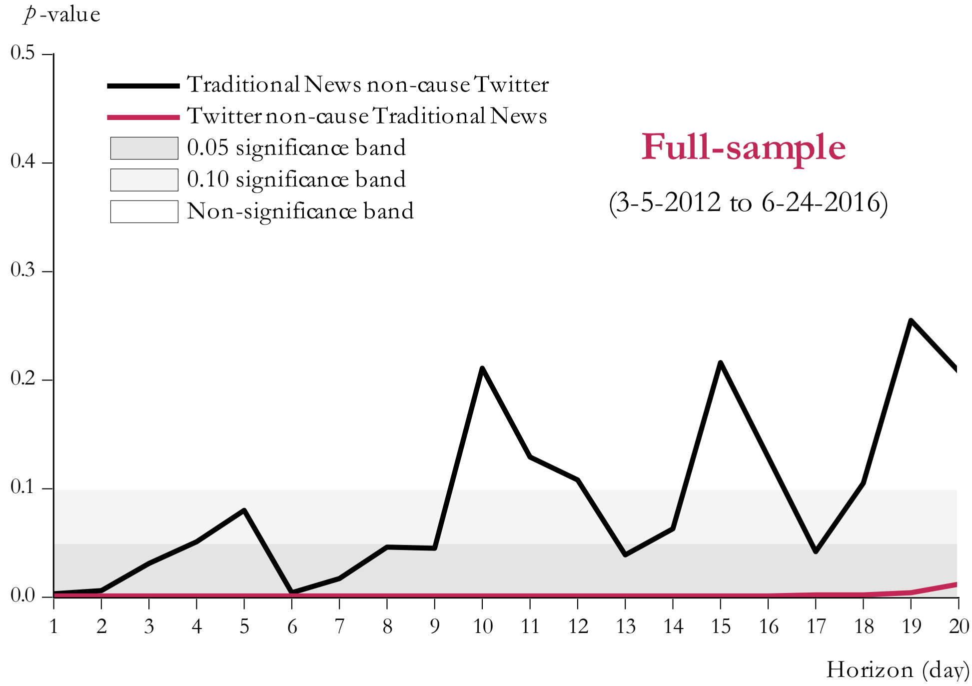

Figure 8. DPR causality tests p-values.

---

[20] The values of the DPR statistics for selected horizons are illustrated in Table A3.1 in Appendix 3.

[21] From Figure 8, one can notice that the derived p-values (associated to the hypothesis: traditional news non-cause twitter; see back line), illustrate a cyclical pattern. We argue that this cyclical pattern is principally related to the newspaper reader behavior, which differs between weekdays and weekends. In particular, as more news is consumed during weekends (newspaper circulation is lower on weekdays compared to weekends), more articles are published in weekends and therefore more Grexit mentions are expected, on average, to be collected. Given that Twitter users often get news from print newspapers (Shearer and Gottfried, 2017), the cyclical pattern of the Grexit mentions observed on traditional news, is expected to pass on also to the Twitter. We support empirically this argument by fitting a set of periodic regressions to both series (Twitter and traditional news) with the purpose to extract the optimal fitted cycle (7 days in both cases). Once both series have been de-cycled, the Dufour *et al.* (2006) non-causal testing is re-applied in order to obtain the new p-values over the same number of horizons. Although the initially observed cyclical pattern is heavily reduced, the newly derived p-values provide the same qualitative causal inference as in Figure 8. These results are available on request.

We further construct the twenty periods ahead impulse response of $T_t\ (N_t)$ following a generalized one standard deviation shock on $N_t\ (T_t)$. Figure 9 shows how a shock on $T_t\ (N_t)$ is transmitted to $N_t\ (T_t)$. In particular, Figure 9.a provides the response of N_t to a shock on T_t (continuous black line) along with the 95% Scheffe´ confidence interval (grey area) and the 95% conditional confidence interval (blue area). In this case, the conditional confidence band supports consistently a positive trajectory, while the corresponding Scheffe´ confidence band provides additional evidence that the impulse response is expected to fluctuate above zero. Furthermore, at the bottom left hand-side of the same figure, we report the twenty-horizon cumulative sum (C. sum) of the responses with the respective *p*-value for testing the significance (C. sum *p*-value).[22] Moreover, the magnitude of the responses implies that a 1% increase in the activity of T_t would lead to a 1.29% increase in the activity of N_t at the first horizon and to a 15.77% (and significant) cumulative increase after twenty-horizons.

Figure 9.b illustrates the response of T_t to a shock on N_t. The conditional confidence band does not include zero for any but three horizons, whilst the Scheffe´ confidence band suggests a positive trajectory up to horizon nine. The twenty-horizon cumulative impact is significant (at the 0.01 level), while a 1% increase in the activity of N_t leads to a 0.71% increase in T_t at the first-horizon and to a 6.66% cumulative increase. The impulse response results provide evidence in favor of: (i) a positive relationship between Twitter and the traditional news media, (ii) a significant bidirectional cumulative impact, (iii) a positive and significant impact of Twitter on the traditional news outlets, that is more prolonged compared to the respective impact in the reverse direction, and (iv) empirical evidence that the twenty-horizon impact of Twitter (for a 1% increase in mentions) on the traditional news is approximately 2.37 (=15.77%/6.66%) times higher than the reverse impact. Overall, these findings come to enhance the validity of the Dufour *et al.* (2006) testing results.

[22] Notice that *p*-values below the selected value of *a* imply that a shock on Twitter has a significant twenty days cumulative impact on the activity of the traditional news outlets.

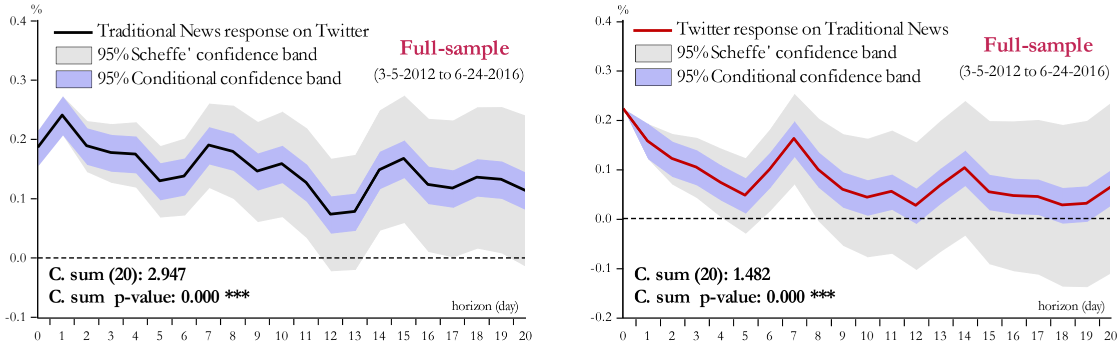

9.a. N_t response on T_t **9.b.** T_t response on N_t

Figure 9. Local projections impulse responses.

Notes : C. sum (20) refers to the twenty-horizon (days) cumulative sum of the impulse responses, while C. sum *p*-value is the resulting *p*-value for testing the hypothesis that the C. sum (20) is equal to zero. Finally, *** denote rejection of the null hypothesis at the 0.01 significance level.

From the executed causal and impulse response analysis (a summary of the results is presented in Figure A2.1 in Appendix 2), we find that the information flow between Twitter and the traditional news outlets is bidirectional which suggests that both Twitter and traditional media set the agenda. Nevertheless, in terms of dynamic interactions, our analysis suggests that the impact of Twitter on traditional news is more prolonged, stronger and more robust (in terms of statistical significance) than the reverse impact. This gives prominence to Twitter as the agenda-setting news source in the context of the Grexit related discussion.

5.2. Twitter, traditional news and sovereign spreads in the GIIPS

The previous section provides evidence of a bidirectional content feed between the two sources of news dissemination; furthermore, Twitter feeds in content the traditional news outlets more systematically than the other way around. In this section, we move on to examine whether the predictive capacity of Twitter towards the bond market is above and beyond the respective capacity of the traditional news media. More specifically, we assess the predictive capacity of the two news sources over the sovereign bond spreads (S_{jt}) for the GIIPS (Greece, Ireland, Italy, Portugal, Spain) and France (hence, $j = 1,\ldots,6$) by estimating $Eq.$ 2 for the multivariate vector $W(t) = (S_{jt}, M_{kt}, L_{jt}, D_{jt}, E_t, G_t)$.[23] We follow Dergiades *et al.* (2015) in capturing country-specific idiosyncratic risk by two types of risk, that is, the credit or default risk, D_{jt}, and the liquidity risk, L_{jt}, while the international risk is quantified by the common Eurozone risk, E_t, and the global financial risk, G_t.

In more detail, the sovereign bond spread is defined as the difference between the 10-year government bond yield in country j and the German government bond yield. All series come from Datastream (see Figure A1.5 in Appendix 1). We construct for each country j the default risk as the difference between the 10-year Credit Default Swap (CDS) premia in country j and the 10-year German CDS premia (all series come from Datastream; see Figure A1.6 in Appendix 1). The liquidity risk for each country j is approximated by the difference between the bid-ask spread of the 10-year bond in country j and the bid-ask spread of the respective German bond (all series come from Datastream; see Figure A1.7 in Appendix 1). As in De Santis (2014), the Euro area common risk factor is identified by the difference between the return on the 10-year KfW (Kreditanstalt für Wiederaufbau) bond and the respective return on the 10-year German government bond (all series come from Datastream; see Figure A1.8 in Appendix 1). Finally, to capture global financial risk we use the Global Financial Stress Index constructed by the Bank of America Merrill Lynch Global Research Division (available from Bloomberg; see Figure A1.8 in Appendix 1).

For each country j, we estimate three different versions of the multivariate vector $W(t)$ depending on the news information source M_{kt} $(k = 1,2,3)$. The first specification contains the "Grexit" mentions in Twitter (that is, $M_{1t} = T_t$). The second specification contains the "Grexit" mentions in the traditional news outlets (that is, $M_{2t} = N_t$). The final specification contains the orthogonalized Twitter variable (that is, $M_{3t} = T_t^{\perp}$), after taking out any effect of N_t from T_t.[24] For all specifications and for every country j, we obtain the DPR statistics up to twenty horizons (or days) ahead by calculating bootstrapped p-values with 1000 replications.

---

[23] The lag-length in all specifications is determined based on the Schwarz Information Criterion.

[24] After regressing the Twitter variable on the traditional news media variable, the orthogonalized variable is obtained from the respective residuals.

The p-values are analytically presented in Figure 10.[25] Starting with Greece, the hypothesis of no-causality running from T_t to the Greek sovereign spread (red-line in Figure 10.a) is rejected at the conventional levels of significance (0.01, 0.05 and 0.1) up to the eighteenth horizon.[26] When the T_t is replaced by N_t (black-line in Figure 10.a), the rejection of the null hypothesis, at the 0.1 significance level, is verified only up to the sixth horizon (exception is the first horizon where the rejection takes place at the 0.05 significance level). Finally, in the case where the orthogonal Twitter variable $(T_t^{\perp})$ is used (dashed red line in Figure 10.a), the rejection of the null hypothesis at the conventional levels of significance persists up to the seventh horizon (exception is the fifth horizon). The testing results reveal that the effect of N_t on the Greek sovereign spread is more short-lived (6 days) compared to T_t (18 days), while T_t continues to cause the Greek spread for several horizons (7 days), even when it is orthogonal to N_t.[27]

To examine possible contagion effects from the news related to the Greek debt crisis towards the remaining countries, we perform the same testing procedure by replacing the Greek sovereign spread with each country's respective sovereign spread. The causality results indicate evidence of contagion, mainly in the case of Portugal. Moreover, in the case of Ireland the evidence is weak, while there is no evidence of causality for the remaining countries. In particular, for Portugal the hypothesis of no-causality running from the T_t to the sovereign spread (red line in Figure 10.d) is rejected at the conventional levels of significance for sixteen out of the twenty horizons. Similar inference is obtained for the N_t variable, with rejection occurring for nineteen out of the twenty periods (black line in Figure 10.d). Furthermore, for the $T_t^{\perp}$ variable, we fail to reject the null hypothesis at any

[25] The values of the DPR statistics for selected horizons are reported in Table A3.2 in Appendix 3.

[26] As the number of followers is available for each tweet, the followers weighted Twitter series (over the raw data) is expected to be a superior metric in explaining bond market movements. The reason for not adjusting each tweet with the respective followers, is the "apples-to-apples" comparison framework between the two news sources. If we were to implement the followers' adjustment for the tweets, an equal treatment approach would be needed to also adjust the traditional news by the respective number of readers (or even the viewers if the text source comes from a TV transcript). Unfortunately, to the best of our knowledge, such data is not available and therefore an equivalent adjustment for the traditional news series cannot be performed. In any case, we find that the followers' weighted Twitter series does not offer apparent superiority in explaining the Greek spreads or the traditional news over the unweighted series. Results based on the followers' weighted Twitter series, are available on request.

[27] Our results persist even when both series are scaled by their respective standard deviations. Additionally, the causal inference remains qualitatively similar when the disaggregated at country level (for selected countries) time series are used. These results are available on request.

horizon, implying that Twitter conveys no additional information relative to the traditional news outlets.

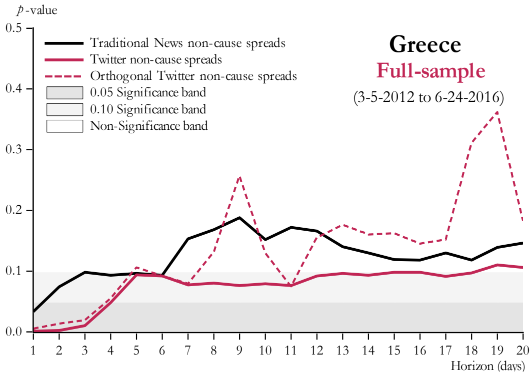

10.a. DPR causality tests p-values

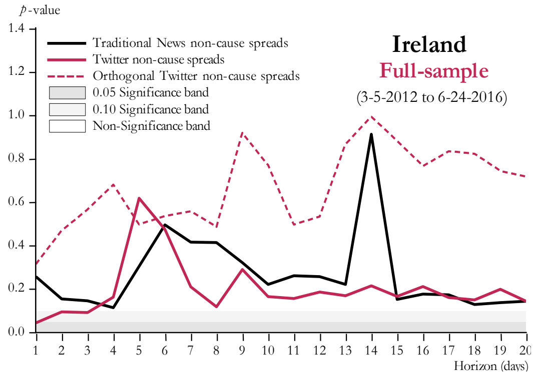

10.b. DPR causality tests p-values

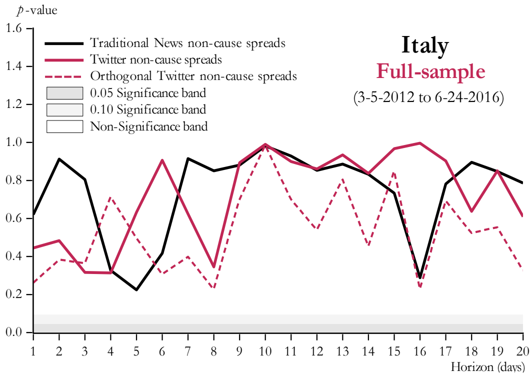

10.c DPR causality tests p-values

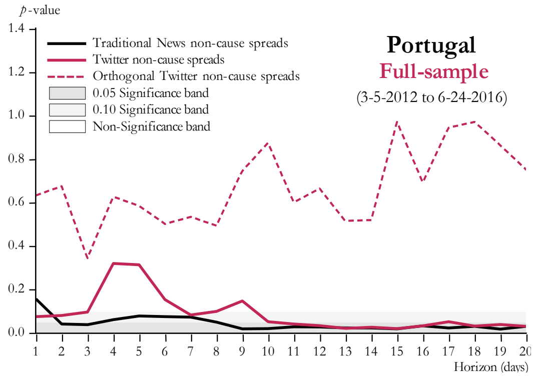

10.d DPR causality tests p-values

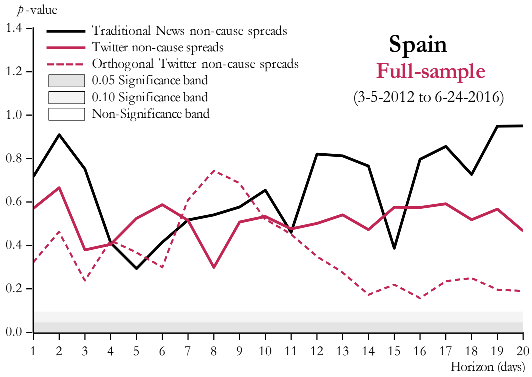

10.e DPR causality tests p-values

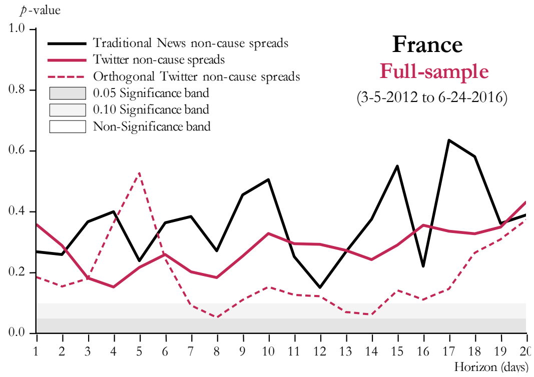

10.f DPR causality tests p-values

Figure 10. Full-sample DPR p-values for the GIIPS and France.

Focusing on Ireland, only the T_t variable appears to predict the Irish spread for the first three horizons (red line in Figure 10.b). In all other cases, we are unable to reject the no-predictability hypothesis. For the remaining three countries (that is, Italy, Spain and France), the derived inference is uniform; no effect on the spreads can be traced for any variable or horizon (Figures 10.c, 10.e and 10.f; a summary of the results is presented in Figure A2.2 in Appendix 2).[28, 29]

Moreover, we construct the twenty periods ahead impulse response trajectory of each country's spread, following a generalized one standard deviation shock on T_t, N_t and $T_t^{\perp}$. Figure 11 shows the transmission of these shocks to each country's spread. All impulse trajectories are accompanied with the 95% Scheffe´ confidence interval (grey area) and the 95% conditional confidence interval (blue area). At the bottom left hand-side of each graph, the twenty-period cumulative sum (Cum sum) of the responses and the resulting *p*-value for testing the significance (Cum *p*-value) are reported.

Starting with Greece, the spread trajectories are consistently positive no matter the origin of the shock $(T_t, N_t \text{ or } T_t^{\perp})$; at the same time, the respective conditional bands do not include zero (see Figures 11.a to 11.c). Focusing on Scheffe´s confidence band, the spread response to a T_t, N_t and $T_t^{\perp}$ shock is positive up to the seventh, third and fifth horizon, respectively (see Figures 11.a to 11.c). In addition, the cumulative impact of the responses is positive and significant at the 0.01 significance level in all three cases. The magnitude of the Greek sovereign spread response following a shock on T_t is an 11 basis points increase at the first horizon and a 329 basis points cumulative increase (see Figure 11.a). Following a shock on N_t, the spread increases by 8 basis points during the first horizon and by 215 basis points cumulatively (see Figure 11.b). Finally, the increase in the spread following a shock on $T_t^{\perp}$ is 7 basis points during the first horizon and 220 basis points cumulatively (see Figure 11.c).

[28] Exception is France, where sporadic rejections (for four horizons) take place only when the orthogonalized Twitter variable is used (red-dashed line in Figure 10.f). It is worth mentioning that the cumulative twenty periods impulse response is negative, suggesting reduction in the spread (see Figure 11.r below) which points to flight-to-safety considerations.

[29] As an additional validation of the agenda setting hypothesis, we have executed an exercise focusing on the stock markets of GIIPS and France. Acting within the same methodological framework, we find that the results are not qualitative any different from the respective findings in the sovereign bond markets. We intend to return to this issue in a separate academic paper.

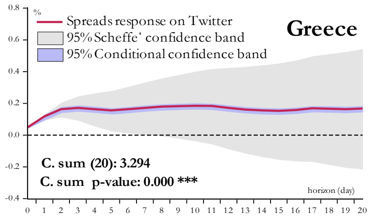

11.a S_t response on T_t **11.b.** S_t response on N_t **11.c.** S_t response on $T_t^{\perp}$

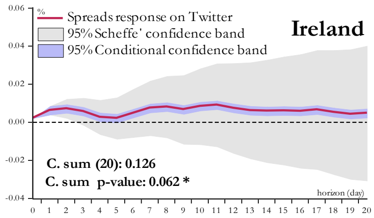

11.d. S_t response on T_t **11.e.** S_t response on N_t **11.f.** S_t response on $T_t^{\perp}$

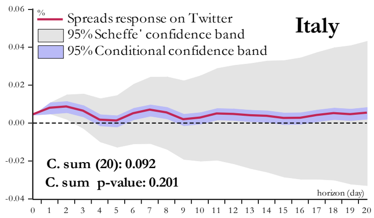

11.g. S_t response on T_t **11.h.** S_t response on N_t **11.i.** S_t response on $T_t^{\perp}$

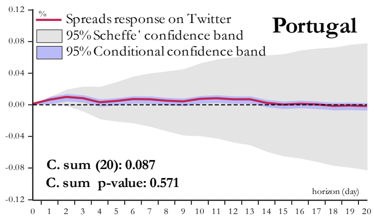

11.j. S_t response on T_t **11.k.** S_t response on N_t **11.l.** S_t response on $T_t^{\perp}$

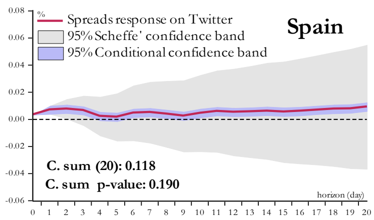

11.m. S_t response on T_t **11.n.** S_t response on N_t **11.o.** S_t response on $T_t^{\perp}$

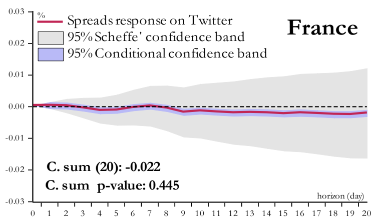

11.p. S_t response on T_t **11.q.** S_t response on N_t **11.r.** S_t response on $T_t^{\perp}$

Figure 11. Impulse responses.

Notes : C. sum (20) refers to the twenty-horizon (days) cumulative sum of the impulse responses, while C. sum p-value is the resulting p-value for testing the hypothesis that the C. sum (20) is equal to zero. Finally, *, ** and *** denote rejection of the null hypothesis at the 0.1, 0.05 and 0.01 significance level, respectively.

The impulse response analysis for the rest of the countries confirms the derived causal inference discussed above. For the case of Italy and Spain, although the impulse trajectories in all occasions are primarily positive, the Scheffe´s confidence bands embrace zero immediately after the second horizon (see Figures 11.g to 11.i and Figures 11.m to 11.o, respectively). Moreover, for both countries and all instances, the cumulative effects (ranging between 1 and 12.9 basis points) are statistically insignificant. In Ireland, (see Figures 11.d and 11.e), both sources of news dissemination affect the spread in a comparable fashion as: (i) the impulse trajectories evolve similarly; (ii) the Scheffe´s confidence bands after the second horizon imply insignificance and (iii) the $T_t^{\perp}$ delivers responses that are indistinguishable from zero. Finally, the cumulative impact of T_t on the Irish spread is 12.5 basis points (being significant at the 0.1 significance level), while the respective impact of N_t is 7.9 basis points (being insignificant).

For Portugal and after the first horizon, the Scheffe´s confidence bands support, in all cases, responses that wiggle around zero (see Figures 11.j to 11.l). Moreover, all cumulative effects on the Portuguese sovereign spread are insignificant. Finally, France is the only country where all the derived trajectories for the spread are in principle negative, although not different from zero (see Figures 11.p to 11.r). The twenty periods cumulative impact of both news sources is negligible and insignificant (-2.2 and -1.9 basis points), while significance occurs only after a shock on $T_t^{\perp}$ (see Figure 11.r and footnote 28).

Overall, we may argue that the impact of Twitter on the Greek sovereign spread is positive and of higher magnitude than that of the traditional news outlets (both appear to significantly affect the Greek spread in the short run). Further, the predictive power of Twitter still persists even when we account for the effects of the traditional news outlets. Finally, the combined inference from the causality testing and the impulse response analysis indicate some weak contagion effects in the case of Portugal and Ireland.

6. Discussion of results and conclusions

This paper considers the relationship between social media (Twitter in particular) and traditional news with an application to the Eurozone bond market to reach the following conclusions. First, there is a bidirectional information flow between Twitter and traditional news outlets, suggesting not only that both types of news serve as important empirical predictors for the sovereign bond market, but also that the ‘old’ (traditional news) and the ‘new’ (Twitter) media are connected. In addition, the impact of Twitter on the traditional news is more prolonged, stronger and more robust in terms of significance, which points to the dominance of Twitter (over traditional media) as the agenda-setting news source in the context of the European sovereign bonds market. Second, the impact of Twitter's "Grexit" mentions on the Greek sovereign spread is positive and of higher magnitude than that of the traditional news outlets; in addition, the predictive power of Twitter persists even when we take out the effects of the traditional news (by orthogonalizing the Twitter variable on the traditional news variable). Third, our analysis shows weak contagion effects from the informational content in Twitter and traditional news for the case of Portugal and Ireland.

Our evidence of weak contagion effects might be related to the exposure of banks to Greek public and private debt. Recall that we find some evidence mainly for Portugal[30] and Ireland. Figure 12 reports Bank of International Settlements (BIS) data which shows that prior to the crisis, Portuguese banks had the highest exposure to Greek public and private debt (reaching 6.79% of their total exposure around the world in 2010Q1). In terms of timing, we observe that Irish banks decided to reduce notably their exposure to Greek debt earlier than everybody else in 2010Q4, that is, when Ireland itself was bailed-out for €85bn.[31]

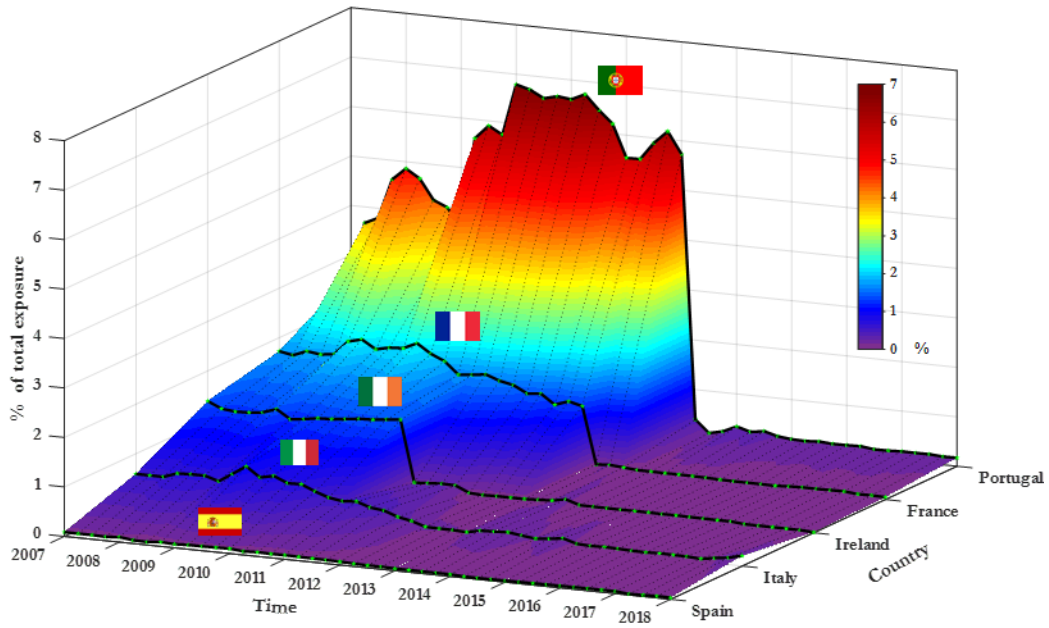

Figure 12. Exposure of IIPS and France to Greek debt.

Notes: The data for the exposure of the IIPS banks to the Greek public and private debt (% of their total exposure around the world) come from the Bank of International Settlements (BIS) and cover the period 2007Q1 to 2018Q1.

[30] Using a composite indicator that measures multidimensional sovereign bond market stress in the euro area from September 2000 to August 2018, Garcia-de-Andoain, and Kremer (2018) find spill-over effects from Greece to Portugal and vice versa.

[31] See: https://www.bbc.co.uk/news/av/uk-northern-ireland-11859578/85-billion-euro-bail-out-agreed-for-irish-republic.

Our results highlight the importance of social media platforms, Twitter in particular, in predicting bond market movements over and above the impact of economic/financial fundamentals and more so, compared to traditional news. The dominance of Twitter over traditional media in affecting bond spreads suggests that financial markets are affected through the transmission information channel and more so by Twitter.

Let us elaborate on this. Although crowd-sourced media, like Twitter, provide valuable signals for pricing assets in financial markets, these very signals may be blurred by the 'many-to-many' principle. Almost any individual can create and share content in real time (continuous information flow) which implies that Twitter, as a source, has a low signal-to-noise ratio. Hence, to disentangle the signal from noise, a considerable information processing capacity is essential. On the other hand, as traditional information sources deliver more curated news in discrete time intervals (discrete information flow), the signal-to-noise ratio increases compared to Twitter and thus, signal extraction demands lower processing capacity. Given the institutional investors' ability to collect and process a large amount of raw data from multiple sources in real time fast and effectively, it is reasonable to assume that their processing capacity is, on average, superior to that of individual investors. It is thus harder for the latter type of investors to process effectively signals arriving from low signal-to-noise ratio sources. Hence, investors located at the lower percentiles of the investors' processing capacity distribution are expected to draw information primarily from sources with higher signal-to-noise ratio. This suggests that information sources with low signal-to-noise ratio and continuous transmission frequency benefit investors with high processing capacity, while sources with high signal-to-noise ratio and discrete transmission frequency benefit mainly individual investors.

Under the realistic assumption of a heterogeneous information processing capacity among the different types of investors, a non-uniform impact on the pricing of the underlying asset is expected for each news source, even in cases where they disseminate the same information. For instance, consider that information with valuable pricing content (whether good or bad news) is released. Typically, this signal is first broadcasted by news sources with continuous transmission frequency (e.g. Twitter) and attracts the immediate attention of all market participants (individual and institutional investors). But does this attention imply instantaneous reaction? On average, we expect that investors need time to process the content of the information before updating their information set and reacting accordingly. Since institutional investors have the capacity to process the information faster, they will be the first ones to react. As the initially broadcasted information is reproduced by sources with discrete transmission frequency and relatively higher signal-to-noise ratio (e.g. newspapers), individual investors will then also process the signal and eventually react.

The above reasoning suggests that institutional investors react, on average, to news coming from all types of sources (continuous/discrete transmission and low/high signal-to-noise ratio). On the other hand, retail investors mainly react to news arriving from sources with discrete transmission frequency and relatively higher signal-to-noise ratio due to their lower information processing capacity. Overall, given that institutional investors account for a majority of the transactions in the bond market, their reaction is expected to have a greater impact on bond spreads than that of retail investors, therefore explaining the stronger influence of crowd-sourced media news.[32]

It is also important to note that by considering the instantaneous manner in which social media information is spread, social media contribute to the efficient functioning of financial markets. Unless, of course, misinformation finds its way through social media platforms. In a recent paper, for instance, Fan *et al.* (2019) find that automated Twitter accounts (known as 'bots') can pump out messages that have the ability to affect public opinion and the stock market. This raises the important issue of regulating social media. The responsibility lies with media companies, journalists, and governments. Worldwide, there seems to be consensus among consumers that media businesses, journalists and companies like Google or Facebook need to do more to combat misinformation. When it comes to government intervention, however, a mixed picture emerges, with sentiment towards government intervention being stronger in Europe than in the US. This raises the issue of how effective government intervention might turn out to be, in the absence of coordinated governmental actions across the world.[33]

References

Agrawal, S., P. Azar, A. Lo and T. Singh (2018). Momentum, Mean-Reversion and Social Media: Evidence from StockTwits and Twitter. Journal of Portfolio Management 44, 85-95.

Altavilla, C., M. Pagano and S. Simonelli (2017). Bank Exposures and Sovereign Stress Transmission. Review of Finance 21, 2103-2139.

[32] According to ECB data (https://www.ecb.europa.eu/press/pr/date/2017/html/pr170202_1.en.html), monetary and financial institutions held close to €5 trillion of Euro-area debt over the 2013-2016 period. Over the same period, households and non-profit institutions held about €1 trillion of Euro-area debt.

[33] According to The Reuters Institute Digital News Report 2018, 60% of responders in Europe favor increased government intervention compared to 41% in the US (see: http://www.digitalnewsreport.org/).

Araujo, T. and T.G. van der Meer (2018). News Values on Social Media: Exploring What Drives Peaks in User Activity About Organizations on Twitter. Journalism, November Issue, 1-19.

Augustin, P., H. Boustanifar, J. Breckenfelder and J. Schnitzler (2018). Sovereign to Corporate Risk Spillovers. Journal of Money, Credit and Banking 50, 857-891.

Azar P. and A. Lo (2016). The Wisdom of Twitter Crowds: Predicting Stock Market Reactions to FOMC Meetings via Twitter Feeds. Journal of Portfolio Management 42, 123-134.

Barberá, P., A. Casas, J. Nagler, P.J. Egan, R. Bonneau, J.T. Jost and J.A. Tucker (2019). Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data. American Political Science Review 113, 883-901.

Beckers, S. (2019). Do Social Media Trump News? The Relative Importance of Social Media and News Based Sentiment for Market Timing. Journal of Portfolio Management 45, 58-67.

Bedendo, M. and P. Colla (2015). Sovereign and Corporate Credit Risk: Evidence from the Eurozone. Journal of Corporate Finance 33, 34-52.

Ben-Rephael, A., Z. Da and R.D. Israelsen (2017). It Depends on Where You Search: Institutional Investor Attention and Underreaction to News. The Review of Financial Studies 30, 3009-3047.

Boudoukh, J., R. Feldman, S. Kogan and M. Richardson (2019). Information, Trading, and Volatility: Evidence from Firm-Specific News. The Review of Financial Studies 32, 992-1033.

Boutet, A., H. Kim and E. Yoneki (2012). What's in Your Tweets? I Know Who you Supported in the UK 2010 General Election. In: Proceedings of the International AAAI Conference on Weblogs and Social Media (ICWSM), 2012. From: https://hal.archives-ouvertes.fr/hal-00702390/

Chadwick, A. (2013). The Hybrid Media System: Politics and Power. New York, NY: Oxford University Press.

Chen, H., P. De, Y. Hu and B-H Hwang (2014). Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media. The Review of Financial Studies 27, 1367-1403.

Conway, B.A., K. Kenski and D. Wang (2015). The Rise of Twitter in the Political Campaign: Searching for Intermedia Agenda-setting Effects in the Presidential Primary. Journal of Computer-Mediated Communication 20, 363-380.

Dergiades, T., C. Milas and T. Panagiotidis (2015). Tweets, Google Trends and Sovereign Spreads in the GIIPS. Oxford Economic Papers 67, 406-432.

De Santis, R.A. (2014). The Euro Area Sovereign Debt Crisis: Identifying Flight-to-Liquidity and the Spillover Mechanisms. Journal of Empirical Finance 26, 150-70.

Dufour, J.M., D. Pelletier and E. Renault (2006). Short-run and Long-run Causality in Time Series: Inference. Journal of Econometrics 132, 337-362.

Dufour, J.M. and E. Renault (1998). Short-run and Long-run Causality in Time Series: Theory. Econometrica 66, 1099-1125.

Fan, R., O. Talavera and V. Tran (2019). Social Media Bots and Stock Markets. European Financial Management (early view) 1-25.

Garcia-de-Andoain, C. and M. Kremer (2018). Beyond Spreads: Measuring Sovereign Market Stress in the Euro Area. European Central Bank Working Paper No. 2185.

Gayo-Avello, D. (2013). A Meta-analysis of State-of-the-art Electoral Prediction from Twitter Data. Social Science Computer Review 31, 649-679.

Gennaioli, N., A. Martin and S. Rossi (2014). Sovereign Default, Domestic Banks, and Financial Institutions. Journal of Finance 69, 819-866.

Gennaioli, N., A. Martin and S. Rossi (2018). Banks, Government Bonds, and Default: What do the Data Say? Journal of Monetary Economics 98, 98-113.

He, Z. (2017). The Contractionary Effect of Bad Economic News. Journal of Money, Credit and Banking 49, 1339-1384.

Heredia, B., J.D. Prusa and T.M. Khoshgoftaar (2018). Social Media for Polling and Predicting United States Election Outcome. Social Network Analysis and Mining, December Issue, 8-48.

Johansen, S. (1991). Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models. Econometrica 59, 1551-1580.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95, 161-182.

Jordà, Ò. (2009). Simultaneous Confidence Regions for Impulse Responses. Review of Economics and Statistics 91, 629-647.

Jungherr, A. (2014). The Logic of Political Coverage on Twitter: Temporal Dynamics and Content. Journal of Communication 64, 239-259.

Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux Publishers, New York.

Kilian, L. and Y. Kim (2009). Do Local Projections Solve the Bias Problem in Impulse Response Inference? CEPR Discussion Paper Series 7266.

Kruikemeier, S., K. Gattermann and R. Vliegenthart (2018). Understanding the dynamics of politicians' visibility in traditional and social media. The Information Society 34, 215-228.

Kümpel, A.S., V. Karnowski and T. Keyling (2015). News Sharing in Social Media: A Review of Current Research on News Sharing Users, Content, and Networks. Social Media and Society 2, 1-14.

MacKinnon, J.G., A.A. Haug, and L. Michelis (1999). Numerical distribution functions of likelihood ratio tests for cointegration. Journal of Applied Econometrics 14, 563–577.

Meyer K.M. and T. Tang (2015). #SocialJournalism: Local News Media on Twitter. International Journal on Media Management 17, 142-257.

Mitchell, A., T. Rosenstiel, and L. Christian. What Facebook and Twitter Mean for News, Pew Research Center Report, March 2012.

Moon, S.J. and P. Hadley (2014). Routinizing a New Technology in the Newsroom: Twitter as a News Source in Mainstream Media. Journal of Broadcasting and Electronic Media 58, 289-305.

Neuman, R.W., L. Guggenheim, S.M. Jang and S.Y. Bae (2014). The Dynamics of Public Attention: Agenda-Setting Theory Meets Big Data. Journal of Communication 64, 193-214.

Shearer, E. and J. Gottfried (2017). News Use Across Social Media Platforms 2017. Pew Research Center Report, September 2017.

Schroeder, R. (2018a). Towards a Theory of Digital Media. Information, Communication & Society 21, 323-339.

Schroeder, R. (2018b). Rethinking Digital Media and Political Change. Convergence: The International Journal of Research into New Media Technologies 24, 168-183.

Soo, C.K (2018). Quantifying Sentiment with News Media across Local Housing Markets. The Review of Financial Studies 31, 3689-3719.

Sprenger, T.O., A. Tumasjan, P.G. Sandner and I.M. Welpe (2013). Tweets and Trades: The Information Content of Stock Microblogs. European Financial Management 20, 926-957.

Tetlock, P.C. (2015). The Role of Media in Finance. In Handbook of Media Economics, edited by S.P. Anderson, J. Waldfogel and D. Strömberg, Chapter 18, 701-721. Amsterdam: North-Holland.

Toda, H.Y. and T. Yamamoto (1995). Statistical inference in vector autoregressions with possibly integrated processes. Journal of Econometrics 66, 225-250.

Tumasjan, A., T.O. Sprenger, P.G. Sandner and I.M. Welpe (2010). Predicting Elections with Twitter: What 140 Characters Reveal About Political Sentiment. International AAAI Conference on Weblogs and Social Media (ICWSM) 10, 178-185.

Turner, J.D., Q. Ye and C.B. Walker (2018). Media Coverage and Stock Returns on the London Stock Exchange, 1825-70. Review of Finance 22, 1605-1629.

Wolski, M. (2018). Sovereign Risk and Corporate Cost of Borrowing: Evidence from a Counterfactual Study. EIB Working Papers 2018/05, European Investment Bank (EIB).

Zhao, W.X., J. Jiang, J. Weng, J. He, E-P. Lim, H. Yan and X. Li (2011). Comparing Twitter and Traditional Media using Topic Models. European Conference on Information Retrieval ECIR 2011: Advances in Information Retrieval, 338-349.

Appendix 1. Data

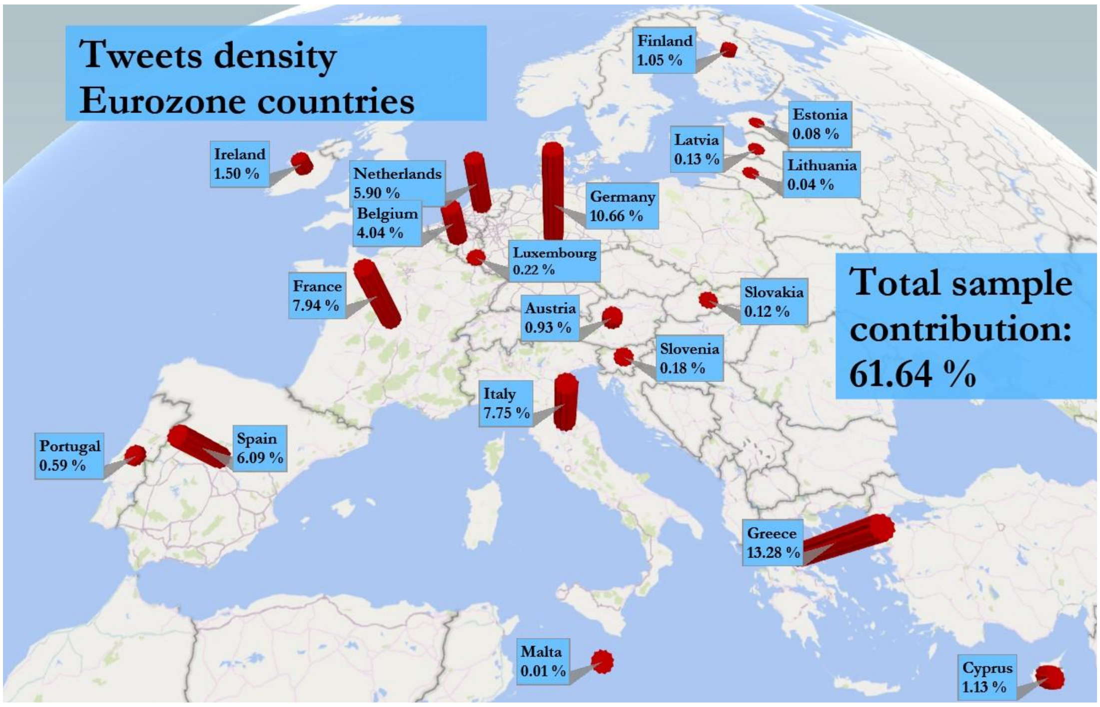

Figure A1.1. Tweets density in Eurozone (country disaggregation level).

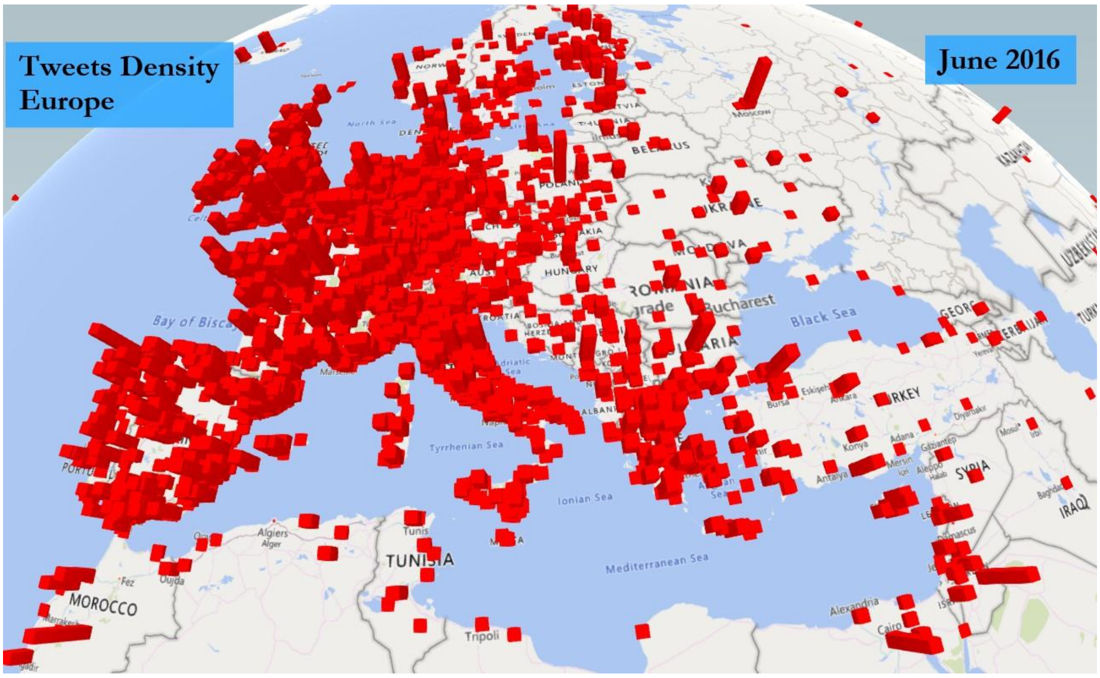

Figure A1.2. Tweets density in Europe (at the higher disaggregation level).

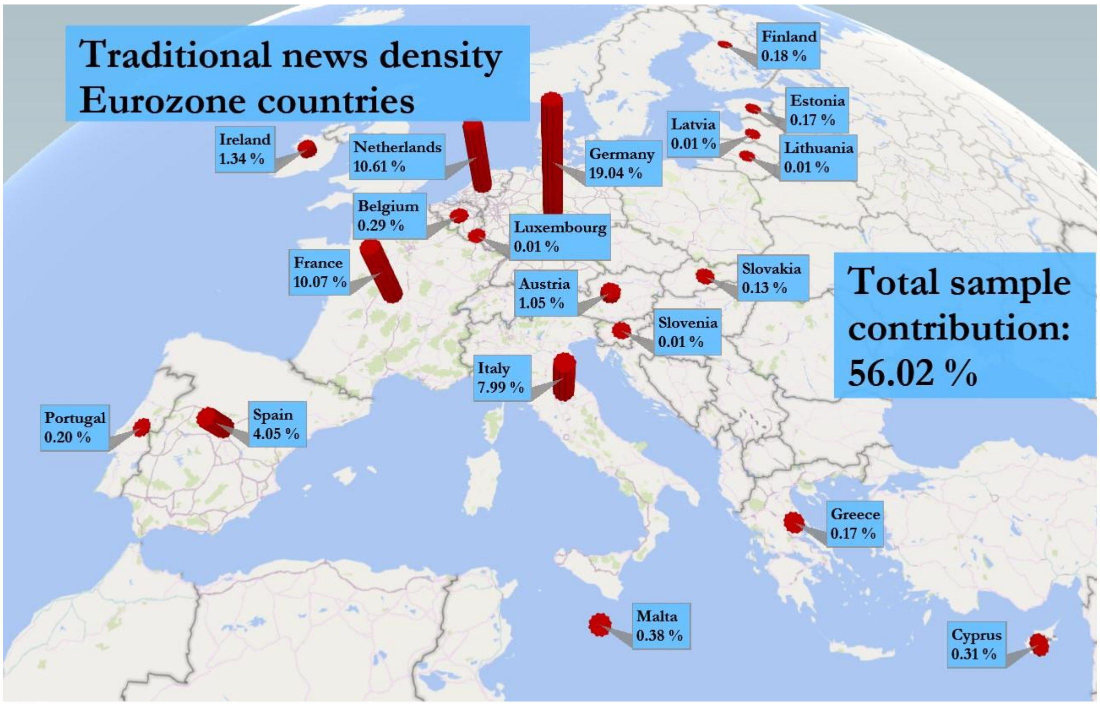

Figure A1.3. Traditional news density in Eurozone (country disaggregation level).

Between correlation ↘	**Worldwide**	Germany	UK	Italy	Netherlands	France	Australia	Ireland	US	Spain	Canada	All other
Worldwide	**0.89**	0.90	0.87	0.88	0.87	0.92	0.82	0.79	0.90	0.89	0.76	0.95
Germany	0.97	**0.77**	0.72	0.85	0.86	0.90	0.68	0.71	0.74	0.85	0.64	0.87
UK	0.96	0.95	**0.67**	0.75	0.70	0.76	0.74	0.74	0.83	0.75	0.70	0.84
Italy	0.96	0.94	0.96	**0.75**	0.75	0.92	0.69	0.71	0.78	0.86	0.63	0.92
Netherlands	0.96	0.95	0.96	0.95	**0.62**	0.80	0.72	0.68	0.72	0.78	0.70	0.81
France	0.95	0.95	0.92	0.94	0.94	**0.84**	0.71	0.72	0.79	0.90	0.65	0.95
Australia	0.95	0.91	0.96	0.95	0.92	0.89	**0.66**	0.74	0.77	0.78	0.74	0.79
Ireland	0.93	0.96	0.95	0.93	0.96	0.93	0.89	**0.54**	0.71	0.73	0.65	0.78
US	0.92	0.91	0.84	0.90	0.87	0.92	0.83	0.87	**0.63**	0.76	0.76	0.88
Spain	0.90	0.92	0.91	0.90	0.95	0.88	0.84	0.93	0.80	**0.60**	0.66	0.89
Canada	0.89	0.84	0.79	0.87	0.83	0.85	0.80	0.80	0.96	0.76	**0.54**	0.73
All other	0.97	0.97	0.95	0.94	0.95	0.97	0.91	0.94	0.90	0.90	0.84	**0.82**

Wihtin correlation (Traditional news)

Within correlation (Twitter)

Figure A1.4. Within and between correlations by country.

Notes: The **blue area** encloses within correlations for the traditional news (e.g. the value at line 1 and column 2 suggests that the traditional news time series of Germany has a correlation of 0.90 with the worldwide traditional news time series). The **green area** encloses within correlations for Twitter (e.g. the value at line 8 and column 3, suggests that the Twitter time series of Ireland has a correlation of 0.95 with the Twitter time series of UK). The **black** main diagonal illustrates the between correlations for the two news sources. (e.g. the value at line 4 and column 4, suggests that the Twitter and the traditional news time series of Italy have a correlation of 0.75.

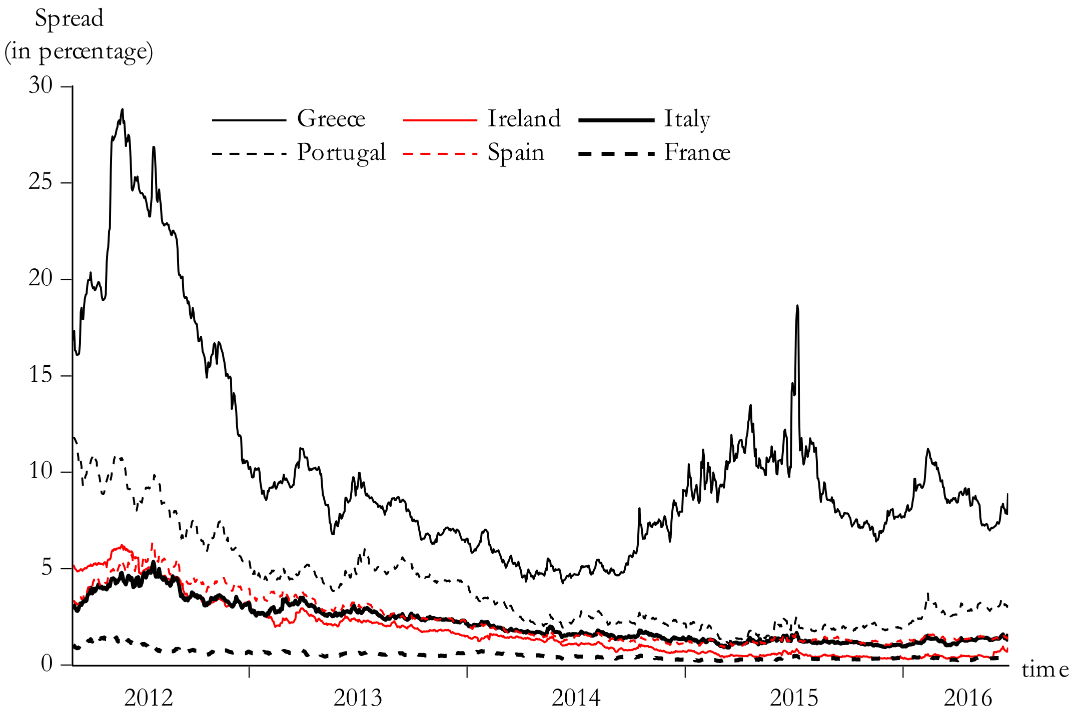

Figure A1.5. Sovereign Spreads for GIIPS and France

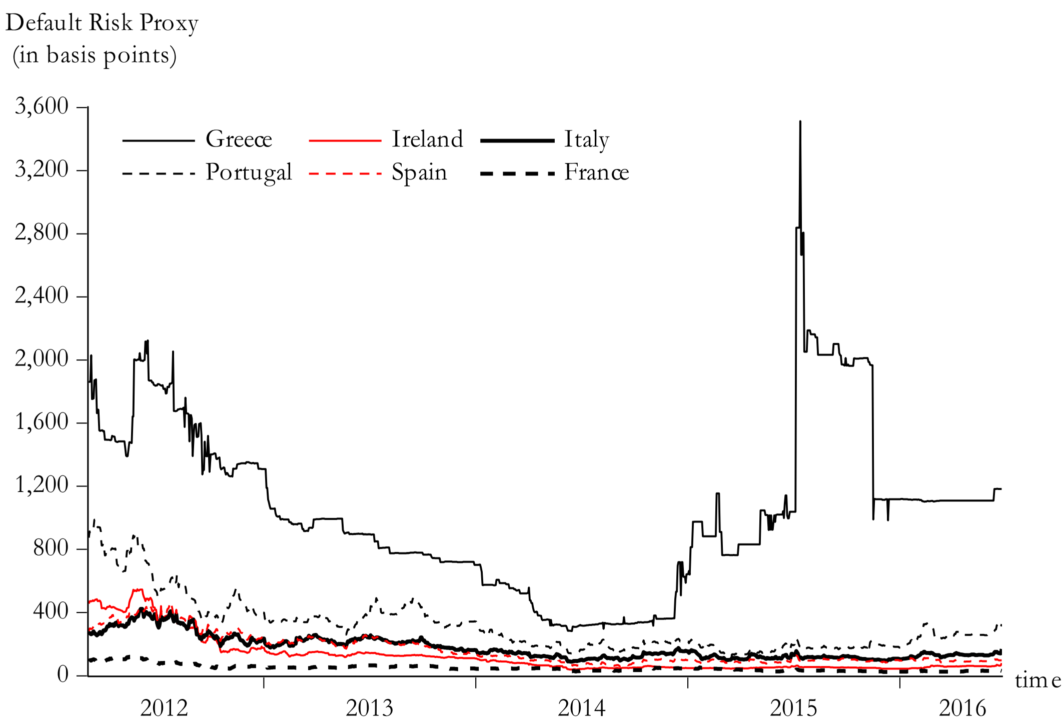

Figure A1.6. Default risk proxy for GIIPS and France

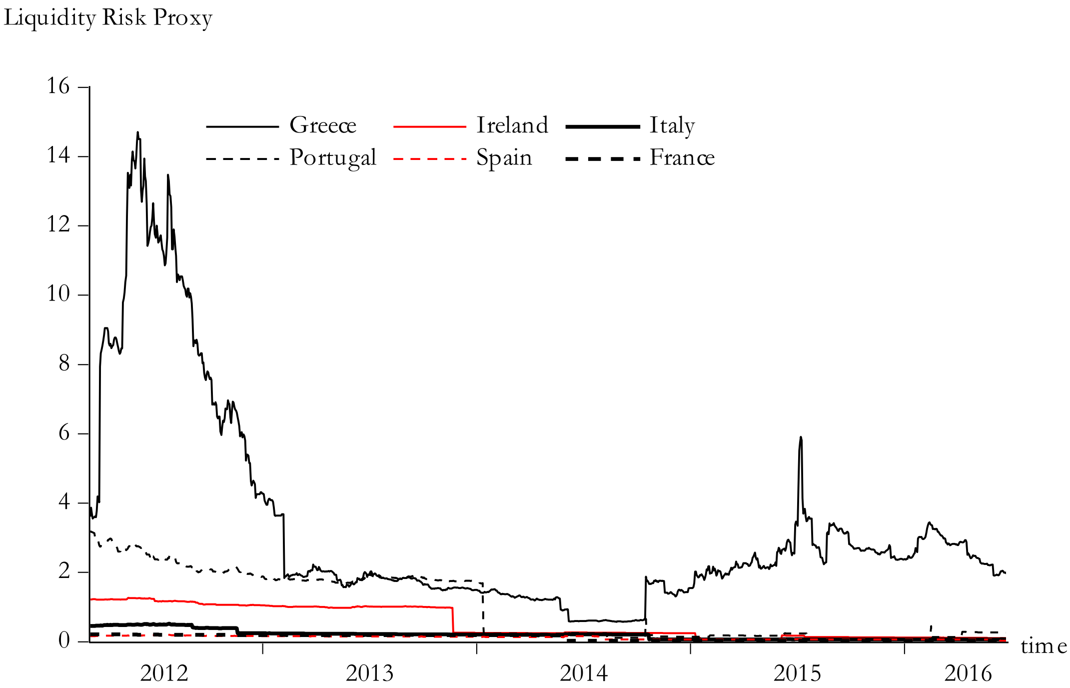

Figure A1.7. Liquidity risk proxy for GIIPS and France

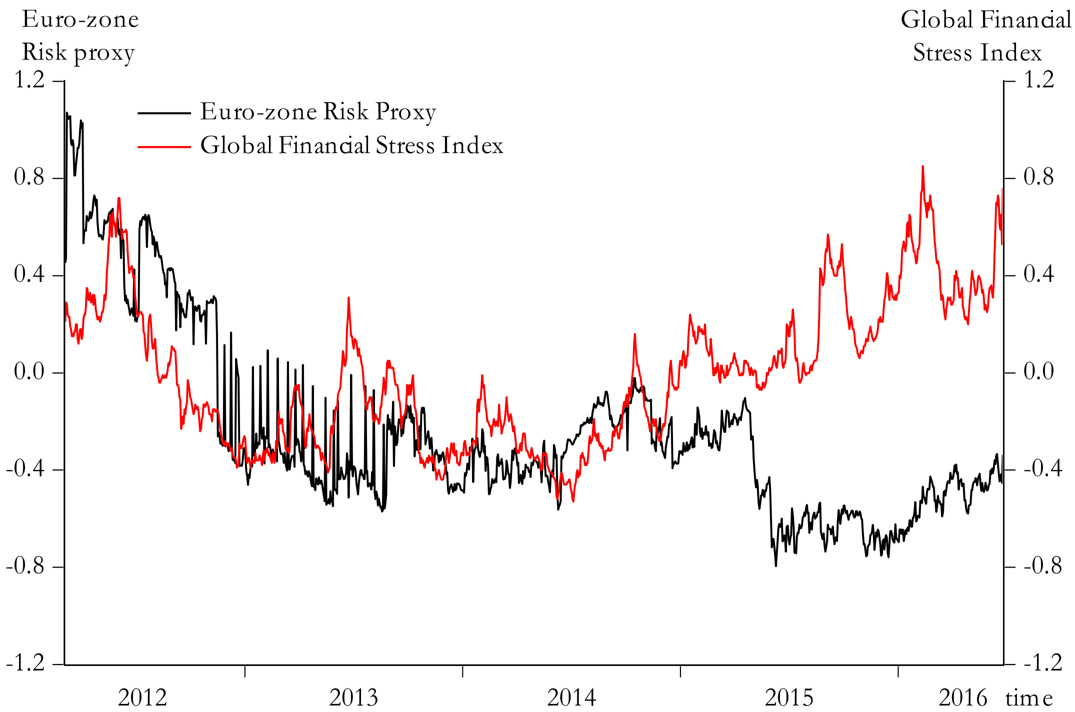

Figure A1.8. Eurozone risk proxy and Global Financial Stress Index

Notes: The Euro area common risk proxy is identified by the difference between the return of the 10-year KfW (Kreditanstalt für Wiederaufbau) bond and the respective return of the 10-year German government bond. The Global Financial Stress Index is provided by the Bank of America Merrill Lynch Global Research Division.

Appendix 2. Summary of the results

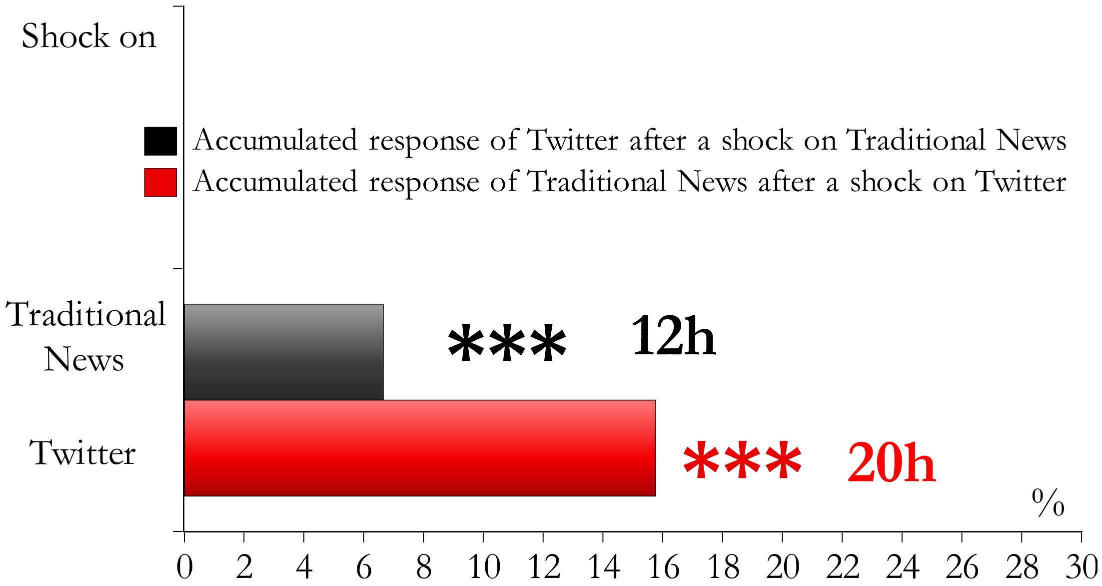

Figure A2.1. Summary results for section 5.1.

Notes: The bar length shows the twenty-periods cumulative response of one news source after a shock on the other news source; *** indicate the 0.01 significance of the cumulative response; finally, the number right before the letter h, imply the number of horizons for which the null hypothesis of no causality is rejected at the conventional levels of significance.

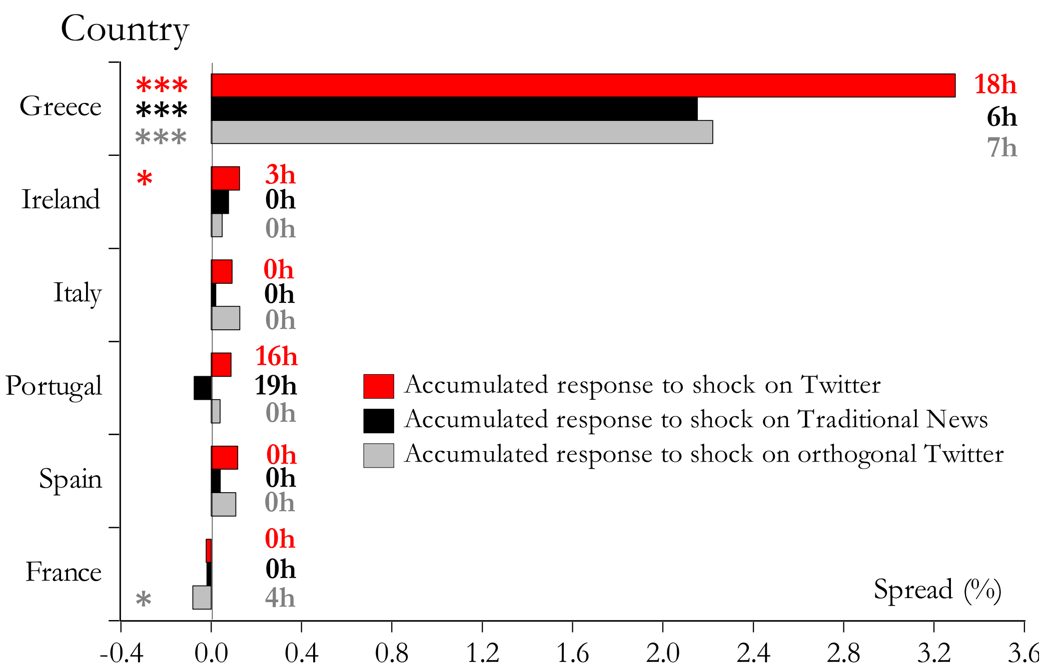

Figure A2.2. Summary results for section 5.2.

Notes: The bar length shows the twenty-periods cumulative response of the spreads after a shock on the respective news source; ***, ** and * indicate the 0.01, 0.05 and 0.1 significance of the cumulative response; finally, the number right before the letter h, imply the number of horizons for which the null hypothesis of no causality is rejected at the conventional levels of significance.

Appendix 3. Causality testing results

Table A3.1. Twitter activity predictability from the respective activity on Traditional news outlets and vice versa (Sample: 3-5-12 to 6-24-16).

Panel A

Sample period	Hypothesis	Horizon									
		h = 1	h = 2	h = 3	h = 4	h = 5	h = 6	h = 7	h = 8	h = 9	h = 10
Full-sample (3-5-12 to 6-24-16)	$T_t \nrightarrow N_t$	160.02***	115.22***	84.02***	83.39***	57.27***	51.30***	59.96***	59.86***	54.24***	49.51***
	$N_t \nrightarrow T_t$	23.66***	23.28***	15.84**	14.39*	12.31*	22.82***	16.90**	14.87**	14.35**	9.54

Panel B

Sample period	Hypothesis	Horizon									
		h = 11	h = 12	h = 13	h = 14	h = 15	h = 16	h = 17	h = 18	h = 19	h = 20
Full-sample (3-5-12 to 6-24-16)	$T_t \nrightarrow N_t$	44.68***	36.37***	36.62***	34.87***	39.30***	36.85***	31.87***	34.65***	30.66***	23.38**
	$N_t \nrightarrow T_t$	12.03	12.45	15.57**	14.15*	10.16	12.11	16.51**	13.29	9.91	10.78

Notes: Dufour *et al.* (2006) Wald statistics (DPR) are reported. The symbol $\nrightarrow$ denotes the null hypothesis of non-causality that runs from the left-hand variable to the right-hand variable. T is the logarithm of the Grexit mentions in Twitter and N is the logarithm of the Grexit mentions in the traditional news outlets. The sequence of stars (*, ** and ***), signify rejection of the null hypothesis at the 0.1, 0.05 and 0.01 significance level, respectively.

Table A3.2. Sovereign spreads predictability in the GIIPS from Twitter and Traditional news (Sample: 3-5-12 to 6-24-16).

Country	Hypothesis	Horizon									
		$h=1$	$h=2$	$h=3$	$h=4$	$h=5$	$h=6$	$h=7$	$h=10$	$h=15$	$h=20$
Greece	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	24.32***	19.61***	14.00***	8.42***	6.16***	6.18***	6.55***	6.91***	6.30***	6.61***
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	8.56**	7.21*	6.53*	6.26*	6.07*	5.88*	5.11	5.07	5.74	5.67
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	12.63***	11.85***	10.48**	6.74*	5.11	5.89*	5.70*	4.99	4.50	4.33
Ireland	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	7.85**	6.10*	5.46*	3.95	0.60	1.13	2.96	4.11	3.93	4.52
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	2.94	4.62	4.40	4.28	2.41	1.14	1.52	3.51	4.40	4.98
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	3.80	2.35	1.61	1.29	1.95	1.73	1.63	0.88	0.52	1.09
Italy	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	2.75	2.43	3.46	3.26	1.45	0.41	1.32	0.14	0.18	1.76
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	1.85	0.55	0.89	3.05	4.17	2.56	0.41	0.16	1.15	1.12
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	4.08	2.93	2.94	1.19	1.91	2.76	2.41	0.14	0.62	3.40
Portugal	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	7.16*	7.56*	6.45*	3.22	3.32	4.82	5.92*	8.08**	9.58**	11.52**
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	5.39	8.78**	9.10**	7.72*	6.28*	6.52*	6.95*	11.62**	12.05**	11.99**
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	1.81	1.54	3.03	1.46	1.55	1.96	1.64	0.51	0.22	0.95
Spain	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	2.06	1.67	3.02	2.72	1.91	1.65	1.93	1.88	1.84	2.56
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	1.36	0.53	1.03	2.74	3.30	2.46	1.65	1.38	2.65	0.33
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	3.56	2.52	3.70	2.40	2.59	2.86	1.18	1.70	3.75	4.41
France	$T_t \nrightarrow S_t \mid_{E.G.L.D.}$	2.25	2.88	3.75	4.06	3.23	2.61	2.73	2.60	2.80	1.65
	$N_t \nrightarrow S_t \mid_{E.G.L.D.}$	3.20	3.25	1.99	1.66	2.84	1.87	1.75	1.09	0.82	2.06
	$T^{\perp} \nrightarrow S_t \mid_{E.G.L.D.}$	5.04	5.31	4.72	2.75	1.95	3.46	5.68	4.77	5.58	2.83

Notes: Dufour *et al.* (2006) Wald statistics (DPR) are reported. The symbol $\nrightarrow$ denotes the null hypothesis of non-causality that runs from the left-hand variable to the right-hand variable. T is the Grexit mentions in Twitter (logarithm), N is the Grexit mentions in the traditional news outlets (logarithm), $T_t^{\perp}$ is the orthogonal Twitter and S is the sovereign spreads. $E.G.L.D.$ suggests conditioning on euro-zone risk, global financial risk, liquidity risk and default risk. Finally, *, ** and *** signify rejection of the null hypothesis at the 0.1, 0.05 and 0.01 significance level, respectively.