VU Research Portal

The shifting discourse of the European Central Bank: Exploring structural space in semantic networks

Nerghes, A.; Lee, J.; Groenewegen, P.; Hellsten, I.

published in

Proceedings of the Tenth International Conference on Signal Image Technology & Internet Based Systems
2014

DOI (link to publisher)

[10.1109/SITIS.2014.13](https://doi.org/10.1109/SITIS.2014.13)

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Nerghes, A., Lee, J., Groenewegen, P., & Hellsten, I. (2014). The shifting discourse of the European Central Bank: Exploring structural space in semantic networks. In A. Dipanda, & K. Yetongnon (Eds.), *Proceedings of the Tenth International Conference on Signal Image Technology & Internet Based Systems* (pp. 447-455). IEEE Computer Society. <https://doi.org/10.1109/SITIS.2014.13>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

The shifting discourse of the European Central Bank: Exploring structural space in semantic networks

Adina Nerghes	Ju-Sung Lee	Peter Groenewegen	Iina Hellsten
Vrije Universiteit Amsterdam	University of Twente	Vrije Universiteit Amsterdam	Vrije Universiteit Amsterdam
Amsterdam, The Netherlands	Enschede, The Netherlands	Amsterdam, The Netherlands	Amsterdam, The Netherlands
adina.nerghes@vu.nl	j.s.lee-1@utwente.nl	p.groenewegen@vu.nl	i.r.hellsten@vu.nl

Abstract—Convenient access to vast and untapped collections of documents generated by organizations is a valuable resource for research. These documents (e.g., press releases, reports, speech transcriptions, etc.) are a window into organizational strategies, communication patterns, and organizational behavior. However, the analysis of such large document corpora does not come without challenges. Two of these challenges are 1) the need for appropriate automated methods for text mining and analysis and 2) the redundant and predictable nature of the formalized discourse contained in these collections of texts. Our article proposes an approach that performs well in overcoming these particular challenges for the analysis of documents related to the recent financial crisis. Using semantic network analysis and a combination of structural measures, we provide an approach that proves valuable for a more comprehensive analysis of large and complex semantic networks of formal discourse, such as the one of the European Central Bank (ECB). We find that identifying structural roles in the semantic network using centrality measures jointly reveals important discursive shifts in the goals of the ECB which would not be discovered under traditional text analysis approaches.

I. INTRODUCTION

The increasing availability of textual information opens new venues for large-scale research. In particular, numerous text documents are generated daily by organizations across the world regarding their activities and objectives. However, large corpora of such text documents are difficult to analyze without proper methods which are at least semi-automated.

Another challenge for research is that such texts are manifestations of highly formalized discourse, which is characterized as redundant, structured and even predictable [1]. Discourse ‘acts as a powerful ordering force in organizations’ [2] because meaning is negotiated in organizations, and these meanings shape organizational practices [3]. As ‘a carrier of power through its ability to order and constitute the social world’ [2], discourse harbors the potential to signal consequential information to other organizations and society in general. Its timely analysis may be crucial, yet it is often challenging.

We propose an approach for dealing with complex semantic networks generated from large text corpora of formal organizational discourse. More precisely, the method assesses dynamic discursive shifts in complex semantic networks, in an ample

and comprehensive manner.

Firstly, to analyze the large text corpora collected, we use semantic network analysis. Semantic network analysis is one of the areas of research that has gained popularity in recent years. This type of analysis maps networks of concepts (i.e., a concept being a word or multiple words) in the form of networks of meaning. Although language can be suitably represented as a network [4], semantic networks are often large and complex and exhibit highly intricate network structures at all levels [5]–[7]. Some posit these networks to exhibit stylized topologies such as *small-world* or *scale-free* [4], [6]–[9]. Such networks, however, provide insights into how language serves as a framework for representing and communicating information. The complexity of large semantic networks arises not only from the size of the corpora, but also from an array of global and local features, which in turn emerge from the structure of links between the concepts. In this paper, we use semi-automated coding of concepts to be included in the semantic networks [10]–[12].

Secondly, our paper proposes an approach for assessing dynamic shifts in formal discourse through the structural positions of semantic network nodes. The structural space method we propose in this paper combines two classic social network analysis structural measures to create four structural roles for network nodes. The two structural measures we employ are total degree centrality (i.e., popularity) and betweenness centrality (i.e., connectivity).

The idea of structural roles in social networks has been explored through various approaches over the years. A few examples would be structural holes [13], equivalence [14]–[17], blockmodels [18], [19], and role structure [20]. However, the identification of structural roles through the combination of structural measures has not been widely explored. One such effort comes from Carley and Kaufer [10], and it combines density, conductivity and consensus to explore connectivity in semantic networks. The paper of Huang et al. [21] proposes a combination of multiple strongly correlated social network analysis (SNA) metrics to evaluate only those top ranked nodes in undirected binary networks. For visualisations purposes, NodeXL offers the possibility of plotting nodes based on their

actual centrality scores but without identifying different roles [22]. The distinctive feature of our method is the identification of four structural roles based on the combination of two structural measures, and thus it is not merely focused on high ranking nodes.

The corpora used in this study comprises the press releases issued by the European Central Bank (henceforth ECB) between 2006 and 2013. The ECB determines the monetary policy for the whole euro area. Established by the Treaty of Amsterdam in 1 June 1998 [23], the ECB is the formal successor of the European Monetary Institute. As one of the seven institutions of the European Union, the ECB is the central bank for the euro and administers the monetary policy of the 17 EU member states, which constitute the Eurozone, one of the largest currency areas in the world.

The ECB distributes large volumes of information (e.g., policy deliberations, public speeches, annual reports etc.) as one of their key policy tools. Because the ECB's only formal instrument, through which they can exert an (indirect) effect on asset prices (of key importance for the economy), is the overnight interest rate, their communications become a powerful tool. These can impact developments in financial markets [24]–[27], directly influence private sector expectations, and are used to signal interest rate increases [28], [29]. The communications of the ECB also increase the predictability of interest rate decisions [30], being generally considered trustworthy and understandable by the public [31].

In the following section we describe our data set, and the approach we are proposing. Chapter 3 presents the results of our analyses. Finally, Chapter 4 summarizes our overall findings and discusses the benefits and limitations of our method.

II. DATA AND METHODS

For this study, 825 press releases issued by the ECB between January 2006 and December 2013 have been collected from their web archive. These press releases have been divided in four time periods each spanning a period of two years. The aggregation of data in these four periods was motivated by our aim of assessing the impact of the financial crisis on the ECB discourse. The first sub-sample, containing 184 text documents, covers the period just prior to the financial crisis: January 2006 until December 2007. We label this period *pre-crisis*.¹ The second sub-sample (*crisis*) includes the 203 press releases issued by the ECB between January 2008 and December 2009. The third sub-sample contains the 210 press releases issued by the ECB between January 2010 and December 2011, and represents the *post-crisis* period. Lastly, the fourth sub-sample includes 228 press releases issued by the ECB between January 2012 and December 2013, further referred to as *recovery*.

Each of the data samples (*pre-crisis*, *crisis*, *post-crisis*, and *recovery*) has been pre-processed with AutoMap [32]. The pre-

¹Although the global financial crisis emerged in the United States at the end of August 2007, we assume its effects were not visible in the European Union until the beginning of 2008.

processing removed all the noise words (e.g., numbers, verbs, extra spaces etc.) in the data and prepared it for the generation of semantic networks. Four semantic networks were generated using AutoMap, one for each sub-sample (see Table I for the descriptive statistics of each network). The generation of networks is based on Carley's approach to coding texts as cognitive maps [33] and Danowski's approach to proximity analysis [34]. Semantic networks translate text into networks of concepts and the links between them, where a concept can be a word or a phrase (i.e., n-gram) [35]. The links between concepts are formed based on co-occurrence. For example, if two words co-occur within the specified window size and stop unit, a link (or semantic network edge) will be formed. The window size and the stop unit determines the range in which connections will be made between words [36], while the stop unit determines the point where the window size ends. A window size of two and a stop unit of one sentence (used in this study) will create a link between each two consecutive words within the limits of one sentence. The value of strength for each link is determined by frequency of co-occurrence [37].

As mentioned above, a concept in our semantic networks can be a single word or an n-gram. N-grams are created by replacing the spaces between words with an underscore [32]. An example of such conversion is the phrase 'interest rate' which becomes 'interest_rate'. This procedure helps us identify the most common multi-word expressions used in text documents. Thus, when we refer to key concepts, we refer to single words as well as n-grams.

TABLE I: Descriptive statistics of each semantic network generated*

Measure	<i>Pre-crisis</i>	<i>Crisis</i>	<i>Post-crisis</i>	<i>Recovery</i>
Node count	580	628	648	755
Link count	200848	228874	265572	341806
Density	0.597	0.580	0.632	0.600
Density, Weighted	0.021	0.019	0.024	0.018
Clustering coefficient	0.784	0.775	0.794	0.781
Degree centralization	0.136	0.140	0.165	0.120

*Each network is undirected, symmetric and valued; Only nodes with frequencies ≥ 10 have been included in the networks.

The descriptive statistics of our networks (see Table I) show that even after employing a frequency threshold (≥ 10) the resulting networks are complex and dense with high link counts; in fact, the densities far exceed those of most human social networks. The combination of the complexity of these networks and the formal character of the documents from which they have been extracted poses a challenge for the analyst. To overcome this challenge, we propose the structural space method that considers total degree centrality and betweenness centrality of concepts in semantic networks, concurrently.

A. Centrality in networks

Even after decades of social network research, the current thinking about network centrality is still mostly defined by the work of Freeman [38] and Bonacich [39]. In 1977, Freeman developed a set of centrality measures based on betweenness

[40]. In a follow-up article two years later, Freeman [38] elaborates on three concepts of centrality in a social network, which have since been further developed into degree centrality, closeness centrality, and betweenness centrality. The fourth commonly used measure, eigenvector centrality, was studied by Bonacich [39]. We now define and briefly elaborate on (total) degree centrality and betweenness centrality, the focal measures in this paper.

1) *Total degree centrality*: Total degree centrality is one of the most commonly used centrality measures in social network analysis [41]. The degree centrality of a node in a network reflects the number of other nodes to whom the focal node is tied [38](or, in the case of weighted networks, the sum of the weights of all the links a node has), and thus measures the involvement of a node in its local network. Nodes with low total degree centrality are potentially more peripheral to the network [42], unless they are connected to popular others. In semantic networks, total degree centrality may represent the ‘importance’ of a concept or its key concept status. A key concept with high degree centrality is able to activate many other key concepts; thus, it functions as a hot topic’s central key concept [36]. Using only the local structure to calculate the degree centrality of a node, this measure does not take into consideration the position of the concept within the global structure of the network. In this paper, we employ the weighted version of total degree centrality.

2) *Betweenness centrality*: Betweenness centrality (C_B) is the sum of the proportions of the shortest paths a node lies on for every pair of nodes (out of all shortest paths for each pair). The formulation for unweighted betweenness is:

$$C_B(i) = \sum_{s \neq i \neq t} \frac{\sigma_{s,t}(i)}{\sigma_{s,t}}$$

where $\sigma_{s,t}$ indicates the count of shortest paths between nodes s and t . For weighted betweenness, the shortest paths are computed using the inverse of the edge weight since heavier edges should warrant greater flow (and higher betweenness). We employ this inversion as most of the edges between concepts are valued (i.e., weighted). More broadly, betweenness centrality represents the frequency with which a particular node is on the geodesic path between any other two nodes in the network [11]. As such, betweenness centrality captures one aspect of a node’s position in the graph, thus taking into account the global structure of the network. The betweenness centrality of a concept within a semantic network is a direct indicator of its influence [37], [43], [44]. A key concept with high betweenness centrality controls access to other key concepts in the network [41], [45]–[47], and thus serves as a gatekeeper between different domains [48]. For semantic networks, it is presumed that a node with high betweenness centrality has a higher likelihood to get activated or activate when connections across domains are activated.

B. Structural roles

By combining popularity and connectivity of concepts in semantic networks, we expect to capture emerging topics

Fig. 1: The four quadrants of the structural space

within the texts and subtle shifts in formal discourse through the classification of nodes according to their structural roles. Because the ECB’s discourse is highly formal and the resulting networks are complex, looking separately at 1) the top most frequent concepts, 2) the top most central concepts, or 3) concepts having the highest betweenness centrality will not be very informative. These top concepts are very similar across the four periods (as shown by the example in Table II) and constitute the core issues under discussion.

Concept frequency is arguably a more parsimonious metric than popularity (i.e., total degree centrality). However, since we are interested in the semantic structure, focusing of popularity over frequency is appropriate. While a naïve Pearson correlation between the two metrics is high ($r \approx 0.7$) for all of the periods, a closer inspection reveal significant variance in their relationship and that a log-linear association appears only for those concepts having higher than average frequency and total degree.

In order to explore both in-depth and orthogonally informative dimensions of the ECB discourse, we characterize the discourse using two distinct measures, building on the manner in which popular and connecting concepts play different roles in the structure and dynamics of semantic networks. Combining the popularity (i.e., total degree centrality) and connectivity (i.e., betweenness centrality) dimensions allows for the identification of four structural roles. This combination positions the concepts within this structural role space.

TABLE II: Concepts with the highest total degree centrality in each network

Pre-crisis		Crisis		Post-crisis		Recovery	
Concept	Deg.	Concept	Deg.	Concept	Deg.	Concept	Deg.
ecb	8.2	ecb	9.8	ecb	11.3	ecb	12.7
european	5.6	market	6.2	financial	6.9	bank	10.4
eu	5.2	central bank	6.2	euro area	6.3	european	9.4
system	4.9	eurosystem	5.9	market	6.2	financial	8.2
eurosystem	4.7	eu	5.5	bank	6.0	include	7.5
euro area	4.6	euro area	5.2	include	5.9	market	7.4
central bank	4.5	operate	4.8	system	5.8	monetary	7.3
include	4.5	national	4.5	eu	5.8	eu	7.2
market	4.3	include	4.2	central bank	5.3	central	6.9
operate	4.0	increase	4.1	economic	5.2	area	6.9

Note: concepts are color-coded to highlight their similarity; the degree centrality values are in units of 1000.

In order to connect the concepts with these structural roles, each concept in the network has been ranked based on its total degree centrality (C_D) and its betweenness centrality (C_B). For these rankings, we first consider the *set* of unique, unordered values X derived from some vector (or bag) of measures X_b . The order set of X is then:

$$X_{\text{ordered}} = \{x_1, \dots, x_n | x_i \in X; n = |X|; x_1 < \dots < x_n\}$$

where n is the number of unique measure values. We also define an index set J such that $x_j^{\text{ordered}} | j \in J$ is the j^{th} element of ordered set X_{ordered} . We now have a bijection $X_{\text{ordered}} \rightarrow X_{\text{rank}}$. For degree centrality, we replace X_b with the degree centrality measures C_D and obtain X_{ordered} which contains the unique, ordered degree centralities. For each node i , $C_D(i)$, we obtain the normalized degree centrality rank $C_D^{\text{rank}}(i)$:

$$C_D^{\text{rank}}(i) = 100 \cdot \frac{i}{n} | (x_j^{\text{ordered}} = C_D(i)).$$

The rankings are normalized and rescaled to the $[0,100]$ interval so that we can easily compare rankings across semantic networks. The rankings for betweenness centrality are obtained in a similar fashion (i.e., using C_B for X_b). In sum, we rank the total degree centrality and betweenness centrality scores for the concepts from each time period network into a normalized range between 0 and 100. Ranking was employed because a) the networks are of different sizes and densities hence, we want to be able compare across time periods and b) using the raw centrality scores produces less compelling and readable visualizations due high skewness of the distributions.²

Based on this ranking, we expose four structural roles, as four quadrants of the structural space (see Figure 1). For the sake of brevity, in the figure, we consider the ranks normalized in the $[0,1]$ interval. The Globally Central (GC) role includes concepts with high degree centrality and high betweenness centrality $[C_D^{\text{rank}} \times C_B^{\text{rank}}]$, where C_D^{rank} and C_B^{rank} are the normalized rankings of total degree centrality and betweenness centrality. These are very popular and highly connective concepts. The Locally Central (LC) role contains concepts with high degree centrality and low betweenness centrality $[C_D^{\text{rank}} \times (100 - C_B^{\text{rank}})]$. LC concepts are very popular concepts that do not have a strongly connective role. The Gatekeeper (G) role incorporates concepts with low degree centrality and high betweenness centrality $[(100 - C_D^{\text{rank}}) \times C_B^{\text{rank}}]$. These types of concepts are highly connective concepts that aren't very popular. Lastly, the Marginal (M) role includes concepts with low degree centrality and low betweenness centrality $[(100 - C_D^{\text{rank}}) \times (100 - C_B^{\text{rank}})]$. M concepts are neither popular, nor connective but they have the potential of becoming emergent concepts.

Figure 2 illustrates a layout example for the four structural roles described above using empirical centrality ranks from one of our sub-samples. The darker the red shade of the nodes, the higher embedded these nodes are in the region of the specific structural role.

²Alternatively, we could have employed normalized centrality scores. However, these exhibit the same skewness and still require transformation. Our approach is mathematically similar to using ranks of normalized scores.

Alternatively, the structural role scores could have been computed by simply adding the total degree centrality and betweenness centrality score components (i.e., the multipliers). However, this addition produces inaccurate role mappings. For example, Globally Central (GC) concepts become classified also as Gatekeepers (G) due to their high betweenness centrality irrespective of their high total degree centrality. Similarly, Marginal (M) concepts can appear as Gatekeepers due to their extremely low total degree centrality. We find multiplication of the role components to parsimoniously produce distinct role assignments.

While in semantic networks total degree centrality represents the popularity of a concept and betweenness centrality represents the links between two different thematic areas, the combination of these two measures has the potential to uncover more subtle structural properties of concepts, and thus a set of changes in discourse over time. A GC concept is a central key concept of a hot topic because not only is it highly connected to other concepts but it also serves as a bridge between different parts of the network. An LC concept is the central key concept of a local hot topic because it is highly connected but does not serve as a bridge in the network. A G concept is influential in the network because although it is not highly connected, it acts as a bridge in the network, linking different themes or topics. Such a concept can mark the emergence of merging themes. Lastly, an M concept, which is not well connected and does not serve as a linking concept, is common in discourse and may have the potential to become an emerging topic.

After identifying the four roles and the nodes that belong to each role, additional visual dimensions can be added by sizing, shaping, and/or colouring the nodes based on other measures. This approach provides other ways to explore each

Fig. 2: Example of structural roles

role individually, or the structural space as a whole in terms of distinct or subtle patterns. For example, other network metrics may be highly correlated within just a single quadrant.

For the purposes of this article, we will focus on a set of key concepts within the ECB discourse. Three of the key concepts have been selected because they reflect the main objectives of the ECB, as stated by the Treaty of the European Union [23], namely: ‘stability’, ‘growth’, and ‘interest rate’. The other key concepts have been included in the set as crisis-oriented terminology, namely: ‘crisis’, ‘debt’, ‘inflation’, ‘lend’, ‘loan’, ‘longer term refinancing operations (LTRO)’, ‘main refinancing operation (MRO)’, ‘refinancing’, and ‘risk’.³

III. RESULTS

1) *Structural roles*: We will begin by discussing each of the four periods by highlighting the observed variations in discourse as it develops across the different phases of the financial crisis. Below, we plot our semantic networks by using the structural roles and focusing on the selected key concepts. In Figures 3 to 6, we introduce an additional dimension (as earlier discussed in this paper) by colouring the nodes based on their raw frequencies of occurrence in the text data. The color spectrum ranges from dark **blue** (low frequency) to dark **red** (high frequency).

We also display edges among the focal concepts; that is, the subgraph induced by the node set comprising these concepts. The edges are weighted and represent the count of co-occurrences (within the two sentence moving window) between these focal nodes. Since we are interested in the actual volume of activity of these concepts and their co-occurrences, the edge weights are left unnormalized. We note that these edges do not represent the total activity of the focal concepts but just the activity among themselves (for presentation purposes). Finally, at the top of each graph, we report 1) the number of prominent nodes n (i.e., having a raw frequency greater than ten); 2) the number of distinct edges of the subgraph of key concepts $|E|$; and 3) the sum of the edge weights of that subgraph Σw .

Figure 3 shows that even in the *pre-crisis* period, before the end of 2007, crisis-oriented key concepts are present, some of them having relatively high total degree centrality (i.e., ‘lend’, ‘mro’, and ‘risk’) and being connected to the main objectives of the ECB. The globally central (GC) position of ‘risk’ as well as the “on-the-fence” position of ‘mro’ (which borders the locally central (LC) and the marginal (G) quadrants) could indicate that some of the ECB’s attention was focused on the emerging financial crisis before the end of 2007. We also observe that unlike ‘interest_rate’ and ‘stability’, which are highly ranked GC concepts, ‘growth’ (one of the ECB’s main objectives) is a highly ranked LC concept. This indicates that, during the *pre-crisis* period, ‘growth’ was a popular concept but not a very connective one.

³MRO’s serve to drive short-term interest rates, to manage the liquidity situation and to signal the monetary policy stance in the euro area, while LTRO’s provide additional, longer-term refinancing to the financial sector.

Fig. 3: Structural space of the *Pre-crisis* semantic network

In the *crisis* period (see Figure 4), most of the crisis-related key concepts are becoming more prominent. Concepts such as ‘inflation’, ‘loan’, ‘mro’, and ‘debt’ are ranked higher in total degree centrality and in betweenness centrality than in the previous period, suggesting they became more central and connective of different domains in the discourse of the ECB during the *crisis* period. At the same time, the betweenness and degree centralities of ‘interest rate’ and ‘stability’ noticeably decrease, suggesting once again that the main objectives of the ECB lose rhetorical ground against the full-blown financial crisis. The concept ‘risk’ is higher ranked in the GC category during the *crisis* becoming one of the ‘hottest’ topics of the ECB discourse. We also observe the emergence of ‘ltro’, a concept that was not present in the *pre-crisis* period. ‘ltro’ enters the discourse of the ECB as a very highly ranked G

Fig. 4: Structural space of the *Crisis* semantic network

Fig. 5: Structural space of the *Post-crisis* semantic network

concept, indicating that it connected disparate topics during the *crisis*. The link weights also increase in the *crisis* period, highlighting the increased co-activity among these concepts in the ECB press releases during this period. An interesting finding is that ‘crisis’ remains a marginal concept during the *crisis* period. The similar marginal position of ‘crisis’ during the *pre-crisis* as well as the *crisis* period raises questions regarding the discursive practices employed by the ECB. Had the ECB avoided the use of such concepts to prevent panic among stakeholders? Or had the ECB denied or ignored the existence of the crisis?

Figure 5 plots the structural space of the *Post-crisis* semantic network, showing all the main objectives of the ECB in a GC position. While in the *pre-crisis* only two of the main objectives were in GC positions and in the *crisis* period the centrality of these two concepts decreased, in the *post-crisis* period all the three main objectives return to being globally central, GC. At the same time we observe that the betweenness centrality of ‘crisis’ increases, while the betweenness of ‘debt’ and ‘inflation’ decreases significantly. These positional changes suggest that the shift in the ECB discourse could be at least partly explained by their struggle to deal with the aftermath of the crisis, while at the same time refocusing on their core objectives. The link weights show increased activity for ‘ltro’, ‘mro’, and ‘refinancing’, lending further evidence to the ECB’s resumed focus on the aftermath of the crisis.

Figure 6 plots the structural space of the *recovery* semantic network, revealing significant changes in the discourse of the ECB beyond the crisis. Compared to the previous *post-crisis* period, ‘interest_rate’ is now an LC concept. This concept, signifying one of the main objectives of the ECB, maintained a GC position in all the three previous periods analyzed. While ‘ltro’ suffers a drastic decrease in betweenness centrality (becoming an M concept), ‘mro’ and ‘refinancing’ become G concepts. We also note the positional change of ‘loan’,

moving from the LC quadrant to the GC quadrant. Based on all these structural changes, we argue that the *recovery* period exhibits a clear shift towards a discourse directed at dealing with the aftermath of the crisis. By assessing the width of the links, we see that the positional changes described above are also reflected in the co-occurrence levels. While ‘mro’ and ‘refinancing’ show increased activity, ‘ltro’ co-occurs less often with the other key concepts.

As for the graph-level, structural indicators, we observe that the count of nodes (i.e., n , the count of non-infrequent concepts) increases almost exactly linearly to the word counts of the collected documents for each period. These word counts are 28155, 30991, 33538, and 42892 (from *pre-crisis* through *recovery*). However, the activity in the focal concept subgraph does not follow suit. Specifically, the edge count increases initially and then stabilizes at ~ 53 , and the sum of edge weights peaks at *post-crisis* and then decreases. We surmise that the ECB discourse becomes expansive with the inclusion of additional topics (not identified in this paper). Hence, a naïve analysis using simple, relative frequencies of these key concepts would only diminish their importance. On the other hand, our structural role analysis reveals that some of the concepts associated with ECB’s objectives (here, ‘stability’ and ‘growth’) in fact remain prominent.

2) *MRQAP*: As the last part of our analysis, and in light of the findings above, we performed QAP correlations and multiple regressions (MRQAP); QAP is the acronym for the quadratic assignment procedure. These methods compare one or more networks using edges and their weights as data points while controlling for their dependencies [49]. This type of analysis is appropriate for our networks because we are using valued data and we can characterize each of the four periods as a function of its previous periods. The regression coefficients from an MRQAP are identical to those of least squares regression; however, their significance scores (i.e.,

Fig. 6: Structural space of the *Recovery* semantic network

p -values) are derived by comparing the estimates against their distributions obtained from applying the same regression model to a large sample of permutations ($m = 1000$) of the node structure thereby controlling for autocorrelation [49]. The same applies to the computation of a QAP Pearson correlation. The networks are first conformed by node count as the networks sizes need to be identical as required by standard regression.

The Pearson correlations reported in Table III are moderate to high despite the complexity of the four semantic networks. Their patterns show what we would nominally expect: proxi-

TABLE III: MRQAP correlations

	<i>Pre-crisis</i>	<i>Crisis</i>	<i>Post-crisis</i>	<i>Recovery</i>
<i>Pre-crisis</i>	–	0.758***	0.697***	0.593***
<i>Crisis</i>	–	–	0.758***	0.627***
<i>Post-crisis</i>	–	–	–	0.851***

*** indicates significance at $p < 0.001$

mal time periods bear the most resemblance while those farther apart differ the most. For example, the *pre-crisis* network's correlations diminish with more recent periods. Interestingly, the *post-crisis* and *recovery* periods exhibit more similarity to one another than any other adjacent pairs of periods, suggesting these periods are not as distinct as those other pairs and that *recovery* was likely already underway during the *post-crisis* period. Because the *pre-crisis* and *recovery* networks are least similar, we can argue that the *recovery* network is a transition phase in the ECB discourse towards a new state and not a resumption of the status quo of the *pre-crisis* period.

TABLE IV: MRQAP coefficients

Dependent	Intercept	<i>Pre-crisis</i>	<i>Crisis</i>	<i>Post-crisis</i>
<i>Recovery</i>	0.361***	0.037 [^]	−0.071***	0.968***

Adj- $R^2 = 0.725$; [^] = $p < 0.1$, *** = $p < 0.001$

We are also interested in how much of the first three periods constitute the *recovery* period. Conforming the four periods' networks by the intersection of their common nodes yields 382 nodes per network. In Table IV, we show the results of the MRQAP regression for predicting the *recovery* network from the earlier periods' networks. As suggested by correlations in Table III, the *post-crisis* period is the most predictive of *recovery*. Similarly, the *pre-crisis* period adds very little to the *recovery* period; however, the coefficient remains positive indicating a contribution to similarity. On the other hand, the negative (and significant) coefficient for *crisis*'s prediction on *recovery* indicates a slight reversal in the semantic structure from that period. That is, semantic associations of high prominence in *crisis* appear less prominently in *recovery*, controlling for the effects from the other two periods. In other words, the ECB seemed more inclined to focus less on the financial crisis and more on the subsequent recovery, an observation supported by the shifts in the structural roles (Figures 3 to 6). A similar regression analysis (not shown in this paper) using *post-crisis* as the dependent network shows that this reversal had not occurred yet in *post-crisis*, thereby qualifying the assertion we made earlier that the *post-crisis* and *recovery* periods were highly similar.

In light of the findings of the structural space analysis and the MRQAP, we can conclude that the *recovery* period is a different state in the discourse of the ECB. This new state in their discourse exhibits elements characteristic to the *post-crisis* and moves further away from the *crisis* period.

IV. CONCLUSION

The goals of the present article were three fold: 1) we sought to overcome two common challenges in text analysis, namely the size of the text corpora and its formal character; 2) we aimed to explore the benefits of the structural space dimensions; and 3) we wanted to investigate how the discursive practices of the ECB have been affected by the recent financial crisis.

The structural space method employed by this study revealed substantial and imperative shifts in the ECB discourse, demonstrating that it could be a valuable instrument for change detection in formal discourse. As shown, looking at the obvious most central concepts in formal discourse does not always reveal the underlying and subtler shifts across the periods investigated. Formal discourse such as that of the ECB contains repetitive top key concepts, indicative of the obvious and perhaps uninformative central topics of an organization. The structural space analysis proved more explanatory regarding the shifts and changes in formal discourse, by combining structural measures and looking beyond the core of the network structure. At the same time, structural roles of key concepts may be good predictors of emerging topics and the dynamics of discursive change.

In recent years, an increasing number of researchers have been focusing on the importance of central banks' communication [27], [31], [50]–[52]. However, not much research has addressed the communications of the ECB from a discourse analysis perspective. As the central bank for the euro, the ECB does not have many instruments to directly influence the markets. Therefore, their communications become a key policy instrument. In other words, the ECB used communications especially to signal interest rate increases, and consequently directly influence private sector expectations [28], [29]. The importance of the communications issued by the ECB is thus understandable and is a valuable source of information for financial market participants. This being said, in times of crisis, the ECB's role in guiding financial market expectations through communication is particularly important due to higher market uncertainty. While previous research showed that ECB communications increase the predictability of interest rate decisions [30], our focus was directed at uncovering the shifts and adaptations of the ECB discourse in a time of crisis.

The structural space dimension of the selected key concepts exposed significant changes in the ECB discourse. Below, we summarize our main findings for each of the examined periods.

Pre-crisis: Our analysis showed that crisis-oriented key concepts were already present, suggesting that even before the end of 2007 the ECB's discourse shifted towards crisis terminology, and their focus may have been on the impending financial crisis.

Crisis: The crisis-related key concepts detected by our method in the *pre-crisis* period became more prominent in the *crisis* period. We showed that the key concepts associated with the main objectives of the ECB lost ground in front of the effects of the full-blown financial crisis. Also, we noted the emergence of longer-term-refinancing ('ltro') as a highly ranked connective (G) concept and the increase in betweenness for main refinancing operations ('mro'). These changes denote the focus of the ECB on refinancing operations during the *crisis* period. Interestingly, the 'crisis' concept has not been a highly connective, nor a popular concept during this period. The marginal position of this concept in both the *pre-crisis* period and the *crisis* period could denote an intentional attempt of the ECB to minimize panic reactions among the stakeholders, or it could be explained by a narrow focus of the ECB's discourse towards the overwhelming market defaults and not towards the crisis as a whole.

Post-crisis: This period revealed the ECB's discourse in a distinct state, where all the main objectives of the ECB are in a prominent position, while 'crisis' itself became a more connective concept. The changes observed in this period point towards a focus of the ECB's discourse towards dealing with the aftermath of the crisis.

Recovery: While in the *crisis* and the *post-crisis* periods 'ltro' is the highest ranked G concept, during the *recovery* period it suffers a drastic decrease in betweenness centrality ranking, appearing now as a marginal concept (M). At the same time, we show 'loan' becoming more popular, and 'mro' and 'refinancing' becoming more connective. These structural changes in the semantic structure show the shift towards dealing with the aftermath of the crisis more clearly than in the previous period.

Finally, our findings revealed that by the end of 2013, the discourse of the ECB had in no way returned to the *pre-crisis* levels, but perhaps advanced to a 'new state' altogether. This 'new state' could be explained by the role of the ECB in dealing with the aftermath of the financial crisis. Also, during the *recovery* period the ECB seemed to focus less on the financial crisis and more on the subsequent recovery process. This particular finding was supported by the structural space analysis as well as the MRQAP coefficients.

Although the method we have employed in this study revealed important findings, one of its limitations is the fact that we have only used it with a single data set (divided in four periods). Nevertheless, based on the relevance of the results uncovered by this method, we advocate for further development and testing of the structural space as a method for analysis of semantic networks.

Future research employing this method should also explore the inclusion of other structural measures, such as closeness centrality, clique counts or clustering coefficient. During our preliminary analysis, we tested the inclusion of eigenvector centrality in the structural space, which proved to be highly correlative to total degree centrality, and thus did not add anything to the informative value of the structural roles.

While our method of classifying nodes into one of the four

structural roles was used to highlight only a handful of key concepts, the classification may easily be broadened to identify lists of top concepts (e.g., top ten) within each of the roles. This enumeration of the roles offers a more complete depiction of the roles and their evolution.

Our naïve treatment of weighted degree centrality, while typical in network research, raises some concerns. Specifically, weights and the number of distinct ties ought to be considered separately as the same total degree centrality score of a node can arise from starkly different ego-centric structures. While the exploration of this issue is beyond the scope of this paper, we hope (and expect) that future research will improve the use of weighted degree centrality in semantic and social network analysis.

As for the complex structure of our semantic networks, some diagnostic tests on our networks reveal that they are only mildly small-world and not at all scale-free, contrary to the findings of other work. Still, further investigation (outside the scope of this paper) would be required to determine if these inconsistencies are due to the type of semantic network or the exact nature of semantic network extraction or simply that semantic networks can vary widely in their topologies. As for metric comparisons with other research, our within-network correlations for our two centrality measures were modestly high, echoing other findings, e.g., [53], [54] and also highlighting the usefulness of the structural role approach in identifying outliers in the G and LC roles.

Our use of centrality ranks as opposed to actual centrality scores warrants additional, future inquiry. We suspect that in order to compare them more precisely across networks of varying sizes and densities, tighter controls must be exerted. We envision highly robust comparative indices that account for both the relative or ranked centrality score as well as the absolute score.

Also, our study aggregated the data in four periods of two years each. Arguably, smaller data time slices could potentially reveal subtler aspects in the dynamics of discourse and fluctuations in terminology.

To conclude, we can argue that our approach proved beneficial for the analysis of large corpora or formal organizational discourse. We anticipate our noteworthy results to open new avenues for semantic network research dealing with formal discourses and beyond the context of the financial crisis.

ACKNOWLEDGMENT

Research funded by the NWO-Aspasia 015-007-047 grant.

REFERENCES

- [1] J. T. Irvine, "Formality and Informality in Communicative Events," *American Anthropologist*, vol. 81, no. 4, pp. 773–790, 1979.
- [2] M. Alvesson and D. Kärreman, "Varieties of discourse: On the study of organizations through discourse analysis," *Human Relations*, vol. 53, no. 9, pp. 1125–1149, 2000.
- [3] D. Grant and C. Hardy, "Introduction: Struggles with Organizational Discourse," *Organization Studies*, vol. 25, no. 1, pp. 5–13, 2004.
- [4] J. Borge-Holthoefer and A. Arenas, "Semantic Networks: Structure and Dynamics," *Entropy*, vol. 12, no. 5, pp. 1264–1302, 2010.

- [5] M. E. Bales and S. B. Johnson, "Graph theoretic modeling of large-scale semantic networks." *Journal of biomedical informatics*, vol. 39, no. 4, pp. 451–64, 2006.
- [6] M. Steyvers and J. B. Tenenbaum, "The large-scale structure of semantic networks: statistical analyses and a model of semantic growth." *Cognitive science*, vol. 29, no. 1, pp. 41–78, 2005.
- [7] E. Postma, F. Wiersman, and H. van den Herik, "Small-world semantic networks," in *Proceedings of the Twelfth Belgium-Netherlands Artificial Intelligence Conference*, 2000, pp. 191–198.
- [8] R. Ferrer I Cancho and R. V. Solé, "The small world of human language." *Proc. of Biological Sciences*, vol. 268, no. 1482, pp. 2261–5, 2001.
- [9] R. V. Solé, B. Corominas-Murtra, S. Valverde, and L. Steels, "Language networks: Their structure, function, and evolution," *Complexity*, vol. 15, no. 6, pp. 20–26, 2010.
- [10] K. M. Carley and D. S. Kaufer, "Semantic connectivity: An approach for analyzing symbols in semantic networks," *Communication Theory*, vol. 3, no. 3, pp. 183–213, 1993.
- [11] J. A. Danowski, "Analyzing Change over Time in Organizations' Publics with a Semantic Network Include List: An Illustration with Facebook," in *ASONAM*. IEEE, 2012, pp. 954–959.
- [12] J. Diesner, "From Texts to Networks: Detecting and Managing the Impact of Methodological Choices for Extracting Network Data from Text Data," *Künstliche Intelligenz*, vol. 27, no. 1, pp. 75–78, 2013.
- [13] R. Burt, *Structural holes: The social structure of competition*. Cambridge, MA: Harvard University Press, 2009.
- [14] R. S. Burt, "Cohesion Versus Structural Equivalence as a Basis for Network Subgroups," *Sociological Methods & Research*, vol. 7, no. 2, pp. 189–212, 1978.
- [15] R. Burt, "Detecting Role Equivalence," *Social Networks*, vol. 12, no. 1, pp. 83–9, 1990.
- [16] J. P. Boyd and M. G. Everett, "Relations, residuals, regular interiors, and relative regular equivalence," *Social Networks*, vol. 21, no. 2, pp. 147–165, 1999.
- [17] S. P. Borgatti and M. G. Everett, "Notions of position in social network analysis," *Sociological Methodology*, vol. 22, no. 1, pp. 1–35, 1992.
- [18] H. C. White, S. A. Boorman, and R. L. Breiger, "Social structure from multiple networks. I. Blockmodels of roles and positions," *American journal of sociology*, pp. 730–780, 1976.
- [19] C. J. Anderson, S. Wasserman, and K. Faust, "Building stochastic blockmodels," *Social Networks*, vol. 14, no. 1, pp. 137–161, 1992.
- [20] S. A. Boorman and H. C. White, "Social structure from multiple networks. II. Role structures," *American Journal of Sociology*, pp. 1384–1446, 1976.
- [21] S. Huang, T. Lv, X. Zhang, Y. Yang, W. Zheng, and C. Wen, "Identifying node role in social network based on multiple indicators." *PloS one*, vol. 9, no. 8, p. e103733, 2014.
- [22] D. Hansen, B. Shneiderman, and M. A. Smith, *Analyzing Social Media Networks with NodeXL: Insights from a Connected World*. Morgan Kaufmann, 2010.
- [23] European Union, "Treaty of Amsterdam amending the Treaty on European Union, the Treaties establishing the European Communities and related acts," *Official Journal of the European Communities*, p. 144, 1997.
- [24] B. Bernanke, V. Reinhart, and B. Sack, "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment," *Brookings papers on economic Activity*, vol. 2, pp. 1–100, 2004.
- [25] R. S. Gürkaynak, B. Sack, and E. Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, vol. 1, no. 1, pp. 55–93, 2005.
- [26] R. S. Gürkaynak, B. Sack, and E. Swanson, "The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models," *The American Economic Review*, vol. 95, no. 1, pp. 425–436, 2005.
- [27] D. L. Kohn and B. P. Sack, "Central Bank Talk: Does It Matter and Why?" 2003.
- [28] D. Jansen and J. Haan, "Is a word to the wise indeed enough? ECB statements and the predictability of interest rate decisions," *De Nederlandsche Bank, Amsterdam, Tech. Rep. 75*, 2005.
- [29] D. J. Jansen and J. Haan, "Were verbal efforts to support the euro effective ? A high-frequency analysis of ECB statements," *European Journal of Political Economy*, vol. 23, no. 1, pp. 245–259, 2005.
- [30] J. Haan, "The effect of ECB communication on interest rates: An assessment," *The Review of International Organizations*, vol. 3, no. 4, pp. 375–398, 2008.
- [31] C. Rosa and G. Verga, "Is ECB communication effective?" 2005.
- [32] K. M. Carley, D. Columbus, and P. Landwehr, "Automap User s Guide 2013," CMU-ISR-CASOS, Pittsburgh, PA, Tech. Rep., 2013.
- [33] K. M. Carley and M. Palmquist, "Extracting , Representing , and Analyzing Mental Models *," *Social Forces*, vol. 70, no. 3, pp. 601–636, 1992.
- [34] J. Danowski, "Network analysis of message content," in *Progress in Communication Science*, W. Richards and G. Barnett, Eds. Norwood, NJ: Ablex, 1993, ch. XII, pp. 197–222.
- [35] R. Popping, "Knowledge Graphs and Network Text Analysis," *Social Science Information*, vol. 42, no. 1, pp. 91–106, 2003.
- [36] J. Diesner, "Uncovering and Managing the Impact of Methodological Choices for the Computational Construction of Socio-Technical Networks from Texts," Ph.D. dissertation, Carnegie Mellon University, 2012.
- [37] S. Wasserman and K. Faust, *Social network analysis: Methods and applications*. Cambridge: Cambridge University Press, 1994, vol. 24, no. 1.
- [38] L. Freeman, "Centrality in social networks conceptual clarification," *Social Networks*, vol. 1, no. 3, pp. 215–239, 1979.
- [39] P. Bonacich, "Power and Centrality : A Family of Measures," *American Journal of Sociology*, vol. 92, no. 5, pp. 1170–1182, 1987.
- [40] L. Freeman, "A Set of Measures of Centrality Based on Betweenness," *Sociometry*, vol. 40, no. 1, pp. 35–41, 1977.
- [41] G. Henderson, D. Iacobucci, and B. Calder, "Brand diagnostics: mapping branding effects using consumer associative networks," *European Journal of Operational Research*, vol. 111, no. 1998, pp. 306–327, 1998.
- [42] D. Iacobucci, G. Henderson, A. Marcati, and J. Chang, "Network analyses of brand switching behavior," *International Journal of Research in Marketing*, vol. 13, pp. 415–429, 1996.
- [43] V. Hill and K. M. Carley, "An approach to identifying consensus in a subfield: The case of organizational culture," *Poetics*, vol. 27, no. 1, pp. 1–30, 1999.
- [44] C. J. Hooper, N. Marie, and E. Kalampokis, "Dissecting the Butterfly: Representation of Disciplines Publishing at the Web Science Conference Series," in *WebSci*, 2012, pp. 197–200.
- [45] U. Brandes and S. R. Corman, "Visual unrolling of network evolution and the analysis of dynamic discourse," *Information Visualization*, vol. 2, no. 1, pp. 40–50, 2003.
- [46] C. Grebitus and M. Bruhn, "Analyzing semantic networks of pork quality by means of concept mapping," *Food Quality and Preference*, vol. 19, no. 1, pp. 86–96, 2008.
- [47] R. C. Hulst, "Introduction to Social Network Analysis (SNA) as an investigative tool," *Trends in Organized Crime*, vol. 12, no. 2, pp. 101–121, 2008.
- [48] P. Gloor and J. Krauss, "Web science 2.0: Identifying trends through semantic social network analysis," in *International Conference on Computational Science and Engineering*. Institute of Electrical and Electronics Engineers, 2009, pp. 215–222.
- [49] D. Krackhardt, "QAP Partialling as a Test of Spuriousness," *Social Networks*, vol. 9, pp. 171–186, 1987.
- [50] M. Ehrmann and M. Fratzscher, "Communication by Central Bank Committee Members : Different Strategies , Same Communication by Central Bank Committee Members : Different Strategies , Same Effectiveness ?" *Journal of Money, Credit and Banking*, vol. 39, no. 2, pp. 509–541, 2007.
- [51] B. M. Friedman, "The use and meaning of words in central banking: Inflation targeting, credibility, and transparency," 2002.
- [52] J. Sturm and J. D. Haan, "Does central bank communication really lead to better forecasts of policy decisions? New evidence based on a Taylor rule model for the ECB," 2011.
- [53] T. W. Valente, K. Coronges, C. Lakon, and E. Costenbader, "How Correlated Are Network Centrality Measure," *Connections*, vol. 28, no. 1, pp. 16–26, 2008.
- [54] J.-S. Lee and J. Pfeffer, "Estimating centrality statistics for complete and sampled networks: Some approaches and complications," in *HICSS (forthcoming)*, 2015.