Community detection in the European Parliament: A network approach

Candidate number: 1003457

September 4, 2016

Abstract

In this dissertation, we study networks of legislators, based on their voting behaviour. We describe different network-construction methods and we review the state of the art of how scholars transform voting data into networks describing the voting similarity of legislators. We describe how one can detect dense sets of nodes called "communities" and we discuss previous analyses of voting behaviour in the European Parliament.

In addition to our review material, we include original research in a novel analysis of polarization. We use modularity optimization to measure polarization and detect communities in the 7th Assembly (2009–2014) of the European Parliament (EP). Our results suggest that polarization levels in the EP are different for different policy areas. We find relatively "robust" partitions that split the EP in two communities. It depends on the policy area whether we can assign a label to each community. For some policy areas ("environment", "gender equality"), we find "left wing" and "right wing" communities, but for other policy areas ("agriculture", "industry"), there appear instead to be an "extreme" community and a "centrist" community. Hence, our results suggest that one should take into account the differences between policy areas when one wants to study community structure in the European Parliament.

Acknowledgements

Contents

Chapter 1

Introduction

"Politics is a network".
– Metz and Jäckle (2016) –

Background

In 1927, Rice researched whether one can detect "blocs" of legislators within small legislative bodies based on voting data [102]. He defined a bloc as a set of (at least two) legislators such that there is a ratio of voting agreement higher than a particular (arbitrary) threshold for every pair of legislators [102]. Beyle [13] continued Rice's work [102, 103] and described a method (that is now known as the "Rice–Beyle method") to detect voting blocs and measure their voting agreement [66]. Rice [103] assumed the existence of voting blocs and then measured their voting agreement, whereas the Rice–Beyle method simultaneously identifies blocs and measures their agreement [66]. In the past decades, many scholars have analyzed voting behaviour from different points of view [6, 13, 37, 43–46, 66, 102, 103, 105]. Recently, scholars have used several tools from network science to study topics in political science [63, 70, 123]. Particularly, some have used networks to analyze voting patterns and to detect blocs in assemblies [20, 68]. In this paper, we use network analysis to simultaneously identify blocs of legislators in the European Parliament and measure their voting agreement.

Aim of this dissertation

We study networks of legislators, based on their voting behaviour. We describe different network-construction methods and we review the state of the art of how scholars transform voting data into networks describing the voting similarity of legislators. We describe how one can detect dense sets of nodes called "communities" in those networks. We also discuss previous analyses of voting behaviour in the European Parliament. In addition to our review material, we include original research in a novel analysis of polarization. We analyze polarization and communities in the 7th Assembly of the European Parliament (2009–2014).

Research questions

We use tools from network science to measure polarization in the European Parliament (EP). We investigate whether the EP has different levels of polarization for different policy areas (such as "gender equality" or "agriculture"). We seek to examine the extent to which this polarization corresponds to a division based on political parties or countries of origin. We also hope to examine the kinds of communities that are present in the European Parliament.

We investigate whether these communities are "robust" and whether we can assign them a label (e.g. in terms of "left wing" or "right wing").

Outline

In Chapter 2, we introduce some notions from network science that we need throughout the dissertation. In Chaper 3, we give an introduction to roll-call voting and present different types of roll-call voting networks. In Chapter 4, we indicate tersely which one we use and we list several options for the construction of such a network. In Chapter 5, we present the notion of network "communities" and discuss ways that they can be detected algorithmically. In Chapter 6, we elaborate on one method to detect communities; the optimization of modularity. In Chapter 7, introduce the European Parliament, and we discuss our own original research in Chapter 8. We conclude in Chapter 9.

Chapter 2

Preliminaries

In this chapter, we introduce some network-scientific notions that we will use throughout the dissertation. This chapter draws a lot of material from Chapters 1 and 6 in Newman's book (2010) about networks [86].

2.1   What is a network?

A *network* or *graph* $G = (V, E, w)$ is a collection of points that can pairwise be joined by lines. These points are called *nodes* and these lines are called *edges*. The set of notes is denoted by V and the set of edges by E. We work with *undirected* networks, which means that edges do not have a direction. Hence, an edge from node i to node j is the same as an edge from node j to node i. We use $n = |V|$ to denote the number of nodes in a network and $m = |E|$ to denote the number of edges. We use *weighted* networks, in which edges have a weight w_{ij} (e.g. a real number) that represents the strength of the associated tie between node i and j. A *signed* networks only has weights with value of $+1$ or -1.

   If there is an edge e between node i and node j, we say that i and j are *adjacent* to each other. We also say that they are *neighbours*. An edge e and node i are called *incident* in this case (e and j are incident here as well). We study networks that have at most one edge between any pair of nodes. An edge from a node to itself is called a *self-loop*. We assume that networks are *connected*, which means that there is a path between every pair of nodes (so it is possible to reach all nodes if one would travel along edges in the network starting from any node). A *subgraph* is a network $H = (U, F, w)$, such that $U \subseteq V$, $F \subseteq E$, and nodes i, j are adjacent in H only if they are adjacent in G. A network in which there is an edge between every pair of nodes is called *complete*.

2.2   Adjacency matrices

An *adjacency matrix* $\mathbf{A} \in \mathbb{R}^{n \times n}$ contains all information about a network. Its components A_{ij} are given by

$$A_{ij} = \begin{cases} w_{ij}, & \text{if } i \text{ and } j \text{ are adjacent}, \\ 0, & \text{otherwise}. \end{cases} \tag{2.1}$$

   An adjacency matrix corresponding to an undirected network is symmetric, because the existence of an edge from i to j implies that there is an edge from j to i. In a network without self-loops, all diagonal entries A_{ii} are equal to 0.

2.3 Bipartite networks

One special network type is a *bipartite* network. The nodes of a bipartite network are partitioned into two disjoint sets. Every edge runs between those two sets (i.e. every edge is incident to a node in both sets). It is possible to "project" a bipartite network onto a unipartite network, but this leads to a loss of information. We discuss projections of bipartite networks in Section 3.3, when we examine bipartite roll-call voting networks.

2.4 Centrality measures

Centrality measures quantify the importance of nodes or edges in a network. There are many different centrality measures, such as the degree centrality, betweenness centrality, and eigenvector centrality (see [86]). For our computations, we only consider *degree* centrality.

The *degree* d_i of node i is the number of edges incident to it. In a weighted network, the *strength* k_i of a node i is given by the sum of weights of all incident edges:

$$k_i = \sum_{j=1}^{n} A_{ij}. \tag{2.2}$$

Because there are m edges in a network and every edge in an undirected network has two "ends", there are $2m$ ends of edges. The number of ends must be equal to the sum of all degrees, so

$$m = \frac{1}{2} \sum_{i=1}^{n} d_i. \tag{2.3}$$

Similarly, for weighted networks the *total weight* W of the network, which is the sum of all edge weights, is twice the sum of all node strengths:

$$W = \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} A_{ij}. \tag{2.4}$$

2.5 Exponential random graph models (ERGMs)

One uses *exponential random graph models (ERGMs)* to examine the process of network formation [20]. For instance, one can investigate how the degree distribution of a network affects the network-formation process. To find the relevant factors for the process, one can consider an ensemble of possible networks, and measure the similarity between these networks and the original network [20]. We do not consider ERGMs in our research, but we refer to them again in Section 7.7.3 when we discuss studies of the European Parliament.

2.6 Network analysis

One uses network analysis to discover patterns of connections, which can be modelled as a network in which the nodes represent components of the system and edges represent their connections. Hence, a network captures basic connection patterns [86]. An extensive set

of tools for analyzing, modelling, and understanding networks has been developed over the years [86]. One useful tool is network visualization, which allows people to uncover patterns that would be difficult to discover otherwise [57].

There are different scales at which one can analyse networks [98, 106, 127]. For example, one can focus on properties of a single node or edge (e.g. how many edges are incident to a particular node), which are called *micro-scale* properties. Second, if one focuses on an entire network (e.g. how many edges does it take to "travel" between two random nodes), then one is looking at the *macro-scale*. Finally, one can consider *meso-scale* properties to discover features that might not be apparent at the micro-scale or macro-scale levels [106]. For instance, one can examine the way a network breaks down into dense sets of nodes known as "communities" (see Chapter 5).

2.7   Use of networks in political science

Political scientists have also used network analysis [63, 70, 123]. By using network analysis, one can identify which political actors are important and how these actors are connected to each other, even though their relations are often invisible [57] and their performance is often hidden by complicated voting procedures [24]. Network analysis is also useful because it takes into account interdependence of political actors. In politics, flows of influence among actors are important, so studying individual characteristics is insufficient [123].

Another advantage of network analysis is that one does not have to assume subjective judgements in order to uncover structures in a political network nor does it require the researcher to have a political opinion [95]. Finally, network analysis is useful as a means to verify the results from other techniques in political science [4].

2.8   Examples of political networks

Scholars have studied many different political networks in the past decade. It is beyond the scope of this dissertation to review this large body of work in detail, but several reviews exist [63, 70, 109, 117, 123]. We briefly mention three examples to illustrate the power of network analysis in political science. Keller (2015) examined promotion networks in the Communist Party of China (CPC), where she used network centrality measures to identify and predict the inner circle of the CPC five or ten years later [54]. Second, in a series of papers, several authors studied committee-assignment networks in the United States and showed that communities of committees can be identified without *a priori* knowledge about the politicians [95–97]. They used various network types, such as bipartite networks (in which legislators are connected to committees) or weighted unipartite networks (in which edges indicate the number of committees on which two legislators both sit). Finally, a type of political network that has been studied quite a lot in the past decade is a legislation cosponsorship network, in which legislators are connected if they (co)sign the same bill [1, 2, 9, 19, 22, 23, 34, 35, 40, 57, 64, 74, 75, 99, 111, 128].

In Chapters 7 and 8, we analyze voting patterns in the European Parliament and try to detect communities using tools from network science.

Chapter 3

Roll-call voting

3.1 What is roll-call voting?

One of the core tasks of members of parliament (MPs) (which we will also call "legislators") is to vote on bills. A *roll-call vote (RCV)* is a vote in which the names of the legislators (along with their votes) are recorded [20]. Hence, anonymous voting procedures or procedures in which only the total result is recorded are not included in this definition. In some countries, roll-call voting is exercised only upon request from at least a certain number of legislators. There are different reasons why legislators would want roll-call voting to take place. For instance, it might be used to emphasize a parliament's will to pass/reject a bill or to place other legislators under public scrutiny for their voting behaviour [6].

In some countries, such as the United States, all votes are recorded, so a lot of roll-call voting data exist [93]. In many other countries, roll-call data is often unavailable for periods prior to 1990 [3]. Because roll-call data is recently becoming more available [2], the investigation of roll-call voting records has become more popular [68, 73, 93, 94, 124].

Strengths of studying roll-call votes

One strength of studying roll-call voting is that we have information about all legislators (who are present), as each of them is supposed to vote "yes", "no", or "abstain" for every bill. Hence, we can distinguish between legislators who are against a certain bill and legislators who are not in favour (which is not possible for cosponsorship). However, if legislators are absent (which is not unusual in some parliaments [73]), roll-call voting no longer gives information about the opinion of all legislators. Another strength of roll-call voting data is that it can be used to study voting cohesion within a political party (e.g. the extent to which legislators vote in line with their party leader). Furthermore, because newspapers sometimes report how legislators voted, roll-call voting data can also be used to analyse pressure on legislators to vote consistently with the preference of their voters [3].

Weaknesses of studying roll-call votes

A weakness of studying roll-call voting is the availability of information. In some countries, such as the United States, legislative votes are recorded systematically. By contrast, in others (e.g. developing countries), votes are only recorded occasionally and under special circumstances [3]. If this roll-call voting data is unavailable, one could use other types of data (such as cosponsorship data) to get insights about party politics [1].

Another weakness is that voting behaviour is often more subject than cosponsorship to party discipline and agenda setting, as party leaders exert more influence over voting than they do over cosponsorship [19]. Hence, it is hard to determine whether a legislator voted

in a particular way because of his/her own preference or because of other reasons (such as pressure from the party leader).

3.2   Spatial modelling of roll-call votes

There are several methods for analysing roll-call data [67]. We use network analysis, but for comparison we briefly discuss a non-network approach called *spatial modelling*, which has been used often for analysing roll-call voting data [3, 8, 69, 93, 94, 120]. Spatial modelling is a powerful tool to characterize political ideology of legislators by finding structural patterns in voting records. The idea behind spatial models is that vote choices can effectively be represented in a (low-dimensional) space in which a small number of dimensions can generate most individual issue dimensions [67]. Every legislator has his/her own ideal point for a specific bill (his/her most favoured policy outcome). This ideal point is assumed to be a symmetric single-peaked utility function, and a legislator will vote for the option that is closest to his/her ideal point, as it gives him/her the highest utility [3]. The ideal-point estimates are inferences based on statistical estimates, such that they fit as closely as possible with the observed roll-call votes [67].

There has been an extensive literature on legislative behaviour using spatial models [8, 49, 69, 93, 94, 97, 112]. For instance, Alemán et al. (2009) compared roll-call voting to cosponsorship and concluded that the ideal-point estimates from cosponsorship and roll-call voting are correlated with each other for voting data in the United States and Argentina [3].

Waugh et al. argued that spatial models are more appropriate for individual roll-call decisions rather than as measures of system-wide polarization (e.g. in terms of a distance between political parties) [124]. Even though the assumptions in spatial modelling (such as single-peakedness, transitivity, and completeness of preferences) seem to work for fitting individual roll-call votes, these assumptions are too strict when one wants to examine group dynamics on formal party divisions [124]. Because we are interested in examining political group dynamics and measuring polarization in parliament, we thus follow Waugh et al. and use network analysis to analyse the European Parliament.

3.3   Roll-call voting networks

Why should one use a networks approach?

Even though network analysis has become more popular in political science [63, 109, 117], there are still very few studies that have used a network approach for analysing legislators' voting patterns [74]. However, studying roll-call voting from a network perspective has several clear advantages over other approaches.

First, no assumptions about rationality or spatial structures have to be made. Instead of treating polarization as an ideological phenomenon with spatial assumptions, Waugh et al. argued that it is better to focus on the observed voting behaviour and the way that this behaviour divides legislators into distinct groups, by using tools in network science like community detection (see Chapter 5) [124]. Moreover, they argued that the quality function (to assess the goodness of a community assignment) called "modularity" (see Chapter 6) offers a clearer measure of polarization than the existing spatial modelling measures [124].

Second, Metz and Jäckle (2016) argued that network analysis may be able to measure intra-party differences and to "partially open the black box of the political party", even when party discipline is strong [74]. This is useful because other methods do not always work well in that case. (For instance, if one uses spatial models of roll-call voting when party discipline is strong, it may be harder to interpret the dimensions [67]).

Which network representation should one use?

Macon et al. (2012) discussed different approaches to construct a roll-call voting network [68]. Their primary goal was to investigate different ways to turn roll-call voting data into a network and then to use community-detection tools (which are methods that try to find a partition of the network in which there are dense sets of nodes that have few edges between them; see Chapter 5). We use one of their approaches for our computations (a single unipartite network), but we briefly discuss the other options here as well.

Unipartite networks

Several papers have used a single unipartite network to examine roll-call voting data [24, 124]. In such an approach, each node often represents a legislator and every edge between a pair of legislators usually indicates some measure of the extent to which those two legislators voted the same way. The network is unipartite because there is only one type of node (representing legislators), so there can be edges between any pair of nodes. The adjacency matrix of this network is given by a so-called *similarity matrix*, which shows the level of voting agreement between each pair of legislators (see Chapter 4).

Instead of having one unipartite network, it is also possible to construct two separate unipartite networks: a separate agreement and disagreement network [68]. In this approach, one can separately consider voting agreement and disagreement between legislators. Consequently, one can treat forms of disagreement differently (yes–no pairs versus abstain–no pairs), as the yes–no pairs are counted separately. By using separate parameters for the agreement and disagreement network, one can assign different importance to each type of pair [68].

Bipartite networks

One can also use a signed bipartite network to represent how legislators are connected to their votes [68]. With such an approach, one set of nodes represents the legislators, and the other set of nodes represents the bills on which there were votes. Every time legislator i votes in favour of bill j, there is a positive edge e_{ij}. If legislator i votes against a bill, then there is a negative edge. One thereby uses a signed bipartite network to capture the roll-call voting data, in which abstentions are not included.

Although it is possible to transform a bipartite network into one or two unipartite networks (with or without self-loops), the advantage of signed bipartite networks is that no information about bills is thrown away. Legislators in unipartite networks are connected only to each other, but when using a bipartite network, we retain the information about specific bills on which two legislators agree or disagree. Bills may be very different (e.g. they might be related to different policy areas), so one should try to include this source of information in a different way (e.g. by constructing a distinct network for each policy area; see Section 8.2).

Comparison of network representations

Because a bipartite approach uses more of the available information than the unipartite approaches, one could argue that it is better to work with the bipartite network representation [68]. However, Macon et al. (2012) concluded that the unipartite network representations can uncover a more diverse set of robust groups, which complements the information that they uncovered by a bipartite approach. Moreover, they found the predominant grouping of countries in the United Nations General Assembly for all of these different network representations [68]. In light of their experiences with their example, they suggested it is best to consider multiple representations when one wants to obtain a thorough picture of roll-call voting. For our computations, we use a single unipartite network approach because the data we obtained was originally made available as a single unipartite network (see Section 8.1).

Which edge weights should one use?

Unweighted networks

To create an unweighted roll-call voting network, one has to encode all similarity scores as 1 values or 0 values. Scholars have used several approaches to convert different types of data (such as voting data) into an unweighted network [2, 4, 40, 127]. First, one can use a uniform bound δ as a threshold to create an unweighted adjacency matrix [4]. That is, given a weighted adjacency matrix $\mathbf{A}$ and a threshold δ, the unweighted adjacency matrix $\mathbf{A}^\delta$ is

$$A_{ij}^\delta = \begin{cases} 1, & \text{if} \quad A_{ij} \geq \delta, \\ 0, & \text{otherwise}. \end{cases}$$

The choice of threshold δ is subjective. Some authors picked an arbitrary value, such as Amelio and Pizzuti (2015), who put $\delta = 0.6$ so that two legislators are connected in the unweighted network if they vote the same way at least 60% of the time [4]. Another option is to set $\delta = \mu$, where μ is the mean weight of all edges [2]. Gross et al. (2012) chose a threshold less arbitrarily by considering several values of δ and comparing their results. They thereby observed which relationships seem to persist as weaker connections are removed [40]. Zanin et al. (2016) suggested alternative approaches to using a uniform bound as a threshold [127], such as maintaining some fraction of the edges with the highest weights or using a statistical test to select edges that are significant. They also suggested using tools from data mining to create unweighted networks. That is, after applying multiple thresholds, they analyzed the resulting networks and reported "the most relevant features" by using recursive feature elimination methods based on support vector machines and graph kernels [127]. For additional types of thresholds, see [125].

Weighted networks

In a weighted network, one can distinguish the strength of ties based on values (i.e. weights) associated to edges. The disadvantage of using weighted networks is that some measures from network analysis are harder to apply. They might be unclear, as the interpretation of weights may be ambiguous (e.g. reciprocity of weighted networks [113]) or they may not be defined at all for weighted networks. Moreover, roll-call voting networks are expected to be almost complete networks, as there are barely any 0 entries in the adjacency matrix. This has inspired

some authors to use unweighted networks [4]. Other authors decided to display unweighted versions of roll-call voting data (e.g. to illustrate community structure more clearly in their figures), but to keep all information from the weighted network for their calculations and analysis [74].

Despite these issues, we use weighted networks for the following reasons. Alemán and Calvo (2013) argued that different levels of agreement between pairs of legislators are meaningful, so it is important to take advantage of them, rather than throwing away information by reducing all counts to edges with the same weight [2]. Moreover, there are good community-detection methods for weighted networks (see Section 6.5.1) and by using "appropriate" visualization tools, we can present figures that clearly show different edge weights (see Section 8.3).

Chapter 4

How should one construct a similarity matrix?

The edge weights in roll-call voting networks, giving a measure of similarity between pairs of legislators, correspond to the elements of a *similarity matrix* **S** [106]. In this chapter, we examine how to map roll-call voting data to a similarity matrix, which is an important issue that has received little attention [68]. Scholars have constructed similarity matrices in multiple ways [4, 24, 68, 124], and the resulting network can be affected strongly by the choice of method [73].

4.1   Standard assumptions

There are some assumptions, which we call "standard" assumptions, that have been used in several papers. A similarity matrix **S** has size $n \times n$, where n is the number of legislators in a parliament. Legislators i and j who vote the same way as each other are more similar than legislators whose votes conflict. The more RCVs legislator i and j agree upon, the more similar they are and the higher their similarity score S_{ij}. We normalize all similarity scores so that $S_{ij} \in [-1, 1]$ for all i and j, given by

$$S_{ij} = \frac{1}{b_{ij}} \sum_{k=1}^{K} c_{ijk}. \tag{4.1}$$

In equation 4.1, K is the total number of considered bills (see Sections 4.2 and 4.4), c_{ijk} is a certain score based on the outcome of the vote (see Section 4.5), and b_{ij} is used to normalize the similarity score (see Section 4.6).

4.2   Which bills should we include?

Option 1: Include all bills

The most straightforward option is to include all bills regardless of how legislators voted on it. We take this approach, because we want to do an exploration that includes all data instead of throwing some of it away.

Option 2: Only include bills in which the size of the minority surpasses some threshold

One could argue that a bill only provides useful information about voting similarity when there is some substantial disagreement among legislators [68]. Hence, several authors [66, 68, 102]

excluded unanimous RCVs from their analysis. In [124], the threshold was set in such a way that the minority had to be at least 3% of all legislators.

Option 3: Only include bills in which the total turnout surpasses some threshold

One could also argue that a bill only becomes informative when a certain percentage of the legislators actually vote. For example, Attina (1990) only considered bills in which at least 1/3 of the legislators voted [6].

4.3   Which legislators should we include?

Option 1: Include all legislators

We choose to include all legislators in our computations, because we want to do an exploration that includes all data instead of throwing some of it away.

Option 2: Exclude "inactive" legislators

Some authors have ignored "inactive" legislators in their analysis. For instance, Amelio and Pizzuti (2015) removed all legislators who never voted [4], Lyons and Lacina (2009) excluded legislators who participated in very few RCVs [67], and Metz and Jäckle (2016) only examined the top 50% of the most active legislators. If a legislator leaves parliament (e.g. by resigning or dying), Dal Maso et al. (2014) considered the replacement legislator and departing legislator as being a single node [24]. We treat both legislators separately by having one node for each of them (see Section 8.2).

4.4   How should we deal with absence?

Option 1: Treat absence as abstention

The most common approach is to treat absence as abstention [4, 68]. Hence, a legislator who does not vote for a particular bill is treated as a legislator who abstains for that bill.

Option 2: Ignore bills for which at least one of the two legislators is absent

Another approach is to treat absences as a lack of data [120] and exclude RCVs on which at least one of the two legislators is absent [66]. For some situations, this implies that we only have a very small number of bills to consider. Depending on the normalization (see Section 4.6), pairs of legislators could get very high agreement or disagreement scores based on only a small number of bills [73].

Option 3: Assign different scores to absence and abstention

As we discuss in Section 4.5, there are authors who assign nonzero similarity scores to abstaining legislators [24, 73]. In this case, it is possible to distinguish between absence and abstention by assigning a neutral score to absence (e.g. a score of 0) and a non-neutral score to abstention [73]. We will use this approach.

17

Option 4: Predict the vote of absent legislators

One could also try to predict what a legislator would have voted if he/she were not absent using a process called *imputation* [49]. Jakulin and Buntine (2004) used different imputations and compared the results. For instance, they considered the scenario in which legislators do not vote because they expect to agree/disagree with the clear majority, because they lack the information to make a decision, or because they do not want to reveal their opinion [49]. For our computations, we do not have any information about the reasons why legislators were absent, so we will not attempt to impute votes of absent legislators.

4.5   How should we assign similarity scores?

There are sixteen types of voting pairs, as every legislator has four options from which to choose: yes, no, abstain, and absence. Other data sets can have more options, but those are usually converted to one of these four [93]. To assign numerical values to these voting pairs, some (potentially arbitary) choice has to be made about the magnitudes of their values [68]. For instance, one has to decide how many points to assign to abstain–abstain pairs. As noted by Macon et al. (2012), this topic has not been considered much [68]. Below, we discuss the options that other scholars have used.

4.5.1   How should we value yes–yes and no–no agreements?

One could argue that a yes–yes pair is "more similar" than a no–no pair, because a no–no pair is not necessarily always an agreement [68]. Lijphart (1963) illustrated this point with an example in which the United States and the Soviet Union could vote against a proposal in the United Nations General Assembly for entirely different reasons, which thus does not appear to be a "full agreement" between those countries [66]. However, two legislators who vote in favour of a bill might also have completely different reasons for doing so. Moreover, whether a "yes" or "no" vote is confirmative also depends on the way a vote is phrased (e.g. "should the UK leave the EU?" versus "should the UK stay in the EU?"). Hence, we do not make a distinction between yes–yes and no–no pairs in our computations.

4.5.2   How should we treat abstention?

If one ignores abstention like Amelio and Pizzuti (2015) [4], the table with similarity scores looks like Table 4.1. However, we do not want to ignore abstention, so we have to explore other approaches. There is no consensus in the literature about how many points one should assign to abstaining legislators [24, 66, 68, 73]. There are many different reasons why a legislator could abstain, such as neutrality, ignorance, external pressure, or lack of time [73, 120], but one generally does not know which reason causes a legislator to abstain on a particular vote. To create a similarity matrix, one has to decide whether an abstention should be interpreted as a meaningful third option or as a special case.

Option 1: Treat abstention as one of the three meaningful options

Several authors [24, 68, 73, 124] considered abstention as one of the three meaningful options that are available to a legislator. For instance, Dal Maso et al. (2014) and Waugh et al. (2012)

	Yes	No
Yes	1	0
No	0	1

Table 4.1: Similarity scores used by Amelio and Pizzuti (2015) [4]

	Yes	Abstain	No
Yes	1	0	0
Abstain	0	1	0
No	0	0	1

Table 4.2: Similarity scores used by Dal Maso et al. (2014) [24] and Waugh et al. (2012) [124]

used an approach (see Table 4.2) in which two legislators get a full score of 1 on a particular bill if and only if they voted the same way [24, 124]. Mendonca et al. (2015) emphasized the disagreement of a yes–no pair (see Table 4.3) by assigning a negative score to yes–no pairs [73].

	Yes	Abstain	No
Yes	1	0	−1
Abstain	0	1	0
No	−1	0	1

Table 4.3: Similarity scores used by Mendonca et al. (2015) [73]

Option 2: Treat abstention as a different option

Lijphart (1963) made a distinction between "full agreement" and "partial agreement" [66]. If both legislators abstain, they still get the same similarity score as a yes–yes pair and no–no pair, but half of a score is given to yes–abstain pairs (see Table 4.4). Voeten (2000) analyzed the United Nations General Assembly (UNGA) and treated abstention as a "no" vote [120]. His justification was that UNGA resolutions are not binding, so it only matters whether or not a country is supporting a resolution. By contrast, Mendonca et al. (2015) compared both options (i.e. treating abstention as an equally meaningful option (see Table 4.3) and treating abstention as half of an agreement (see Table 4.5)), and concluded that the agreement scores were very similar [73]. Consequently, they decided to present results for one table only. We treat abstention differently from a yes/no vote (see Table 4.6). We assign every yes–yes pair and no–no pair a maximum score of 1, and we assign yes–no pairs the minimum score of −1. Every pair that contains exactly one abstaining legislator is assigned a score of −0.25 and abstain–abstain pairs receive a score of 0.25. If one of the legislators is absent, we give a neutral score of 0 (see Table 4.6).

4.5.3 Other approaches

There are other approaches that one can follow to construct a similarity matrix based on roll-call voting. We briefly discuss two of them here: one based on information theory, and

	Yes	Abstain	No
Yes	1	0.5	0
Abstain	0.5	1	0
No	0	0	1

Table 4.4: Similarity scores used by Lijphart (1963) [66]

	Yes	Abstain	No		Yes	Abstain	No
Yes	1	0.5	-1	Yes	1	-0.5	-1
Abstain	0.5	0.5	0.5	Abstain	-0.5	-0.5	-0.5
No	-1	0.5	1	No	-1	-0.5	1

Table 4.5: Similarity scores used by Mendonca et al. (2015) [73]

the other on exponential weights. We do not use these approaches for our computations, but we will use the information-theoretical notions when we discuss ways to compare "network partitions" in Section 5.4.

Similarity based on information theory

Jakulin and Buntine (2004) computed the distance between two legislators by using notions from information theory and treating similarity as a measure of *interdependence* [49]. For an event X with possible outcomes $\{x_1, ..., x_K\}$ having probabilities $\{p(x_1), ..., p(x_K)\}$, Shannon's *entropy* function [110] is

$$H(X) = -\sum_{k=1}^{K} p(x_k) \log_2 p(x_k). \tag{4.2}$$

The *joint entropy* $H(X, Y)$ is given by

$$H(X, Y) = -\sum_{k=1}^{K} \sum_{l=1}^{K} p(x_k, y_l) \log_2 p(x_k, y_l). \tag{4.3}$$

Jakulin and Buntine (2004) considered two different probabilistic models to describe the voting process: one model assumes independence between the legislators' RCVs and the other model assumes dependence between their RCVs. If there is independence between a set X of RCVs of legislator i and a set Y of RCVs of legislator j, then information about how i has voted does not increase the probability that we can correctly predict what j has voted (and vice versa). Hence, in the model assuming independence, we have

$$H(X, Y) = H(X) + H(Y). \tag{4.4}$$

The difference between the two models' entropies is called the *mutual information* [49]:

$$I(X; Y) = H(X) + H(Y) - H(X, Y). \tag{4.5}$$

The minimum value of 0 is attained when there is no correlation between the voting behaviour of legislators i and j. The higher the mutual information between two legislators, the greater

	Yes	Abstain	No	Absent
Yes	1	−0.25	−1	0
Abstain	−0.25	+0.25	−0.25	0
No	−1	−0.25	1	0
Absent	0	0	0	0

Table 4.6: Similarity scores for our computations

their dependence and their "similarity" [49]. That is, a higher value of $I(X;Y)$ means that we have more knowledge about the vote of legislator j given the vote of i. However, if two legislators always vote in an opposite way (Jakulin and Buntine (2004) only considered "yes" and "no" votes), the maximum value for $I(X;Y)$ is also reached. Hence, "similarity" has a different meaning in this approach than the one that we defined in Section 4.1 (i.e. "legislators i and j who vote the same way as each other are more similar than legislators whose votes conflict").

Exponential weighting

Bertozzi and Flenner (2016) used *exponential weighting* to consider voting records of the United States House of Representatives [12]. They defined the voting record of legislator i by a vector $\mathbf{x}_i \in \mathbb{R}^K$, where K is the number of bills, a "yes" vote is represented by 1, a "no" vote by −1, and each abstention by a 0 in the voting vector. They defined the similarity between two legislators by an adjusted Gaussian similarity function.

4.6 How should we normalize similarity scores?

Option 1: Divide each similarity score by the number of bills in which both legislators vote yes/no

Waugh et al. divided each similarity score by the number of bills in which both legislators vote yes or no [124]. Because we assign a non-zero similarity score to abstaining legislators (see Section 4.5), it is necessary for us to take a different approach. Otherwise, the normalized similarity score of two legislators could be larger than 1 if they both abstained most of the time.

Option 2: Divide each similarity score by the number of bills in which both legislators are present

Rombach et al. (2014) divided each score by the number of bills in which both legislators were present [106]. This option, if combined with assigning a neutral score to absence, is the same as ignoring a bill when at least one of the two legislators is absent. The disadvantage of this approach is that legislators who almost never show up can still get a high normalized similarity score. For instance, a pair of legislators who vote the same way on all ten bills are assigned the same normalized voting score as a pair of legislators who vote the same way once and did not show up the other nine times.

Option 3: Divide each similarity score by the total number of bills in a session

Several authors divided each similarity score by the total number of bills in a session [24, 68, 73]. By using this approach, legislators who do not show up regularly end up having a similarity score close to 0. Because we assign disagreeing legislators a negative similarity score, it seems reasonable to assign absent legislators a neutral score of 0. Hence, we will use option 3 and avoid assigning artificially high agreement scores to legislators with low attendance [68]. A full weight of 1 can only be obtained by two legislators if they are always present, vote exactly the same for all bills, and never abstain [24]. Macon et al. (2012) noted that it does not matter for detecting communities whether one works with a uniformly normalized similarity matrix (in which every similarity score is divided by the same constant b), or whether one uses a similarity matrix that is not normalized at all [68].

Chapter 5

Community detection

5.1   What is a community?

There is no consensus about the definition of a community [31, 33, 98]. One (classical) point of view is that communities are relatively densely connected sets of nodes (e.g. compared to the mean edge density [31]), which are sparsely connected to other dense sets of nodes [33]. Another approach is to assume that nodes of the same community are more likely to form edges with nodes inside their community than with nodes outside their community [33].

Fortunato (2010) mentioned several descriptions of a community, including *local* ones (focusing on one subgraph and neglecting the rest of the network), *global* ones (focusing on an entire network), and descriptions based on *node similarity* (focusing on some reference node property) [31]. Fortunato and Hric (2016) argued that a precise *a priori* definition of community is not really needed, as most community-detection techniques (see Section 5.3) do not require it [33]. Instead, communities are often "algorithmically defined" as the output of a community-detection algorithm [31, 98].

In network science, *hard partitioning* or *hard clustering* refers to the process in which one divides the nodes of a network over *disjoint* sets (communities). This division is called a *partition*. We do not consider a division (called *cover*) into overlapping sets of nodes, which is the result of *soft partitioning* [33]. In our computations, we do not fix the number q of communities *a priori* (for a discussion of the specification of q, see [90]). In real networks there can be *hierarchical structure* with subcommunities nested in communities [60, 62]. Consequently, there can be several meaningful partitions that correspond to different hierarchical levels and it might be hard to discriminate between them [31].

5.2   Communities in real-life networks

In real-life networks, there can be high concentrations of edges within groups of nodes, and low concentrations between them [31]. Scholars have used community-detection tools to uncover unknown structures that might have special functions [14, 31, 98]. For instance, Flake et al. (2002) identified communities of highly related webpages by applying community detection to the World Wide Web [30]. Traud et al. (2012) used Facebook data to find communities in friendship networks and found that at most U.S. universities among the 100 they analyzed, communities are mainly organized by class year, whereas at a few universities, the community assignment is based predominantly on dormitory residence [116]. Scholars have also studied communities in political networks [21, 104]. Porter et al. (2005) found strong links between different committees and subcommittees in the United States House of Representatives (e.g. between the House Rules Committee and the Select Committee on Homeland Security in the 107th Congress) [97]. Macon et al. (2012) identified communities of countries in the United

Nations General Assembly based on vote similarity and found, for instance, communities that correspond to "Western" and "Eastern" communities [68].

5.3 Community-detection methods

There are many different methods to identify communities in networks. In the articles [31, 33, 98], one can find an extensive overview of methods such as spectral methods, dynamic methods, hierarchical clustering methods, and many more alternatives. *Spectral* methods use spectral properties, such as the eigenvalues of an adjacency matrix or of other matrices. *Dynamic* methods use dynamical processes such as random walks to detect communites in a network. *Hierarchical clustering* methods construct a hierarchy of communities based on some "similarity" of nodes.

Even though scholars have tried to compare different community-detection methods [42, 73], there are no clear guidelines on how to compare these methods [31, 98], and there is no universal best method [33]. Therefore, Porter et al. (2009) argued that it is best to use multiple methods to analyze real-world networks such that one can have more trust in the partitions that are similar across multiple methods [98]. In Chapter 6, we focus on one popular method: the *optimization* of a quality function called *modularity*.

5.4 Comparing partitions

One might want to compare the results of different community-detection methods to each other or to some *metadata* [42, 115, 126]. For instance, Newman and Clauset (2016) discussed how metadata can be used to improve community detection [88]. Their method tries to learn whether a correlation exists between some metadata and the detected communities, and the method then uses the metadata if it contains useful information [88]. Hric et al. (2016) introduced a generative model to assess the quality of metadata in its capacity to predict the placement of edges in a network [47]. One can also use their method to improve the prediction of missing nodes. Cherepnalkoski and Mozetič (2016) applied community detection to a parliamentary network and compared the detected communities to information about the legislators [21].

Measures to compare partitions

There are different measures that one can use to compare partitions to each other; these include measures based on *pair counting* or *information theory* [53, 115].

Measures based on pair counting

Measures based on pair counting consider the node pairs that are in the same communities and the ones that are in different communties in two partitions. As before, let V be the set of nodes $\{1, 2, ..., n\}$, let $X = \{X_1, ..., X_r\}$ be a partition of V into r communities, and let $Y = \{Y_1, ..., Y_s\}$ be a partition into s communities. We define the following numbers [53]:

$$a_{11} = \text{number of node pairs in the same communities as each other in } X \text{ and } Y. \tag{5.1}$$

$$a_{00} = \text{number of node pairs in different communities from each other in } X \text{ and } Y. \tag{5.2}$$

$$a_{10} = \text{number of node pairs in the same community in } X \text{ and in different ones in } Y. \tag{5.3}$$

$$a_{01} = \text{number of node pairs in different communities in } X \text{ and in the same one in } Y. \tag{5.4}$$

One pair-counting measure is the Rand index R [100], which indicates the fraction of node pairs on which both partitions agree (i.e. the node pairs that are either in the same community for both partitions or in different communities for both partitions). The index lies between 0 and 1 (where 1 indicates full agreement) and is given by

$$R(X,Y) = \frac{a_{11} + a_{00}}{a_{11} + a_{01} + a_{10} + a_{00}} = \frac{a_{11} + a_{00}}{\binom{n}{2}}. \tag{5.5}$$

Another pair-counting measure is the *Jaccard index* [48]

$$J(X,Y) = \frac{a_{11}}{a_{11} + a_{01} + a_{10}}. \tag{5.6}$$

The Jaccard index divides the node pairs that are in the same community for both partitions by the number of node pairs that are in the same community for at least one partition [31]. The *Jaccard distance*, which satisfies the standard metric axioms (non-negativity, symmetry, and the triangle inequality) [65], measures dissimilarity between partitions:

$$d_J(X,Y) = 1 - J(X,Y) = \frac{a_{01} + a_{10}}{a_{11} + a_{01} + a_{10}}. \tag{5.7}$$

Information-theoretical measures

In Section 4.5.3, we discussed information-theoretical similarity measures. We can also use an information-theoretical framework to compare partitions [26, 42, 73]. The more similar two partitions X and Y are, the less information one needs to infer partition X from partition Y or vice versa [53]. We can measure this amount of information with the *normalized mutual information* (NMI) [114]:

$$\text{NMI}(X,Y) = \frac{2\,I(X;Y)}{H(X) + H(Y)}, \tag{5.8}$$

where $H(X)$ is the Shannon entropy defined in equation (4.2) and $I(X;Y)$ is the mutual information defined in equation (4.5). Another information-theoretical measure is the *variation of information* (VI) [72], which is the sum of the information needed to describe X given Y and vice versa [53]:

$$\text{VI}(X,Y) = H(X) + H(Y) - 2\,I(X;Y). \tag{5.9}$$

The normalized variation of information (NVI), which we will use in Section 8.2, is given by [58]

$$\text{NVI}(X,Y) = 1 - \frac{I(X;Y)}{H(X,Y)}. \tag{5.10}$$

Both the VI and NVI satisfy the properties of a metric [53, 58, 118].

5.5　Robustness of a partition

It is important to examine whether detected communities are *significant*. There are different types of significance [76]. We measure how *significant* communities are by measuring their *robustness* against some (random) perturbation (e.g. the addition or removal of edges in a network) [53]. If an optimal partition of the original network is relatively similar to an optimal partition of a perturbed network (indicated by an NVI score close to 0, for instance), one considers the original partition to be relatively robust. Hence, the robustness approach follows the assumption that a significant partition should not be altered by a small modification of a network [60].

Scholars have used different definitions of robustness and different perturbation methods [38, 53, 54, 60, 68, 76]. Gfeller et al. (2005) changed the weight of every edge e_{ij} by adding some number (equally distributed) between $-\sigma w_{ij}$ and σw_{ij}, where w_{ij} is the weight of edge e_{ij} and $\sigma \in (0,1)$ is fixed [38]. Karrer et al. (2008) rewired a fraction of randomly-chosen edges and compared partitions of the original network to the partitions of perturbed networks [53]. Macon et al. (2012) used different values for the resolution parameter (see Section 6.4.3) to examine the robustness of network partitions [68].

We follow Lambiotte (2010) [60] and use three types of modifications to measure robustness. First, we modify our networks by adding/removing some weight independently to every element of our adjacency matrix. Second, we run our community-detection algorithm R times and use a random ordering of the nodes, so that we can compare the results between different runs. Third, we tune the "resolution parameter" γ (see Section 6.4.3) and compare optimal partitions for different values of γ. We discuss these three modifications in more detail in Section 8.2.

Chapter 6

Modularity optimization

6.1 Quality functions

The most popular community-detection method is the *optimization* of a *quality function*, which assigns a score (e.g. a real number) to each possible partition of a network [33]. Comparing the scores provides one way — though a rather imperfect way — of comparing the quality of the partitions to each other. Within a particular network, partitions with higher (or lower, depending on the quality function) scores are considered to be "better" partitions [31].

However, we show in Section 6.4.3 that these partitions may not correspond to the partitions that one would consider to be the most "intuitive" or the most "preferred" ones. Moreover, whether a partition is considered to be "good" depends on the definition of a community that one uses [31]. One can also use quality functions to compare the output of other community-detection methods. We focus on the most popular quality function, called modularity [89], but for a discussion of other quality functions (such as *performance* and *coverage*), see [31].

6.2 Modularity of a partition

In a very influential paper [89], Newman and Girvan (2004) defined the notion of *modularity* (for an unweighted network) as the fraction of edges that are incident to nodes within the same community minus some expected value of this fraction, if edges were assigned randomly in a network with the same expected degree distribution:

$$Q = \sum_i (E_{ii} - R_i^2) = \text{Tr}[\mathbf{E}] - ||\mathbf{E}^2||\,. \tag{6.1}$$

In equation (6.1), symmetric matrix $\mathbf{E}$ is the normalized adjacency matrix of a network of communities (i.e. in which each community is treated as a single node). Symbol R_i denotes a row of matrix $\mathbf{E}$ (i.e. $R_i = \sum_j E_{ij}$), which is the fraction of edges that are incident to nodes in community i. Notation $||\mathbf{X}||$ denotes the sum of the elements of the matrix $\mathbf{X}$. The trace $\text{Tr}[\mathbf{E}] = \sum_i E_{ii}$ gives the total fraction of edges that connect nodes within the same community. If there is a single community that contains all nodes, Tr $[\mathbf{E}] = 1$ and $Q = 0$. This single community does not give useful information about community structure.

We use a different notation for Q [85]:

$$Q = \frac{1}{2W} \sum_{ij} B_{ij} \delta(C_i, C_j)\,. \tag{6.2}$$

In equation (6.2), W is the total weight of the network (see Equation (2.4)) and C_i denotes the community to which node i is assigned. The Kronecker delta $\delta(C_i, C_j)$ is equal to 1 if node i and node j are assigned to the same community and 0 otherwise. The scalar B_{ij} is an element of *modularity matrix*

$$\mathbf{B} = \mathbf{A} - \mathbf{P}, \tag{6.3}$$

where $\mathbf{P}$ is a *null model*. This null model could, for instance, indicate some expected value of the density of edges (see Section 6.3).

Many authors have used modularity [32, 82, 84, 85] to measure how well a partition divides a network into communities. The modularity Q measures the density of edges within communities as compared to the expected density in some null model (e.g. a network with the same expected node-strength distribution). The value of Q lies in $[-1, 1]$ (for networks with nonnegative weights). For a given network, higher values of Q correspond to partitions in which there are more/stronger edges within communities than one would expect by chance in a certain null model [98]. If $Q = 0$, there are exactly as many edges within communities as one would expect by chance [82].

6.3 Null models

In equation (6.3), we introduced a null model $\mathbf{P}$ that is used to take account of the expected edge weight within communities. The value of modularity Q only has a meaning once $\mathbf{P}$ has been defined [32]. There are many possible null models [11, 60, 101], and one can incorporate specific network features (such as bipartiteness or signed edges) by choosing a particular one [31]. The matrix $\mathbf{P}$ can be derived by some randomization of the original network [33]. Because we work with undirected networks, we require that $P_{ij} = P_{ji}$ for all i, j. Moreover, $Q = 0$ when there is only one community containing all nodes, so equations (6.2, 6.3) imply that [84]

$$0 = \sum_{ij} B_{ij} = \sum_{ij} [A_{ij} - P_{ij}], \tag{6.4}$$

and

$$\sum_{ij} P_{ij} = \sum_{ij} A_{ij} = 2W. \tag{6.5}$$

Consequently, we consider only null models that have the same total edge weight as our original network. Bazzi et al. (2016) differentiated between "null models" and "null networks" [11]. They defined a null model to be a probability distribution on the set of adjacency matrices. Under a specified null, a null network is the expected adjacency matrix (so the same null network can correspond to different null models) [11]. We do not distinguish between null models and null networks in this dissertation.

Uniform null model

The simplest null model is the *uniform* null model (or *Bernoulli random graph*), in which P_{ij} is a constant [84, 101, 104]. To satisfy (6.5), we have for all i, j that

$$P_{ij}^* = \frac{2W}{n(n-1)}. \tag{6.6}$$

A weakness of the uniform null model is its degree distribution. Many real-life networks have right-skewed degree distributions, but the uniform null model has a binomial degree distribution (or Poissonian as $n \to \infty$) [84]. These degree distributions are very different, making the uniform null model a suboptimal model of most real networks [31]. However, Richardson et al. (2009) noted that the uniform null model might be appropriate for the study of similarity networks (such as voting-similarity networks), because of the relative uniformity of edge strength among the nodes and because of the high network density [104]. Bazzi et al. (2016) extended the notion of a uniform null model to a uniform null *network* [11].

Newman–Girvan null model

The most common null model is the *Newman–Girvan (NG)* null model [89]

$$P_{ij}^{**} = \frac{k_i k_j}{2W}, \tag{6.7}$$

which satisfies equation (6.5), as $\sum_i k_i = 2W$. The NG null model gives the expected edge weights in a random network in which edges are placed independently under the condition that the expected node strengths in the random network match the actual node strengths in the original network [68]. In other words, P_{ij}^{**} is an adjacency-matrix element yielding a network in which the probability that there is an edge e_{ij} depends only on k_i and k_j (which are independent of each other) [84]. We can combine equations (6.2, 6.3, 6.7) to obtain Newman–Girvan modularity (Q_{NG}):

$$Q_{\mathrm{NG}} = \frac{1}{2m} \sum_{ij} \left[\left(A_{ij} - \frac{k_i k_j}{2W} \right) \delta(C_i, C_j) \right]. \tag{6.8}$$

Unless specified otherwise, we refer to Newman–Girvan modularity whenever we discuss modularity in the remainder of this dissertation.

6.4 Modularity optimization

One does not only use modularity to compare the quality of network partitions *a posteriori*, but also as an objective function that one wants to optimize *a priori* [14, 89]. There are many community-detection algorithms that try to *optimize modularity* (see Section 6.5), and a partition with the highest modularity $Q_{\max}$ is called an *optimal* partition [39]. We use modularity optimization for our computations (see Section 8.2), so we try to find a partition in which the total intra-community edge weights are optimized relative to the null model [68]. The maximum modularity $Q_{\max}$ depends on the network size (it usually is larger for larger n [39]), so we only compare (maximum) modularity values between networks of the same size. Because $Q = 0$ for a single-community partition, we know that [31]

$$Q_{\max} \geq 0. \tag{6.9}$$

6.4.1 Computational complexity

Brandes et al. (2008) proved that it is an NP-hard problem to find the partition with $Q_{\max}$ [18]. Because the number of possible partitions increases at least exponentially with n, it is not feasible to check all possible partitions [32], even for networks of moderate size. (For instance, in our computations with $n = 840$, there are B_{840} possible partitions, where B_{840} is a Bell number). Hence, we need heuristics or approximation algorithms (see Section 6.5) that can obtain partitions with Q close to $Q_{\max}$ in a reasonable amount of time [14, 98]. Approximation algorithms rarely return an optimal partition, and they are often non-deterministic, as they return different solutions to the same problem [31]. Even though the output of a modularity-optimization heuristic usually has a modularity value lower than the theoretical maximum, we still denote the largest modularity value that we obtain for a network by $Q_{\max}$.

6.4.2 Strengths of modularity optimization

Although other community-detection methods often implicitly assume some kind of null model, modularity-optimization methods make the null model explicit and clear [84]. Consequently, one has more understanding of the method's implications and more control to adjust them, which makes modularity a very intuitive measure [84]. In Section 6.5.1, we discuss the strengths of one particular modularity-optimization heuristic (the Louvain method).

6.4.3 Weaknesses of modularity optimization

Modularity-optimization methods also have several weaknesses. For instance, the maximum modularity Q_{max} depends strongly on the size of a network, the number of communities, and their inter-connectedness [39]. Consequently, one should generally not compare modularity values between distinct networks of different sizes. (There can also be issues with comparing values of Q for networks of the same size, so it is often better to focus on the obtained partition than on the modularity value itself).

In this section, we discuss two major weaknesses of modularity optimization. The problem of the *resolution limit* implies that the optimal partition (with Q_{max}) may not always correspond to the most intuitive partition [32]. The problem of *near-degeneracy* means that there can be (exponentially) many different partitions with Q very close to Q_{max} [39]. Because of these weaknesses (and other weaknesses as well, see [33, 87]), one should be careful with interpreting high-modularity partitions [39].

The resolution limit

Fortunato and Barthélemy (2007) showed that modularity optimization has a resolution limit [32], in that modularity-optimization heuristics might miss communities that are smaller than a certain scale (which depends on n and on the extent of interconnectedness of the communities). Consequently, heuristics might favor partitions that combine communities into larger communities and return partitions that are different from the "intuitive" community assignment. Therefore, the partition with highest modularity value might not always be the desired partition. For instance, one can show that for a ring of K_4 cliques (in which there is a single edge between a clique and two neighbouring cliques), the partition with two cliques in one community has higher modularity than a partition in which every clique is assigned to its own

community (which is counterintuitive to the intuitive idea of a good community assignment) [32, 33].

We describe the resolution limit by using a derivation from [31], adjusted to a weighted network. (Note that weighted networks usually have less severe issues with the resolution limit than unweighted networks [39]). Consider two communities A and B that have a total strength of $k_A = \sum_{i \in A} k_i$ and $k_B = \sum_{i \in B} k_i$ and suppose (for simplicity) that $k_A \approx k_B = \kappa$. The expected total edge weight between community A and B is

$$w'_{AB} = \frac{k_A k_B}{2W}. \tag{6.10}$$

If community A and B are merged, the difference in modularity from the change is

$$\Delta Q_{AB} = \frac{w_{AB}}{W} - \frac{k_A k_B}{2W^2}, \tag{6.11}$$

where w_{AB} is the total edge weight between A and B. If there is a single edge with unit weight between A and B (i.e. $w_{AB} = 1$), one would intuitively expect that the two communities are not merged. However, if

$$\frac{k_A k_B}{2W} < 1,$$

then $\Delta Q_{AB} > 0$, which makes a partition with A and B in the same community a "better" (i.e. higher-modularity) partition. Hence, given that the total edge weight between communities A and B is larger than 1, one gets a higher modularity value when A and B are in the same community if

$$\kappa < \sqrt{2W}. \tag{6.12}$$

Consequently, if the output of a modularity-optimization heuristic includes communities with a total node strength of the order of $\sqrt{W}$ or smaller, we do not know *a priori* whether these communities are single communities or combinations of smaller communities [31].

There are several ways (that are not necessarily inconsistent with each other) to describe the origin of the resolution-limit problem. For instance, some authors [32, 50] described it as a fundamental consequence of the additive form of equation (6.2), which implies that one has to find the best tradeoff between the number of communities and the modularity value of each term in order to maximize modularity. Other authors [31, 33] explained the problem in terms of the null-model assumption that each node can interact with every other node. They argued that it would be more reasonable (especially for large networks) to assume that each node only interacts with a portion of the network.

How should one deal with the resolution limit?

Scholars have proposed different ways to deal with the resolution limit. For instance, one can use multiple heuristics and find a consensus of the high-modularity partitions [28]. One can also perform further subdivisions of the obtained communities by treating each community as a separate network and applying modularity optimization again [32]. Fortunato (2010) argued that the latter procedure (i.e. further subdivision) is not reliable because it requires an arbitrary stopping criterion and because the "local" modularities within communities are inconsistent with each other due to different null models [31]. Moreover, by treating a community as a separate network, one ignores the inter-community edges, so one is examining

a network that is different from the original network, which could affect the results of any subsequent modularity-maximization heuristic [85].

To deal with resolution limit, we consider "multi-resolution modularity" [60], which means that we include a *resolution parameter* γ into the modularity matrix [101]:

$$\mathbf{B} = \mathbf{A} - \gamma \mathbf{P}. \tag{6.13}$$

By considering different values of γ, one can "zoom in and out" to find communities at different resolution scales [4, 98], thereby examining different "background" levels of agreement between legislators. If one increases γ, one places more weight on the randomized term P_{ij}, so it becomes less advantageous to assign i and j to the same community [60]. Consequently, the methods finds smaller communities for larger values of γ. By contrast, a smaller value of γ implies that there is more weight on the observed term A_{ij}, so the method favors larger communities [87].

An issue with using multi-resolution modularity is that one does not *a priori* know the "right" value of γ. Hence, one can compare partitions for a range of resolution parameters and detect the "robust" communities [60, 68]. In section 8.2, we describe how we measure the robustness of the detected partitions by combining several statistical tests [60]. However, we cannot guarantee that we find all small communities [5], as this method does not completely solve the resolution-limit problem [33]. Recently, Newman (2016) described a process in which one repeatedly maximizes modularity (using different heuristics), and one recalculates γ until convergence is achieved [87]. He found that this method to find the correct value of γ converges quickly, but that convergence (or correctness) cannot be guaranteed [87].

The near-degeneracy problem

Good et al. (2010) showed that the modularity function exhibits extreme near-degeneracies, as there can be exponentially many (or more) distinct partitions whose modularity values (local optima) are very close to the global maximum $Q_{\max}$ [39]. In other words, the modularity function is typically not strongly peaked around $Q_{\max}$ for practical examples. They also showed that there may be very little overlap between these partitions (despite the closeness of their modularity values), and that the global (theoretical) maximum $Q_{\max}$ can often not be attained [39].

As the number of communities q gets larger, there are (exponentially) more non-optimal ways to combine them, which makes it harder to find a partition with $Q_{\max}$. Because there are often many partitions with a relatively high modularity score, it is unsurprising that a lot of heuristics tend to return good partitions [39].

How should one deal with the near-degeneracy problem?

One can try to circumvent the near-degeneracy problem in several ways. For instance, one can construct a consensus partition across multiple optimizations [10, 25] to construct a representative partition. We use this approach in Section 8.2. One can also use generative models to identify community structure [39] or one can apply multiple computational heuristics to compare the identified communities [98, 124]. In the latter scenario, one could decide to pick a partition with the highest modularity [124] or to sample a partition from all high-modularity partitions [39].

32

6.5 Heuristics for modularity optimization

There are many distinct heuristics that can find partitions with modularity values fairly close to $Q_{\max}$ in a reasonable amount of time [31]. Sometimes heuristics find partitions that vary a little bit [124] and sometimes they can vary a lot [39]. For our application, we use a locally-greedy heuristic (the Louvain method) that we describe in Section 6.5.1. Other heuristics include other greedy techniques [83], simulated annealing [55, 101], spectral methods [53, 84, 85, 104], extremal optimization [15, 29], and many more (see [31, 39, 98, 124]).

6.5.1 Louvain method

Blondel et al. (2008) introduced a locally greedy modularity-optimization method [14], that is now known as the *Louvain method*, and that has been used by many scholars [21, 27, 39, 68, 78, 91, 106, 124].

The algorithm

We give pseudocode for the Louvain method in Algorithm 1. The algorithm tries to find a network partition with high modularity. The algorithm has two phases. Before phase I, every node i is assigned its own community (so that there are n communities initially). During phase I, the algorithm moves nodes between communities to obtain a locally-maximal increase of modularity. This phase can be repeated often, so nodes can be considered several times. Whenever phase I stops, the algorithm has found a local maximum, as no move from a single node can improve modularity. Phase II then starts. During phase II, a new network is created in which every node represents one of the communities found during phase I. Once this new network has been created, the entire algorithm is applied again until modularity can no longer be increased.

Number of rounds

Every time phase II ends, we call this the end of a "round" of the algorithm. Within a round, phase I is repeated as long as $\Delta Q > 0$. Hence, in every round, the algorithm first optimizes modularity by locally changing communities (phase I), and then it creates a new network of the communities found (phase II). If the modularity value increases during round t, the algorithm starts a new round $t + 1$ after phase II is over. In every round in which modularity increases, the number of communities must decrease, as there is at least one node (representing a community) moved into another community. Consequently, most of the computing time (which is linear in n [31]) is used in the first round. The number of rounds is usually very small [14].

Output of the algorithm

Different runs of the Louvain algorithm can give (slightly) different partitions [106]. In addition to considering the final output of the algorithm, one can also consider the intermediate divisions after each round. Those divisions correspond to local maxima of modularity (at different scales), as it is not possible to increase modularity by moving one node from a community to a neighbouring community [14].

Algorithm 1 Louvain method for community detection (pseudocode)

Input: Weighted adjacency matrix $\mathbf{A}$ (size $n \times n$).
Output: Partition of n nodes and a modularity value Q

Set "increase" to FALSE
for all nodes i **do**
    Assign each node its own community
end for

Phase I:
Set "repeat" to FALSE
for all nodes i **do**
    for all neighbours j of i **do**
        Evaluate gain of modularity (ΔQ) for moving node i to the community of node j.
        Let j^* be a node for which ΔQ is maximum (if there is a tie, use a tie-breaking rule).
        if $\Delta Q > 0$ for j^* **then**
            Node i is moved to community of node j^*.
            Set "repeat" to TRUE
            Set "increase" to TRUE
        end if
    end for
end for

if "repeat" is TRUE **then**
    Repeat phase I.
end if

Phase II — Creating a new network:
for all communities q_i **do**
    Create node N_i to represent community q_i
end for
for all pair of nodes N_i, N_j where $i \neq j$ **do**
    Set the edge weight equal to the sum of edge weights between community q_i and q_j.
end for
for all nodes N_i **do**
    Create a self-loop with weight equal to the sum of edge weights within community q_i.
end for

if "increase" is TRUE **then**
    Use this new network as an input of the Louvain method algorithm.
end if

A user can choose to construct a hierarchical structure in which each level is given by the intermediate division after a round. As noted by Fortunato and Hric (2016), "low" level partitions have less unnatural community mergers (e.g. a ring of cliques can be detected this way), but these partitions have lower values of Q than the "top" levels [33]. Hence, users sometimes have to choose between the "optimal" partition (with highest modularity value) and the most "intuitive" partition.

Increase in modularity

One can easily compute the gain ΔQ in modularity that occurs from moving an isolated node x into a community Y. We repeat the definition of modularity with a Newman–Girvan null model and resolution parameter γ (equation 6.8) here:

$$Q_{NG} = \frac{1}{2W} \sum_{ij} \left[\left(A_{ij} - \gamma \frac{k_i k_j}{2W} \right) \delta(C_i, C_j) \right].$$

By moving node x to community Y, the difference in modularity comes only from the terms in which $i = x$ and $j \in Y$ or $i \in Y$ and $j = x$. By symmetry, we consider only one of those options (i.e. $i = x$ and $j \in Y$) and add the contribution twice:

$$\Delta Q_{NG} = \frac{1}{2W} \sum_{j \text{ in } Y} 2 \left[A_{xj} - \gamma \frac{k_x k_j}{2W} \right] = \frac{1}{W} \left[k_{x,Y} - \gamma \frac{k_x \sum_Y}{2W} \right], \tag{6.14}$$

where k_x is the strength of node x, the symbol $k_{x,Y}$ denotes the weight of all edges incident to x and a node in Y, the symbol $\sum_Y$ denotes the sum of all weights of edges incident to nodes in Y, and W is the sum of all edge weights in the entire network. Equation (6.14) is equal to equation (2) in [14].

Comparison to other modularity-optimization heuristics and to other community-detection methods

Blondel et al. (2008) argued (in the paper where the method was introduced) that the steps of the Louvain method are intuitive and easy to implement [14]. Several scholars have compared the Louvain method to other modularity-optimization heuristics [14, 21, 31, 33, 98, 124]. Most of them confirmed that the Louvain method is computationally very efficient, which makes the method suitable for large networks [21, 31]. Waugh et al. (2012) compared the Louvain method to eight other modularity-optimization heuristics and showed that the Louvain method was the one that found the best partition (with highest Q) in more networks than any other heuristic [124].

Besides the weaknesses that apply to all modularity-optimization heuristics (i.e. the resolution limit [32] and near-degeneracy problem [39]), scholars have also found some weaknesses of the Louvain method or reasons to prefer other heuristics (or other community-detection methods). Porter et al. (2009) stated that other modularity-optimization heuristics such as spectral methods [84, 85, 104] and simulated annealing [55, 101] provide more sophisticated means to find partitions with high modularity [98]. This observation was confirmed by Good et al. (2010) for simulated annealing [39]. Lancichinetti and Fortunato (2009) [61] compared the Louvain method to a community-detection algorithm (that is based on random walks)

called Infomap [107]. They concluded that the Infomap method appears to beat the Louvain method in terms of accuracy on "GN benchmark graphs", "LFR benchmark graphs", and random graphs (for a description of these benchmark graphs, see [61, 62]).

6.6 Significance of high-modularity partitions

Optimization of modularity is supposed to deliver the most modular partition. However, the presence of a high-modularity partition does not always guarantee that there is a clear community structure [10, 31, 53]. For instance, Guimerà et al. (2004) showed that there are high-modularity partitions in random networks (such as the Erdös–Rényi graph) [41]. This is counterintuitive because Erdös–Rényi graphs are not expected to have community structure. (In fact, the notion of modularity is based on the assumption that (some types of) random networks do not have community structure [33]). These non-negative values of Q in random graphs are due to fluctuations in the edge distribution, which can occur even if the edge probability $p_{ij} = p$ for all nodes i, j [31]. Bagrow (2012) demonstrated that trees can have arbitrarily high values of modularity (as $n \to \infty$), which is not in line with the common assumption that trees typically do not possess modular strucure [7].

Hence, we want to know whether our obtained partition(s) represent meaningful community structures or whether these structures may have arisen at random [10, 33]. If there is a low probability that a community is generated randomly, we can be more confident that the communities are meaningful. By comparing the community structure to a reliable null model (see Section 6.3), one can estimate the likeliness of a particular community assignment in a random network. For instance, one can use a *p-value* to measure the significance of a partition. This value indicates the fraction of null-model configurations that yields results (e.g. a community assignment) compatible with the ones measured in the studied network [33]. (It is common to use $p \leq 0.05$ as a threshold for significant results [33]).

Newman (2016) claimed that modularity maximization (in its generalized form with resolution parameter γ) is equivalent (for some appropriate choices of models and parameters) to the method of maximum likelihood applied to the "(degree-corrected) planted partition model" [87]. Consequently, he noted that modularity maximization only returns good partitions if communities in a network are statistically similar (which they may not always be) [87]. For our computations, we measure the significance of partitions by considering their robustness, as we discussed in Section 5.5.

6.7 Modularity as measure of polarization

Modularity is a measure that can be used to quantify community structure in legislatures, it can be used to identify which legislators belong to each community, and it can be used to measure how polarized a parliament is [124]. It does not *a priori* assume that political parties are relevant communities, nor does it require information about political ideologies of the legislators [128]. A higher value of Q indicates that a partition has stronger intra-community edge weight and weaker inter-community edge weight (compared to some null model). Waugh et al. (2012) argued that this corresponds to a parliament in which political groups are more polarized [124]. That is, if the parliament is extremely polarized and legislators always vote with their own party, we can expect the strongest possible edges within communities of legisla-

tors of the same party and weaker edges between members of opposite parties. Consequently, there are partitions of legislators with relatively high modularity values. By contrast, if there is less polarization, legislators are more likely to vote in line with others outside their own group, which decreases the modularity [124].

Many scholars have used modularity in several political applications [4, 24, 78, 98, 104, 124, 128]. For instance, Waugh et al. (2012) examined RCVs in the United States and hypothesized that modularity can serve as an early warning of changing group dynamics and that medium levels of $Q_{\max}$ might lead to instability in the United States Congress [124]. Moody and Mucha (2013) examined RCVs and used modularity to show that the current level of polarization in the US Senate is higher than in the past century [78]. Dal Maso et al. (2014) researched RCVs in the Italian Parliament and hypothesized that modularity optimization could typically split (modern democratic) parliamentary networks into two communities corresponding to the governing coalition and the opposition [24].

One can also compute the modularity of a partition that is based on metadata (e.g. the political party or country of a legislator). If one uses political parties to partition a network, the corresponding modularity value is called party modularity Q_{party} [124, 128]. The ratio $Q_{\text{party}}/Q_{\max}$ represents the relative contribution of formal party divisions to total polarization and can be used to infer information about the role of political parties [124]. For instance, if $Q_{\text{party}}/Q_{\max}$ is close to 1, this could suggest that the natural split of the parliamentary network is aligned closely with party affiliation [128], and it could further imply that polarization is usually driven by formal party division [124]. One can use the same method for partitions based on other types of metadata, such as country modularity Q_{country}), based on the country a legislator is from.

Chapter 7

The European Parliament

7.1   Members of the European Parliament

The *European Parliament (EP)* is a powerful legislature that represents approximately 500 million citizens of the *European Union (EU)* [20]. Voters from all EU countries directly elect a national political group, and the number of delegates from a country corresponds roughly to its population size (see Table 7.1) [20]. We consider the 7th European Parliament, which was elected in 2009 and lasted until 2014. In 2009, there were 736 *Members of the European Parliament (MEPs)*, but the 7th Assembly ended with 766 MEPs, partly because Croatia joined the EU. In the current 8th assembly (which we do not consider in this dissertation), there are 751 MEPs.

7.2   Voting procedures

Votes can take place in three different ways: by showing hands, electronically, or by roll-call voting. Because the first two options do not show how each individual MEP has voted, we only include results of the third option in our analysis. In the European Parliament, only certain votes are required to be taken by roll call (e.g. if there is a tight majority, during the election of the European Commission, or upon request of a certain number of MEPs [20]). In 2009, about 1/3 of the votes were taken by roll call [45]. To increase transparency, VoteWatch [121] has called for all votes in the EP to be taken by roll call, as is the case in the United States Congress [122].

7.3   Policy areas

One can categorize votes based on their topic or policy area. We follow [73] and use the categorization of policy areas from VoteWatch Europe [121]. We give an overview of the number of votes for each policy area in Table 7.2.

7.4   Political groups

Most MEPs are part of a political group in the European Parliament. In the 7th Assembly of the European Parliament, there were seven official groups (see Table 7.3), and there were some MEPs (called "*Non-Inscrits*") who did not belong to any official group. We list the political groups ([20]) from most left-wing to most right-wing (except for the Non-Inscrits (NI)):

Country	MEPs in 2009	MEPs in 2014
Austria	17	19
Belgium	22	22
Bulgaria	17	18
Croatia	0	12
Cyprus	6	6
Czech Republic	22	22
Denmark	13	13
Estonia	6	6
Finland	13	13
France	72	74
Germany	99	99
Greece	22	22
Hungary	22	22
Ireland	12	12
Italy	72	73
Latvia	8	9
Lithuania	12	12
Luxembourg	6	6
Malta	5	6
Netherlands	25	26
Poland	50	51
Portugal	22	22
Romania	33	33
Slovakia	13	13
Slovenia	7	8
Spain	50	54
Sweden	18	20
United Kingdom	72	73
Total	**736**	**766**

Table 7.1: Members of the 7th Assembly of the European Parliament (2009–2014).

Policy	Number of bills
Agriculture	53
Budget	179
Budgetary control	113
Civil liberties, justice, and home affairs	99
Constitutional and inter-institutional affairs	40
Culture and education	19
Development	29
Economic and monetary affairs	128
Employment and social affairs	44
Environment and public health	100
Fisheries	53
Gender equality	28
Industry, research, and energy	51
Internal market and consumer protection	39
Internal regulations of the EP	7
International trade	106
Legal affairs	67
Petitions	5
Regional development	35
Transport and tourism	40
Total	**1235**

Table 7.2: Policy areas (categorization from [121]) and number of bills for every policy area during the 7th Assembly of the EP (2009–2014).

Group	MEPs in 2009	MEPs in 2014
GUE-NGL	35	35
S&D	184	195
Greens-EFA	55	58
ALDE	84	85
EPP	265	274
ECR	54	56
EFD	32	33
NI	27	30
Total	**736**	**766**

Table 7.3: Political groups, ordered from most left-wing (top) to most right-wing (bottom) except for the Non-Inscrits (NI), with their number of MEPs in the 7th Assembly of the EP (2009–2014).

- The *European United Left–Nordic Green Left (GUE-NGL)* consists of radical left-wing, socialist, and communist parties.

- The *Socialists & Democrats (S&D)* consists of social-democratic, socialist and labour parties.

- The *Group of the Greens and European Free Alliance (Greens-EFA)* consists of green and regionalist/separatist parties.

- The *Alliance of Liberals and Democrats for Europe (ALDE)* consists of liberal parties.

- The *European People's Party (EPP)* consists of Christian-democratic, conservative, and centre-right parties.

- The *European Conservatives and Reformists (ECR)* consists of conservative and Eurosceptic parties.

- The *Europe of Freedom and Democracy (EFD)* consists of extreme right-wing and Eurosceptic parties.

- The *Non-Inscrits (NI)* consists of all non-aligned MEPs (the NI is listed last but it is not necessarily the most right-wing party).

There are several reasons why MEPs join political groups. First, the official Parliament Rules of Procedure require the formation of transnational groups, so these groups are awarded resources and procedural advantages [6]. Second, MEPs who share the same ideology on many issues have an incentive to form a group, e.g. so that they can divide tasks [43]. By having a division-of-labour contract with the other MEPs from one's political group, a politician has a larger chance of achieving policy goals at a lower cost, as he/she can just follow the voting instructions of his/her own group and spend most time on one specific policy area [43]. Finally, political groups are also beneficial for the electorate, as voting cohesion among MEPs increases the stability and predictability of voting outcomes in parliament [45]. Nevertheless, European groups have less power than national groups to force members to vote in a particular way, partly because national parties are the ones who select candidates for the European Parliament [45].

7.5   Voting cohesion in the European Parliament

A measure of *voting cohesion* within a group indicates the extent to which legislators from the same party cast the same vote [20]. There are different ways to define and measure voting cohesion [6, 13, 37, 43–46, 66, 102, 103, 105]. Hix et al. (2005) concluded that voting-cohesion levels of European political groups have been rising since the 1990s, as the powers of the EP have increased and the political parties have become more important [43]. However, McElroy and Benoit (2007) noted that some groups (such as the EPP) contain national parties whose positions are out of line with other party group members. Therefore, they warned against methods that automatically assume that national party positions can be inferred from European group positions. This is one of the reasons why a network-based approach is useful; it does not make any *a priori* assumptions about membership of political groups [71].

Hix et al. (2009) applied spatial methods (see Section 3.2) to the European Parliament [45]. They concluded that MEPs voted more along political-party lines than along national lines for two reasons. First, MEPs follow voting instructions from the leader of their political group, because they expect him/her to have the same (or at least a very similar) position on complex legislative issues as they would have if they had time and resources to work out their own positions. Second, MEPs choose to vote with their group to avoid risking future benefits by voting against their European group, especially when there is a clear majority.

A report from VoteWatch (2011) [122] calculated that voting cohesion in the 7th Assembly of the EP lies between 0.8 and 0.95 (where a minimum of 0 indicates that a group is split down the middle in every vote and a maximum of 1 indicates full agreement). Moreover, this report stated that MEPs primarily vote along transnational party lines rather than along national lines [122]. Cherepnalkoski et al. (2016) observed from their results that the Greens-EFA, S&D, and EPP are the most cohesive groups in the EP. The groups with the lowest cohesion are the NI (non-aligned MEPs) and the EFD [20]. They also calculated that the level of voting cohesion is very stable across different policy areas. (They categorized the votes into 9 policy areas [20] instead of our categorization into 20 areas [73, 121]).

7.6   Coalitions in the European Parliament

When MEPs of two or more political groups often agree in their voting, we say that they form a *coalition* [20]. In 2014, a simple majority in the EP corresponded to 384 votes. A coalition between the largest parties (EPP and S&D) could ensure 469 votes, a left-wing coalition (S&D, Greens-EFA, and GUE-NGL) ensured 288 votes, and a right-wing coalition (EPP, ECR, and EFD) ensured 363 votes. Mendonca et al. (2015) noted that the EP is usually split between a coalition including the S&P and a coalition including the EPP [73].

Cherepnalkoski et al. (2016) examined coalition formation in the EP between 2014–2016 and concluded that the degree of co-voting corresponds closely to the seating order in the EP from left (GUE-NGL) to right (EFD) [20]. They also suggested that the two largest parties (EPP and S&D) find it easier to negotiate deals with smaller counterparts than with each other, as was also observed by Hix et al. (2007) [44]. Moreover, Cherepnalkoski et al. (2016) calculated that the strongest "cooperation" in terms of co-voting can be observed between one of the major parties (S&D, EPP) and ALDE [20]. Lastly, they concluded that there are four "blocs of coalitions" in the EP based on co-voting patterns:

- "Left-wing" bloc (GUE-NGL, S&D, Greens-EFA);

- "Centre" bloc (Greens-EFA, ALDE, EPP);

- "Centre-right" bloc (ALDE, EPP, ECR);

- "Far-right" bloc (ECR, EFD, NI).

7.7 Previous network analyses of the European Parliament

7.7.1 Mendonca et al. (2015)

Methods

Mendonca et al. (2015) analyzed the 7th Assembly of the European Parliament using community-detection methods [73]. They first constructed similarity matrices treating abstention either as a full agreement or as half an agreement (see Tables 4.3 and 4.5). These two options give results that are very similar to each other. They then ignored edges to create a network that has positive edges only. That is, they considered *positive networks* in which only the positive edges are kept, and they considered *complementary negative networks (CNNs)*. To construct a CNN, one starts with a complete unweighted network, then removes the edges between pairs of nodes that have a negative edge in the original network, and then takes the complement (i.e. one makes every non-edge an edge and vice versa). Lastly, they applied several community-detection algorithms (including Infomap [107] and a method based on edge betweenness [89]) to both types of networks.

Results

Mendonca et al. (2015) showed that a majority of MEPs tend to vote similarly most of the time [73]. Moreover, the number of communities is relatively low (2–3). They noted that the EP is split in two major political coalitions (one including the EPP and one including the S&D), with some alliances of smaller parties [73]. The results of their two approaches (i.e. considering positive networks and considering complementary negative networks) differ a lot, which confirms that negative edges play a significant role.

What is different in our approach?

Although the network approach in Mendonca et al. (2015) is related to our approach, there are substantive differences in methodology. First, we use the notion of modularity as measure of polarization in the EP (see Section 6.7), and Mendonca et al. do not consider this. Second, we use not only a different heuristic (the Louvain method) for optimizing modularity (Mendonca et al. used a heuristic based on edge betweenness [89]), but we also try to interpret the obtained communities (Mendonca et al. only mentioned the *number* of communities). Finally, instead of throwing away information (by creating a positive network or a CNN), we keep all information by mapping the weights from $[-1, 1]$ to $[0, 1]$ (see Section 8.2).

7.7.2 Cherepnalkoski and Mozetič (2016)

Methods

Cherepnalkoski and Mozetič (2016) analysed Twitter behavior of MEPs between 2014–2015 (i.e. in the 8th Assembly) [21]. They detected communities in the EP with the Louvain method (see Section 6.5.1) and examined how well the communities correspond to the political parties. They also evaluated the correspondence between the communities and the MEPs' countries of origin.

Results

Cherepnalkoski and Mozetič (2016) concluded that most of the communities were dominated by a single political group. They also reported that most MEPs of each political group are contained in a single community (except for MEPs from S&D and EFD who were divided into two communities). They also showed that the detected communities are not organized based on country memberships [21].

What is different in our approach?

Our research uses a different type of data (roll-call voting data instead of Twitter data), and we employ different methods from Cherepnalkoski and Mozetič (2016). For instance, we construct a similarity matrix based on RCVs (see Chapter 4) and use modularity as a measure of polarization (see Section 6.7).

7.7.3   Cherepnalkoski et al. (2016)

Methods

Cherepnalkoski et al. (2016) studied the extent to which MEPs of the same party cast the same vote (voting cohesion), and the extent to which MEPs of different parties cast the same vote (voting coalitions) in the European Parliament during 2014–2016 [20]. They analyzed roll-call votes and Twitter data using exponential random graph models (ERGMs) (see Section 2.5) and Krippendorff's alpha reliability [59] (which compares observed agreement between voters to the agreement expected by chance). Cherepnalkoski et al. explored whether coalitions are formed in the same way for different policy areas, what the relationship is between co-voting and retweeting patterns, and the cohesiveness of each political group in the EP.

Results

Cherepnalkoski et al. (2016) observed a positive correlation between retweeting and co-voting for all policy areas except "economic and monetary system" [20]. Moreover, they calculated that the level of voting cohesion is very stable across different policy areas (see Section 7.5) and that there are four types of coalitions in the EP based on co-voting patterns (see Section 7.6).

What is different in our approach?

Even though we use the same type of data (roll-call votes) as Cherepnalkoski et al. (2016), our methodology is different. We use a community-detection method (modularity optimization) to find "blocs" in the EP, whereas they used ERGMs to investigate the network-formation process and Krippendorff's alpha reliability to measure voting agreement. We also use the notion of modularity to consider polarization in the entire European Parliament, whereas they examine voting cohesion within each single political party. Lastly, the categorization of policy areas that we use (taken from [121]) has 20 categories instead of 9. (Most of the 9 categories in Cherepnalkoski et al. seem to contain our categories as subsets).

Chapter 8

Research methods and results

8.1   Data

Data extraction

The data that we use is from *VoteWatch Europe* [121] and has been made available by Mendonca et al. (2015) [73]. VoteWatch Europe is an international, independent, and non-partisan organization that provides access to all roll-call votes (RCVs) in the European Parliament. In the VoteWatch Europe data, each MEP is described through his/her name, country, political group, and the way that he/she voted for each roll-call vote in the EP [73].

Mendonca et al. (2015) developed a tool in R called *NetVotes* to process the VoteWatch data [73]. They used the igraph R package to perform community detection. Their source code, which we used for constructing the similarity matrix, is available on GitHub [80].

Description

We consider 1235 RCVs (see Table 7.2) in the EP (2009–2014) divided over 20 policy areas[1] as in [73]. Between 2009–2014, some MEPs left the European Parliament and were replaced by others, which gives a total of 840 MEPs in our data set. Each MEP can vote "*for*", "*against*", "*abstain*", or "*be absent*". (Absence can either be "not present", "present but did not vote", or "not present for a justified reason", but we do not differentiate between those options). Mendonca et al. (2015) noted that a majority of MEPs tend to vote similarly and that a certain number of MEPs frequently abstained [73]. The RCVs are categorized based on *policy area* (see Table 7.2) by VoteWatch Europe [121], and we construct a similarity matrix for each policy area (20 matrices in total). We consider RCVs over the whole 5-year term.

Selection bias

About 1/3 of the votes in the European Parliament are taken by roll call, so we have not considered 2/3 of all votes in the EP. McElroy and Benoit (2007) noted that there may be a selection bias, as roll-call votes may be called selectively and strategically for some specific issues [71]. Consequently, we should be careful when interpreting our results, as we do not know how MEPs voted for a majority of the bills.

Number of bills

Some policy areas have only a small number of bills. (For example, there are 5 bills related to "petitions" and 7 bills related to "internal EP regulations"; see Table 7.2.) Consequently,

---

[1]There is one policy area ("foreign and security policy") for which we could not obtain any results due to an error in *NetVotes* [80].

we should be careful when analyzing those policy areas, as we have less certainty that the observed patterns are representative for the policy area in general.

Missing data

Towards the end of our dissertation process (in August 2016), Professor Vincent Labatut informed us that their VoteWatch data [80] is incomplete. A part of the RCVs from 2014 is missing, and some votes concerning amendments (2009–2014) are not included. They are currently rewriting the code, but they have not released the updated version yet. This is another reason why we need to be careful about interpreting the results.

8.2 Methods

Normalization of the similarity matrix

As discussed in Chapter 4, we construct a similarity matrix $\mathbf{S}$ such that all entries $S_{ij} \in [-1, 1]$. We construct a nonnegative adjacency matrix $\mathbf{A}$ by letting

$$A_{ij} = S_{ij} \times 0.5 + 0.5 \,, \tag{8.1}$$

such that $A_{ij} \in [0, 1]$ for all i, j. We thus do not throw away information. MEPs who disagree a lot with each other have an edge weight close to 0 between them, and MEPs who agree a lot have a weight close to 1. A pair of MEPs who never vote for the same bill receive a weight of 0.5 (so we have to keep in mind that we do not treat edges with weight 0.5 as "half of an agreement").

Measuring modularity

We use a function called MODULARITY METRIC [16] (available for MATLAB at [77]) to calculate party modularity Q_{party} and country modularity Q_{country} (which are the values of Q for partitions based on, respectively, party affiliation and country of origin). We also use this function to calculate the modularity of partitions based on "left-wing" versus "right-wing" parties, and on "extreme" versus "centrist" parties. We obtain the values for Q_{max} directly from the Louvain algorithm.

Louvain method

We detect communities by optimizing modularity using a variant of the Louvain method (see Section 6.5.1). We use the function GENLOUVAIN (version 2.0) to apply this method [51]. The nodes are considered in a random order (uniformly at random), so multiple runs usually have (slightly) different results [14, 60]. For every network, we run the algorithm R times (with $R = 100$ for most results), which gives an ensemble of R partitions and R corresponding modularity values. (As we discussed in Algorithm 1, the Louvain method chooses a move that maximally increases modularity, so we do not enable "randmove" in GENLOUVAIN).

Consensus partition

We construct a single *consensus partition* that is most "similar" to the other R partitions (obtained from different runs). We use a function called CONSENSUS SIMILARITY [28] from the Network Community Toolbox [81] to obtain a consensus partition. We measure similarity using the z-score of the Rand coefficient [115]:

$$z_{X,Y} = \frac{a_{00}}{\sigma_{a_{00}}} - \frac{(a_{11} + a_{10})(a_{11} + a_{01})}{\binom{n}{2}}, \tag{8.2}$$

where $\sigma_{a_{00}}$ is the standard deviation of a_{00}, and the other symbols are as in equation (5.5). We calculate the z-score using the function ZRAND from [115] (available at [81]).

Resolution parameter γ

We try to find a range of γ for which the obtained partitions are very similar to each other, such that we can reasonably conclude that the detected communities are robust [68]. We use normalized variation of information (NVI) (equation (5.9)) to measure similarity between partitions, using a function called PARTITION DISTANCE [72] (available at [17]).

Robustness

We use three methods to measure the robustness of our partitions [60]. For the first method, we perturb a network by randomly changing the weight of each edge:

$$A_{ij}^* = A_{ij} + \sigma \times u_{ij}, \tag{8.3}$$

where $\sigma \in [0, 1]$ is a fixed constant and u_{ij} is a random variable uniformly distributed between -1 and 1. For each value of γ, we run the Louvain method once (using a deterministic alphabetical node ordering) for M different realizations of a perturbed network. In this approach, the robustness of partitions is given by

$$\langle \mathrm{NVI} \rangle_{\mathrm{random}}(\gamma) = \frac{2}{M(M-1)} \sum_{m=1}^{M} \sum_{m'=m+1}^{M} \mathrm{NVI}(X_m(\gamma), X_{m'}(\gamma)), \tag{8.4}$$

where $X_m(\gamma)$ is an optimal partition of the m^{th} realization of the perturbed network for resolution paramter γ [60].

For the second method, we fix γ and compare partitions of different runs of the algorithm. In this case, robustness at resolution parameter γ is [60]

$$\langle \mathrm{NVI} \rangle_{\mathrm{algo}}(\gamma) = \frac{2}{R(R-1)} \sum_{i=1}^{R} \sum_{j=i+1}^{R} \mathrm{NVI}(X_i(\gamma), X_j(\gamma)), \tag{8.5}$$

where $X_i(\gamma)$ is the optimal partition in the i^{th} run of the Louvain method.

For the third method, we measure robustness by tuning the resolution parameter γ. In this case, we construct a consensus partition for each γ and compare it to partitions with neighbouring values of γ (i.e. $\gamma + \Delta\gamma$, where $\Delta\gamma = 0.01$). The robustness is given by

$$\langle \mathrm{NVI} \rangle_{\gamma}(\gamma) = \frac{1}{\Delta} \sum_{\tau=1}^{\Delta} \mathrm{NVI}(X(\gamma), X(\gamma + \tau)), \tag{8.6}$$

where $X(\gamma)$ is a consensus partition at γ. When the value of $\langle \text{NVI} \rangle_\gamma(\gamma)$ between two partitions with neighbouring values of γ is small, we know that these partitions are relatively similar, which suggests that we have found robust communities. Nevertheless, Macon et al. (2012) [68] noted that the identification of "robust" communities remains arbitrary to some extent (e.g. the value of $\langle \text{NVI} \rangle_\gamma(\gamma)$ at which we call a partition "robust" is arbitrary).

Visualization

To visualize networks, we use the functions FRKK and GRAPHPLOT2D which are both available at [119]. Function FRKK uses the Fruchterman–Reingold algorithm [36] to find an optimal node placement for a given network with respect to communities. Then, it uses the Kamada–Kawai algorithm [52] to place the nodes within each community. (We set $epsilon = 0.01$, $seed = 1993$, $factor1 = 2$, and $factor2 = 4$).

We use the function GRAPHPLOT2D (with $alpha = 2$) to display our networks using coordinates that are given by FRKK. To make our figures more clear, we only display edges with a weight $w_{ij} \geq 0.85$.

8.3   Results

Polarization for different policy areas

We first compare the (mean) maximum modularity values $Q_{\max}$ (over 100 runs) between all policy areas for a fixed resolution parameter ($\gamma = 1$). In Table 8.1 and Figure 8.1, we show that $Q_{\max}$ ranges between 0.016 ("legal affairs") and 0.061 ("internal regulations of the European Parliament"). Other policy areas with relatively high values of modularity (which suggest a relatively high level of polarization) are "budgetary control", "gender equality", and "petitions".

We also show the country modularity Q_{country} and party modularity Q_{party} in Table 8.1 and Figure 8.1. Policy areas with relatively high $Q_{\text{party}}/Q_{\max}$ ratios are "civil liberties", "employment and social affairs", and "environment and public health". Party affiliation seems to contribute relatively little to the total polarization for "development" and "internal regulations of the EP". The policy area with the highest $Q_{\text{country}}/Q_{\max}$ ratio is "agriculture". This is in line with the observation by VoteWatch Europe [121] that there seem to be some national delegations (such as the French and Scandinavian ones) who vote together instead of voting along political party lines [122].

Lastly, we consider two more partitions. First, we construct a "left wing" community (consisting of GUE-NGL, S&D, Greens-EFA, and ALDE), and a "right wing" community (consisting of EPP, ECR, EFD, and NI). Figure 8.2 shows that this division leads to a relatively high modularity value in a network based on "gender equality" votes.

Second, we construct an "extreme" community (consisting of GUE-NGL, Greens-EFA, EFD, and NI), and a "centrist" community (consisting of S&D, ALDE, EPP, and ECR). Figure 8.2 shows that this division leads to a relatively high modularity value in a network based on "industry" votes.
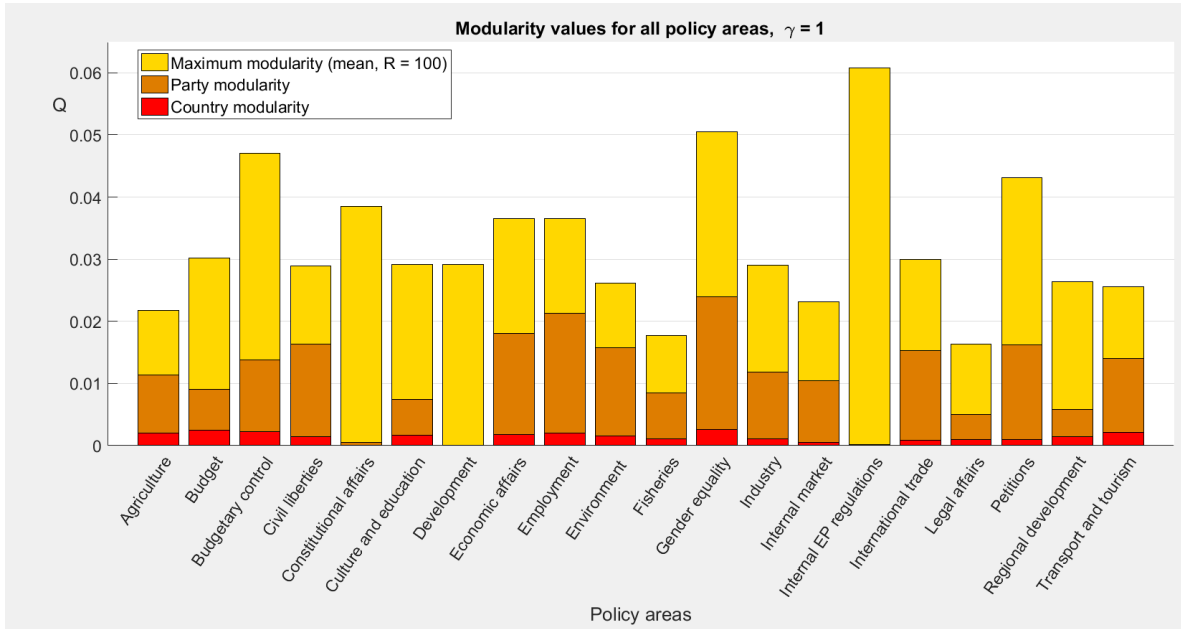
Figure 8.1: Maximum modularity, party modularity, and country modularity values for all policy areas, $\gamma = 1$
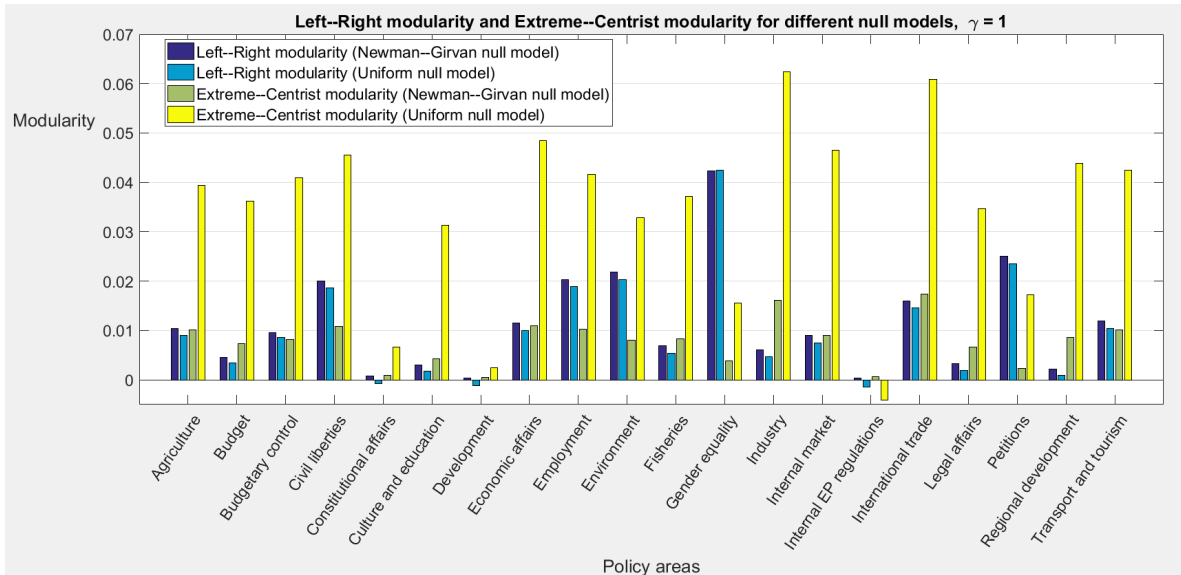

Figure 8.2: Left–Right modularity and Extreme–Centrist modularity for different null models, $\gamma = 1$

Policy	$Q_{country}$	Q_{party}	Q_{max} (mean)	$Q_{country}/Q_{max}$	Q_{party}/Q_{max}
Agriculture	0.0020	0.011	0.022	9.3×10^{-2}	0.52
Budget	0.0024	9.1×10^{-3}	0.030	8.1×10^{-2}	0.30
Budgetary control	0.0023	0.014	0.047	4.9×10^{-2}	0.29
Civil liberties	0.0014	0.016	0.029	4.9×10^{-2}	0.57
Constitutional affairs	-3.4×10^{-5}	4.7×10^{-4}	0.039	-9.0×10^{-4}	1.0×10^{-2}
Culture and education	0.0017	7.4×10^{-3}	0.029	5.7×10^{-2}	0.25
Development	-2.8×10^{-5}	7.3×10^{-5}	0.029	-9.6×10^{-3}	2.5×10^{-3}
Economic affairs	0.0018	0.018	0.037	5.1×10^{-2}	0.49
Employment, social affairs	0.0020	0.021	0.037	5.6×10^{-2}	0.58
Environment, public health	0.0015	0.016	0.026	5.9×10^{-2}	0.60
Fisheries	0.0010	8.5×10^{-3}	0.018	5.9×10^{-2}	0.48
Gender equality	0.0026	0.024	0.051	5.1×10^{-2}	0.48
Industry, research, energy	0.0011	0.012	0.029	3.8×10^{-2}	0.41
Internal market	5.3×10^{-4}	0.010	0.023	2.3×10^{-2}	0.45
Internal EP regulations	7.7×10^{-5}	2.2×10^{-4}	0.061	1.3×10^{-3}	$3.5 * 10^{-3}$
International trade	8.9×10^{-4}	0.015	0.030	3.0×10^{-2}	0.51
Legal affairs	1.0×10^{-3}	5.1×10^{-3}	0.016	6.1×10^{-2}	0.31
Petitions	0.0010	0.016	0.043	2.4×10^{-2}	0.38
Regional development	0.0014	5.8×10^{-3}	0.026	5.4×10^{-2}	0.22
Transport and tourism	0.0021	0.014	0.026	8.2×10^{-2}	0.55
Average	**0.0013**	**0.011**	**0.032**	**4.5×10^{-2}**	**0.37**

Table 8.1: Country modularity Q_{country}, party modularity Q_{party}, (mean) maximum modularity Q_{max}, ratios between modularities for each policy area, with $\gamma = 1$ and $R = 100$ runs of modularity maximization using GenLouvain .
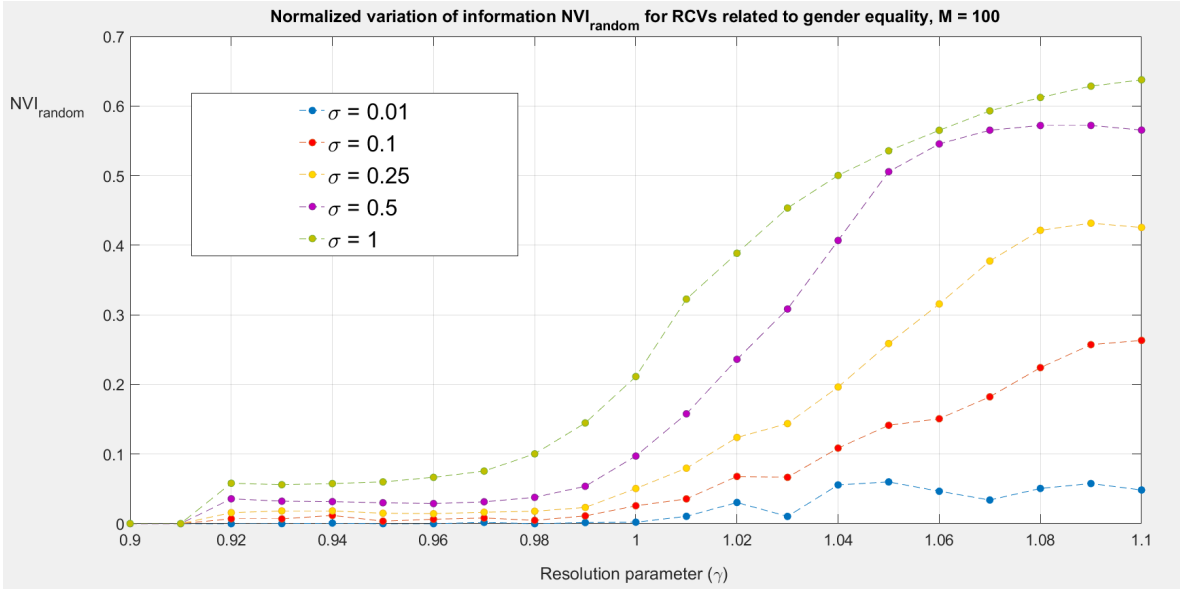
Figure 8.3: Normalized variation of information $\text{NVI}_{\text{random}}$ for RCVs related to gender equality, $M = 100$

8.3.1 Different null models

In this section, we compare the results that we obtain using the Newman–Girvan null model to the results that we obtain using a uniform null model (see Section 6.3). Figure 8.2 shows that the modularity values for those two different null models are close when we consider a left–right partition. When we consider an extreme–centrist partition, the uniform null model gives modularity values that are approximately four times higher than the values we obtain using the Newman–Girvan null model. Both observations (i.e. similarity between modularity values for a left–right partition and a constant factor of 4 for an extreme–centrist partition) seem to hold for all policy areas. Hence, we consider only the Newman–Girvan null model to compare different policy areas.

8.3.2 Different randomizations

Equation (8.3) shows how we perturb a network by randomly changing the weight of each edge (we use $M = 100$ different realizations of a perturbed network). In this section, we compare different values of σ for a RCV network based on gender equality. In Figure 8.3, we observe that higher values of σ lead to higher NVI values. Intuitively, this makes sense, as higher values of σ increase the importance of the random contribution to each edge weight, which leads to more variation between adjacency matrices and more variation between optimal partitions. We use $\sigma = 0.1$ for our next computations, unless stated otherwise.

8.3.3 Communities in the European Parliament

We further investigate five policy areas: the one with the highest $Q_{\max}$ ("internal EP regulations"), the one with the highest $Q_{\text{party}}/Q_{\max}$ ratio ("environment and public health"), the
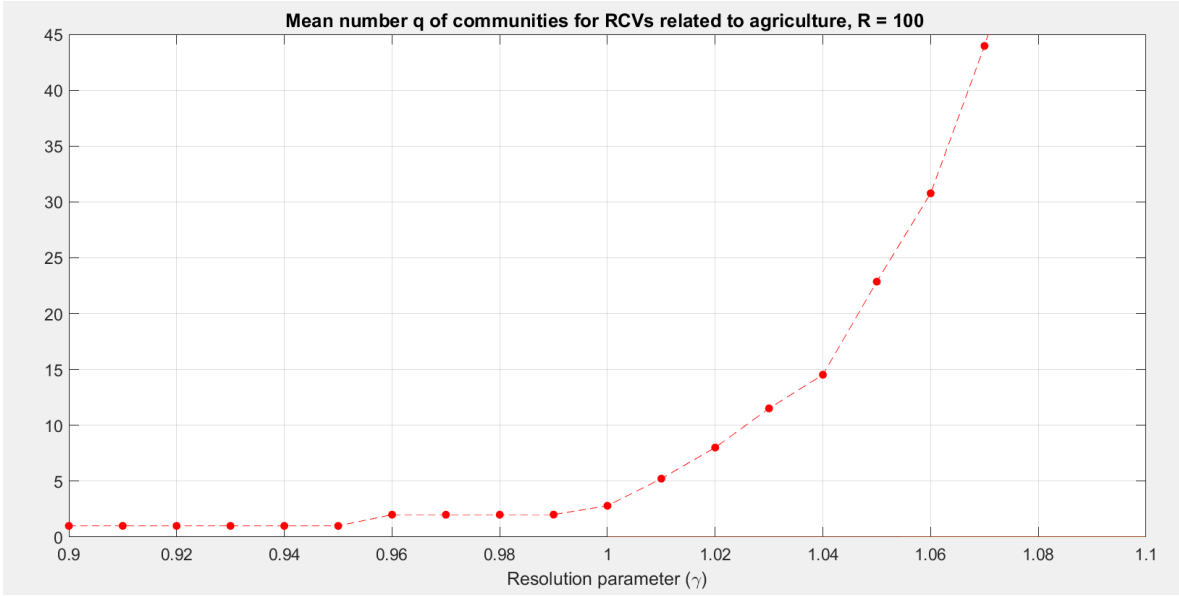
Figure 8.4: Mean number q of communities for RCVs related to agriculture, $R = 100$

one with the highest $Q_{\text{country}}/Q_{\max}$ ratio ("agriculture"), the one with the highest left–right modularity ("gender"), and the one with the highest extreme–centrist modularity ("industry").

Agriculture

Figure 8.4 indicates the mean number of detected communities as a function of γ (over $R = 100$ realizations of modularity maximization with random node ordering). The (mean) number of communities increases quickly for $\gamma > 1$, so we do not expect robust partitions for $\gamma > 1$.

Figure 8.5 shows three types of robustness. The blue line indicates robustness with respect to edge-weight perturbations. A lower value of $\langle\text{NVI}\rangle_{\text{random}}(\gamma)$ indicates that optimal partitions (in $M = 50$ perturbed networks) are more similar to each other. The red line indicates robustness of partitions with respect to multiple runs ($R = 100$) of the Louvain method (for fixed γ). A lower value of $\langle\text{NVI}\rangle_{\text{algo}}(\gamma)$ implies that partitions (from different realizations with random node ordering) are more similar. The yellow line in Figure 8.5 shows the robustness of consensus partitions with respect to neighbouring values of γ (where we use $\Delta\gamma = 0.01$, $R = 100$). A lower value of $\langle\text{NVI}\rangle_{\gamma}(\gamma)$ implies that two partitions (obtained at resolution parameter γ and $\gamma + \Delta\gamma$) are more similar.

Based on Figure 8.5, the most robust partition of the EP network is a division into two communities for $\gamma \in [0.96, 0.99]$ (we do not consider a "division" into a single community). We construct a consensus partition for different values of γ in this "robust" range. We then calculate the fraction of MEPs that are in "community 1" (where "1" is an arbitrary label) for each political party. For instance, if half of the GUE-NGL MEPs are in community 1 and the other half are in community 2, the fraction equals 0.5.

In Figure 8.6, our computations illustrate that for our "agriculture" network, the optimal division is not a strict division between "left wing" and "right wing". The Greens, GUE-NGL,
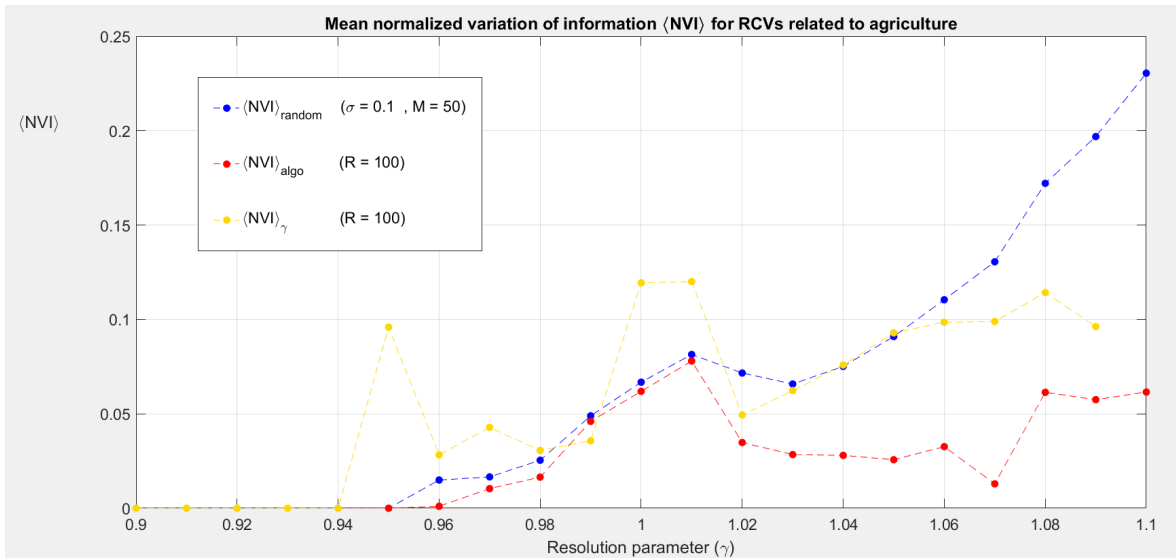
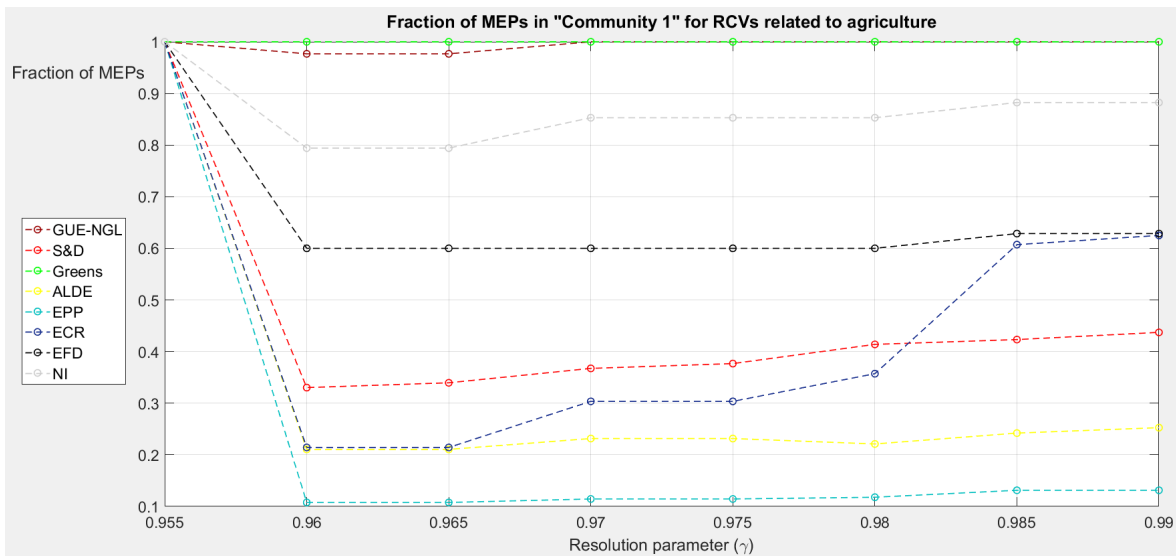Figure 8.5: Mean normalized variation of information NVI for RCVs related to agriculture

Figure 8.6: Fraction of MEPs in "Community 1" for RCVs related to agriculture

and the Non-Inscrits are clearly part of community 1, and the EFD is closer to community 1 as well. This community seems to consists of parties that are further away from the political centre (either to the left or to the right). GUE-NGL consists of radically left-wing MEPs, Greens-EFA consists of green and separist MEPs, the EFD consists of extreme right-wing MEPs, and most of the Non-Inscrits tend to be extreme right-wing as well (e.g. some of them formed a new extreme right-wing group in the 8th Assembly of the EP). The political parties in which most MEPs are part of community 2 (EPP, ALDE, and S&D) are closer to the political centre. We illustrate this partition with a network in Figure 8.7.

Environment

Figure 8.8 shows that the mean number of communities increases quickly for $\gamma > 1.03$, so we do not expect to detect robust communities in that range. Figure 8.9 indicates that we can find a relatively robust partition into two communities for $\gamma \in [0.96, 0.98]$. In figure 8.10, we observe that the partition clearly follows party lines along a left–right dimension. Most MEPs from the left-wing parties are in community 1 (ranging between 75% of the MEPs (ALDE) to 95% of the MEPs (Greens-EFA)), and most MEPs from the right-wing parties are in community 2. Hence, our results suggest that there are "left wing" and "right wing" communities in the EP when MEPs vote upon environmental issues.

Internal EP regulations

In Figure 8.11, we see that the mean number of communities does not increase as rapidly for $\gamma \in [1.00, 1.08]$ as it did for agriculture and environment. The mean NVI values in Figure 8.12 fluctuate more than they did for agriculture and environment. In Figure 8.13, we consider a partition into two communities for $\gamma \in [0.90, 0.98]$. This figure suggests that we cannot attach a "label" to the two communities for the policy area "internal EP regulations". The members of almost all political parties are divided (equally) between the two communities (except for the Non-Inscrits).

Gender equality

Based on Figures 8.14 and 8.15, there are relatively robust partitions for $\gamma \in [0.92, 0.98]$ and for $\gamma \in [1.02, 1.03]$. In Figure 8.17, we zoom in on the partition with two communities and observe that there is a "left wing" community and a "right wing" community. We illustrate this partition in Figure 8.16.

In Figure 8.18, we zoom in on the partition with five communities. A closer look suggests that there are three communities with a substantial number of MEPs. One of them is a "left wing" community (GUE-NGL, Greens-EFA, S&D, and ALDE), one of them is a "right wing" community (ECR, EFD, and NI), and the third community mainly consists of MEPs from the EPP party.

Industry

Figures 8.19 and 8.20 suggest that there are relatively robust partitions for $\gamma \in [0.92, 0.98]$. In Figure 8.21, we observe that there is a "centrist" community (S&D, ALDE, EPP, and ECR) and an "extreme" community (Greens-EFA, GUE-NGL, NI).
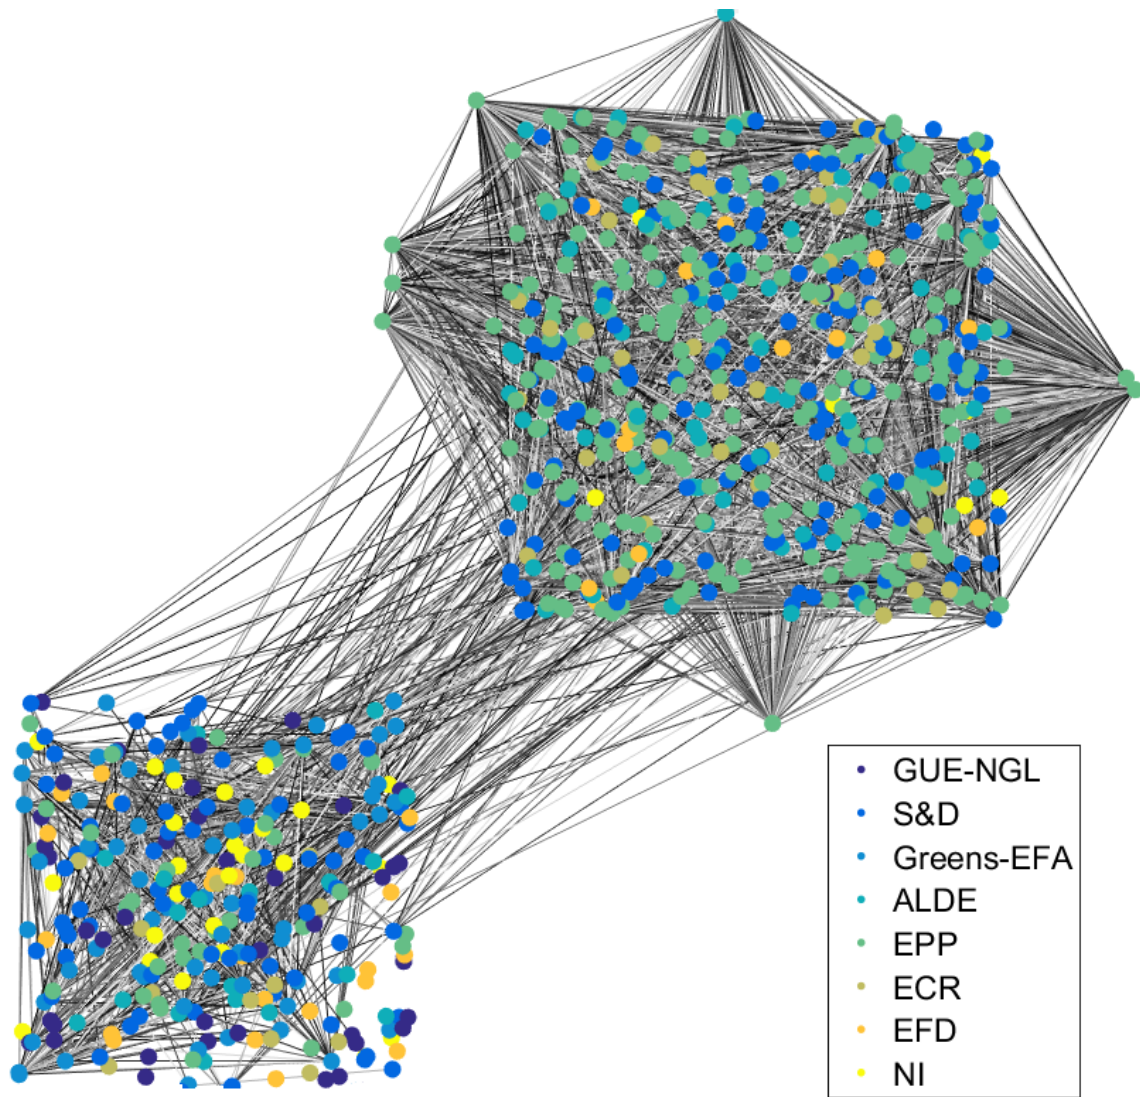
Figure 8.7: Network of MEPs for RCVs related to "agriculture", we only display edges with a weight of 0.85 or higher, $\gamma = 0.96$.
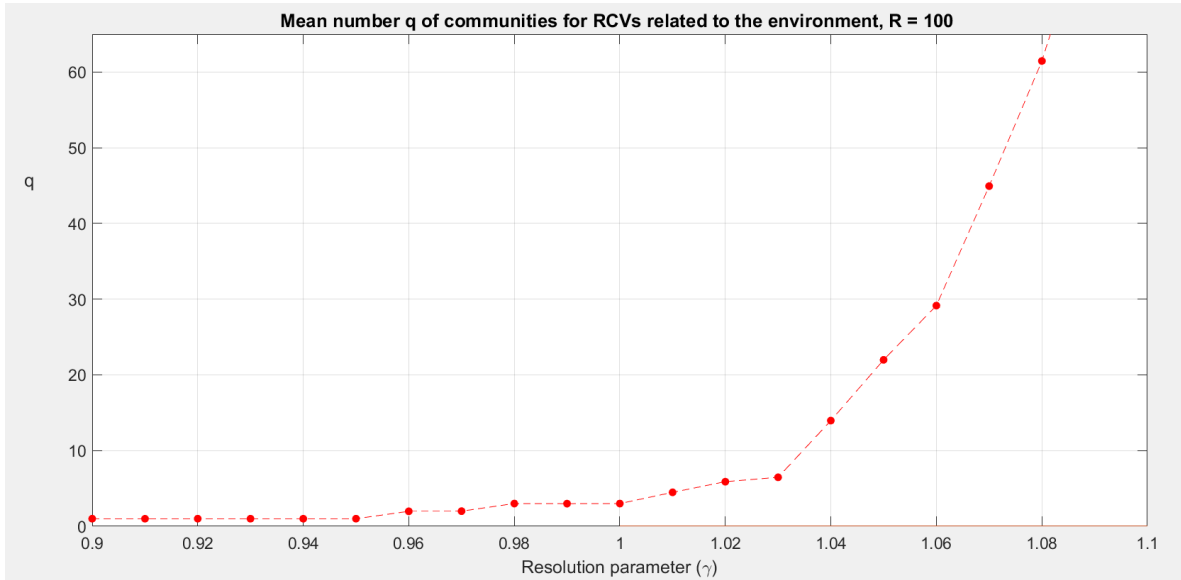
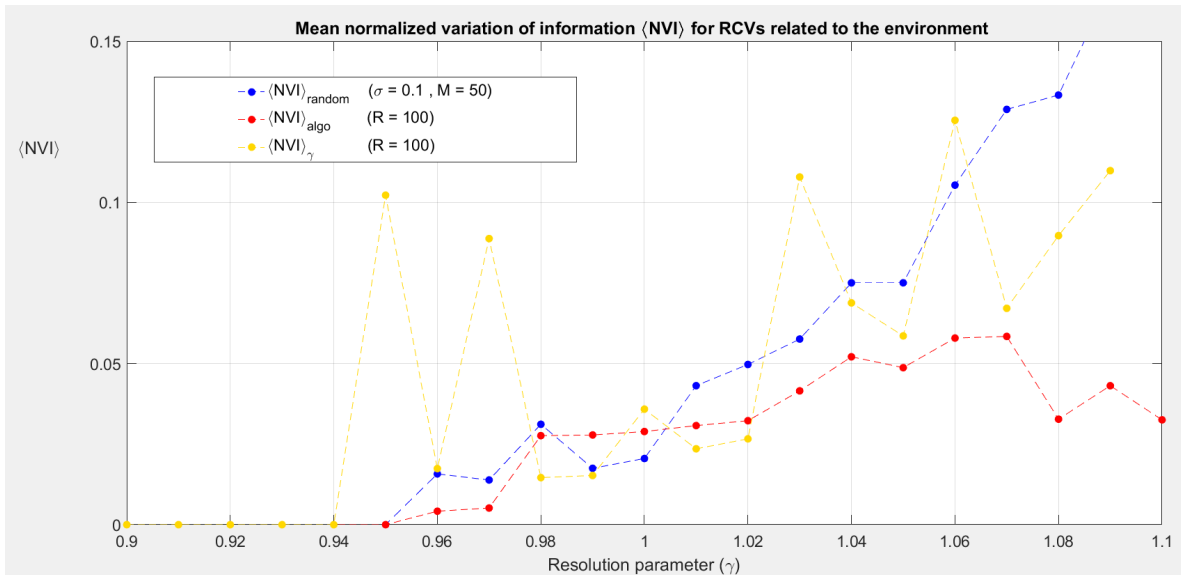Figure 8.8: Mean number q of communities for RCVs related to the environment, $R = 100$

Figure 8.9: Mean normalized variation of information NVI for RCVs related to the environment
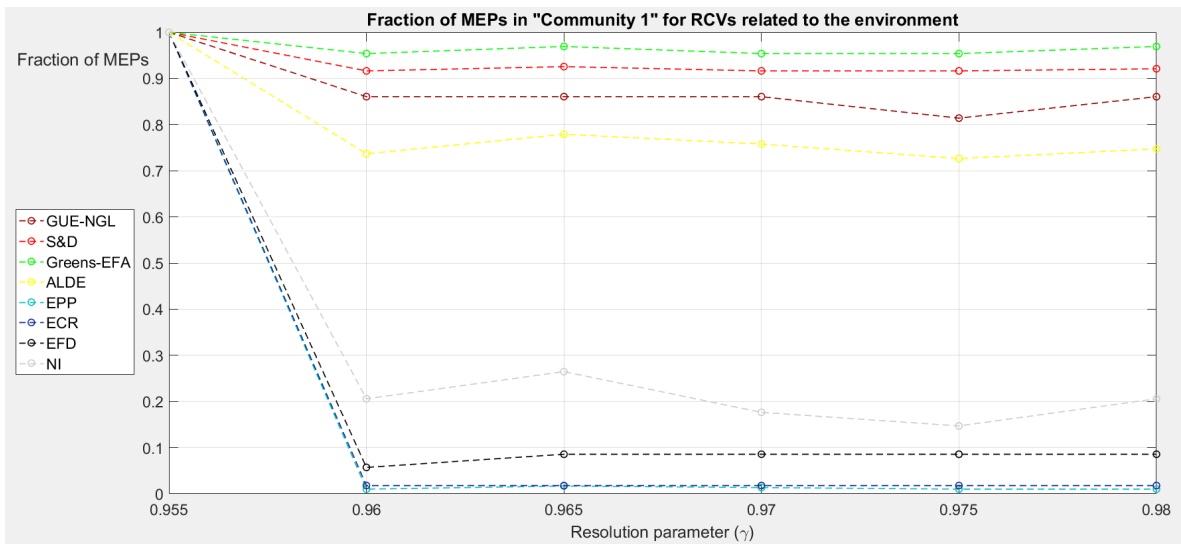
Figure 8.10: Fraction of MEPs in "Community 1" for RCVs related to the environment
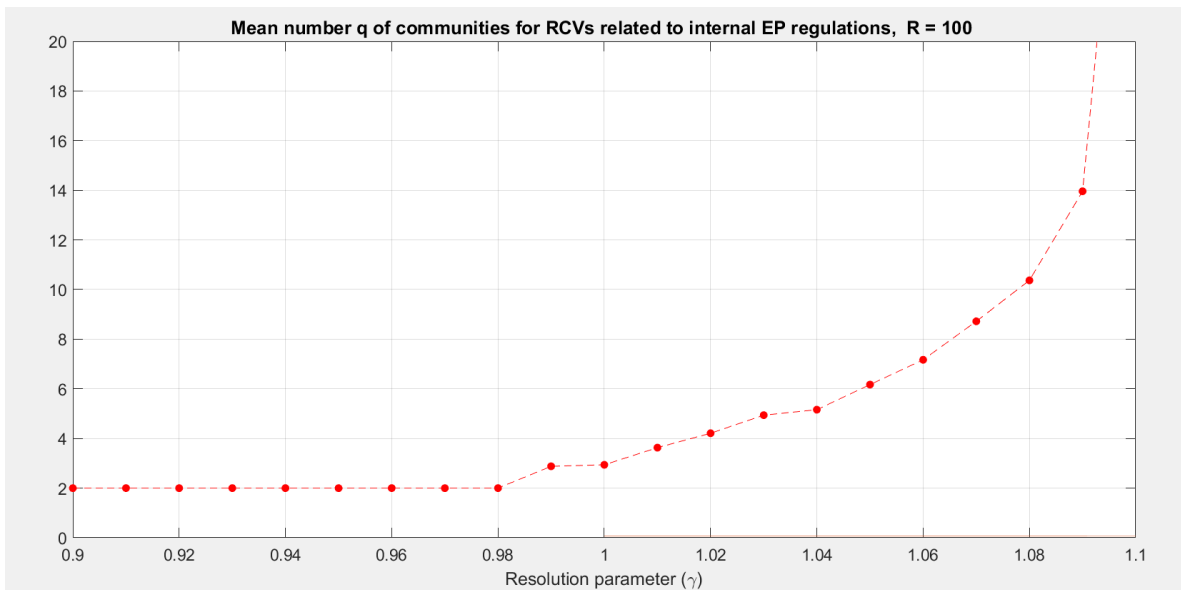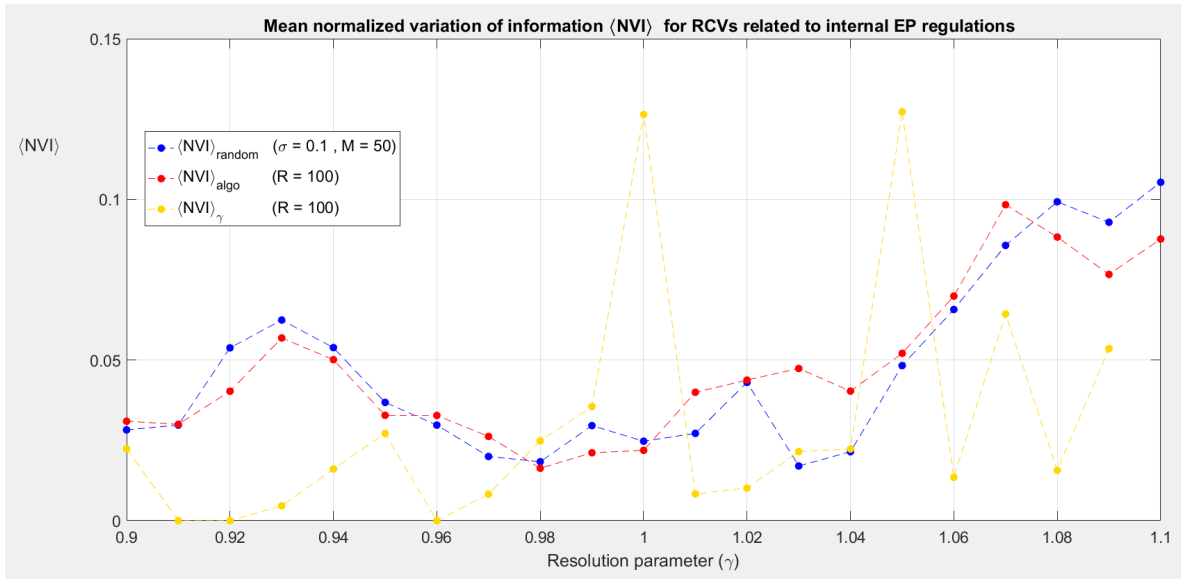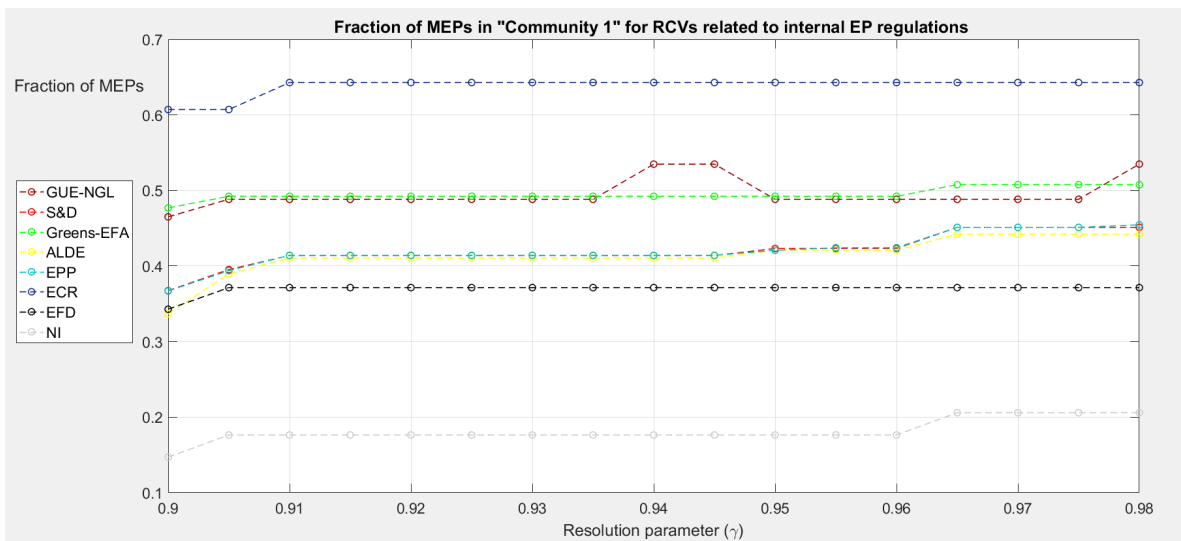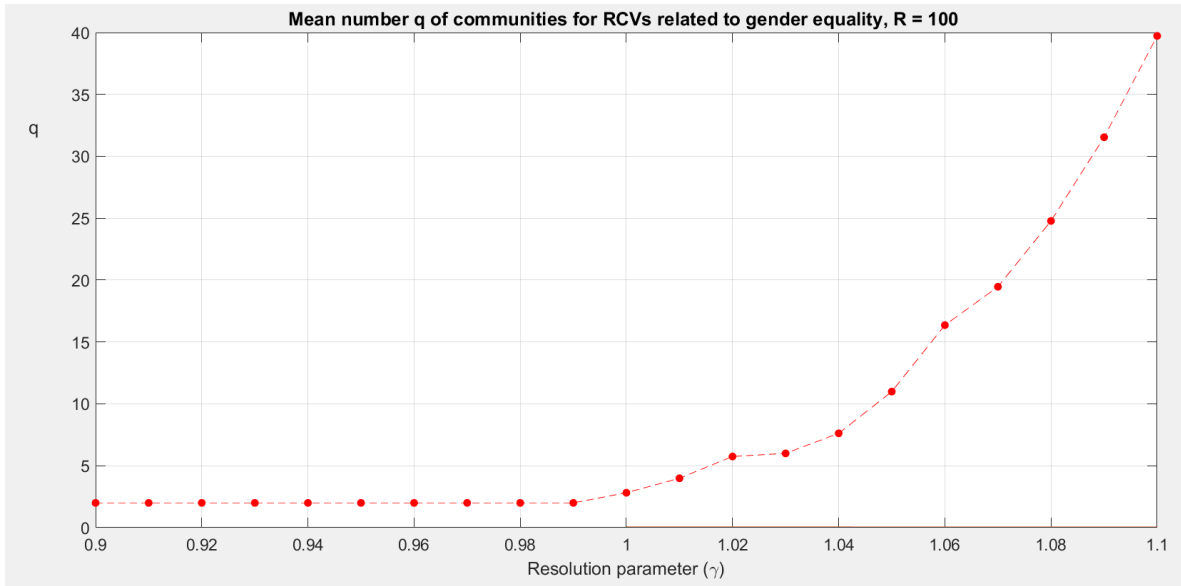

Figure 8.11: Mean number q of communities for RCVs related to internal EP regulations, $R = 100$

Figure 8.12: Mean normalized variation of information NVI for RCVs related to internal EP regulations

Figure 8.13: Fraction of MEPs in "Community 1" for RCVs related to internal EP regulations

Figure 8.14: Mean number q of communities for RCVs related to gender equality, $R = 100$
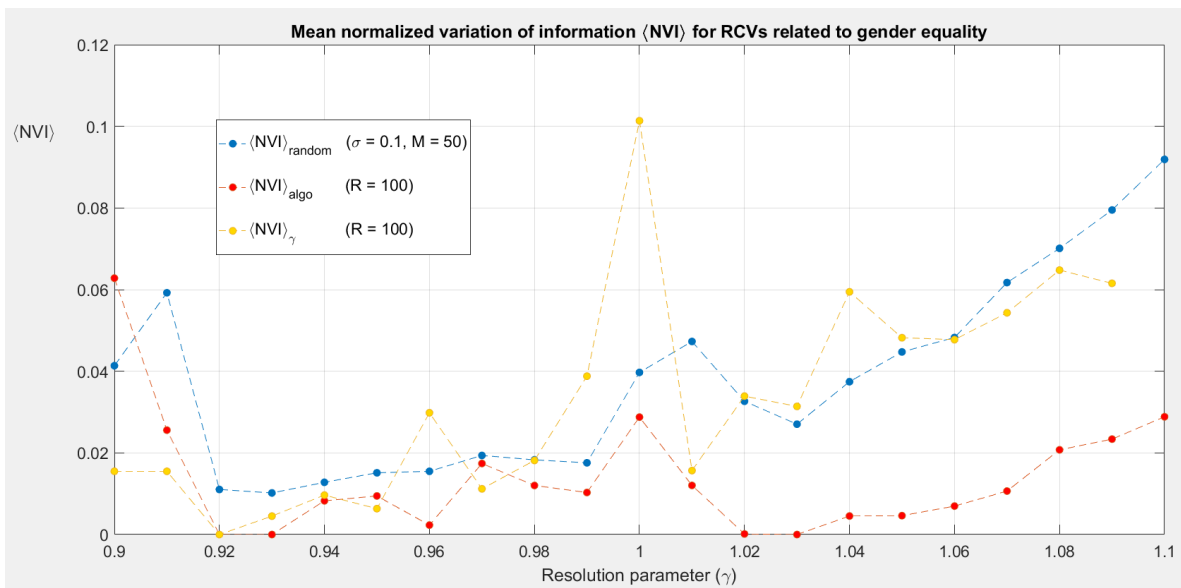

Figure 8.15: Mean normalized variation of information NVI for RCVs related to gender equality
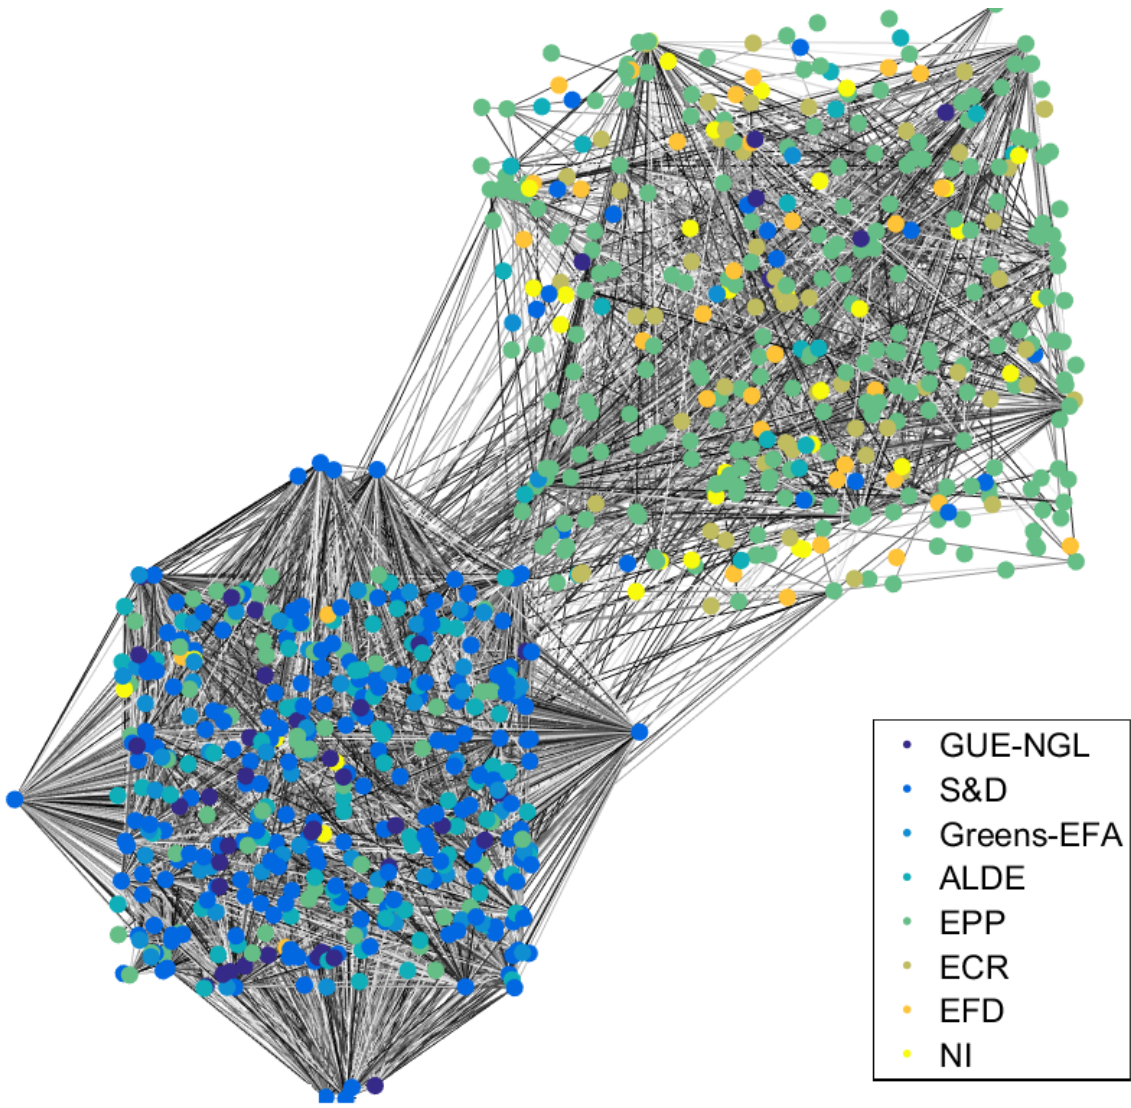
Figure 8.16: Network of MEPs for RCVs related to "gender equality", we only display edges with a weight of 0.85 or higher, $\gamma = 0.93$.
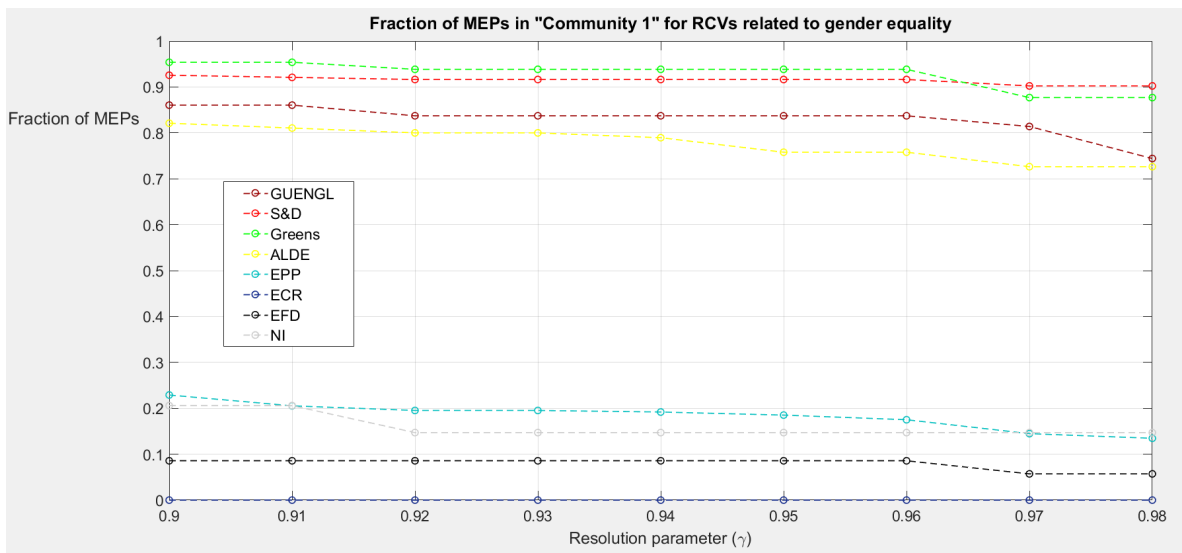
Figure 8.17: Fraction of MEPs in "Community 1" for RCVs related to gender equality
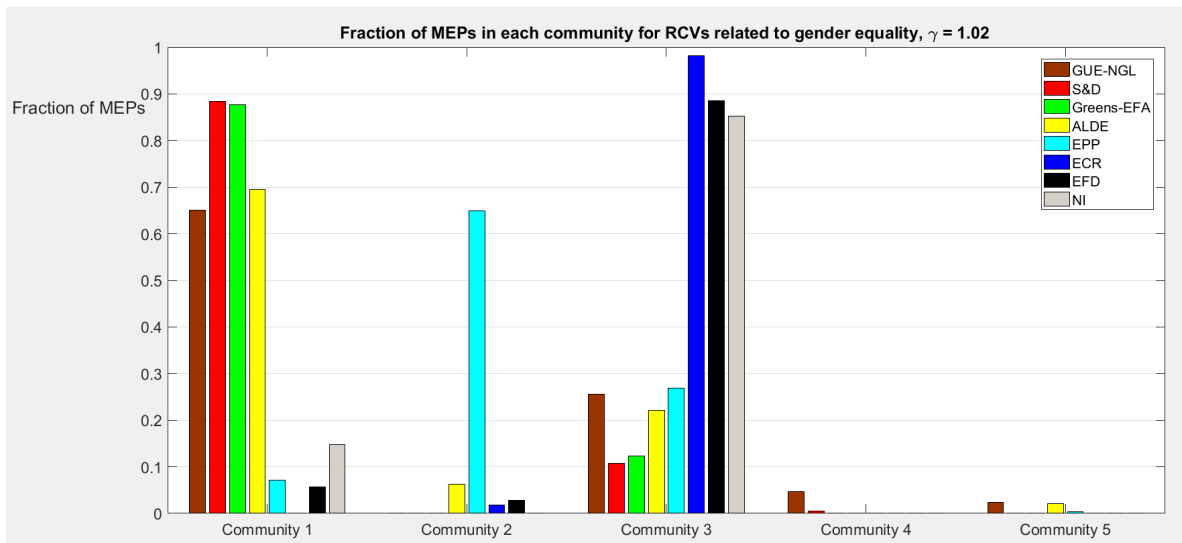

Figure 8.18: Fraction of MEPs in each community for RCVs related to gender equality, $\gamma = 1.02$
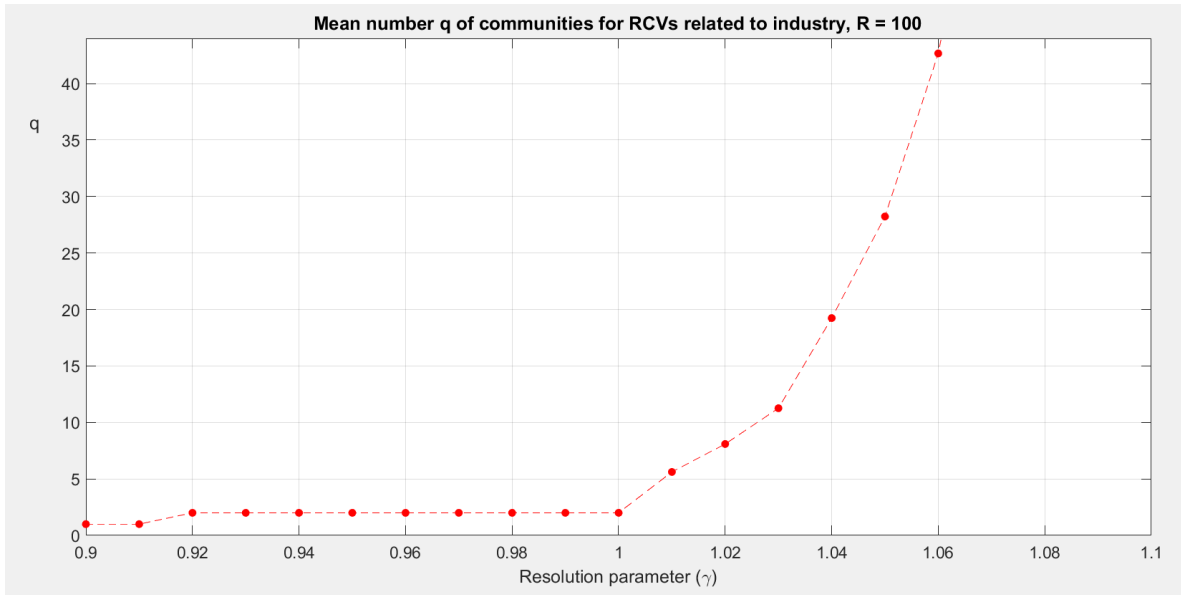
Figure 8.19: Mean number q of communities for RCVs related to industry, $R = 100$
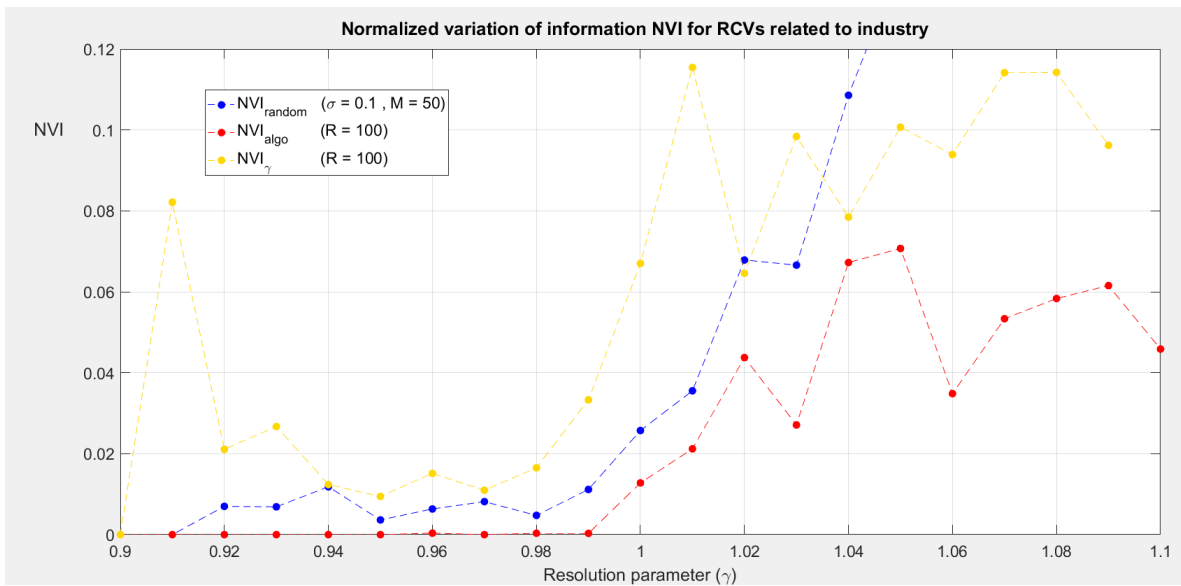

Figure 8.20: Mean normalized variation of information NVI for RCVs related to industry
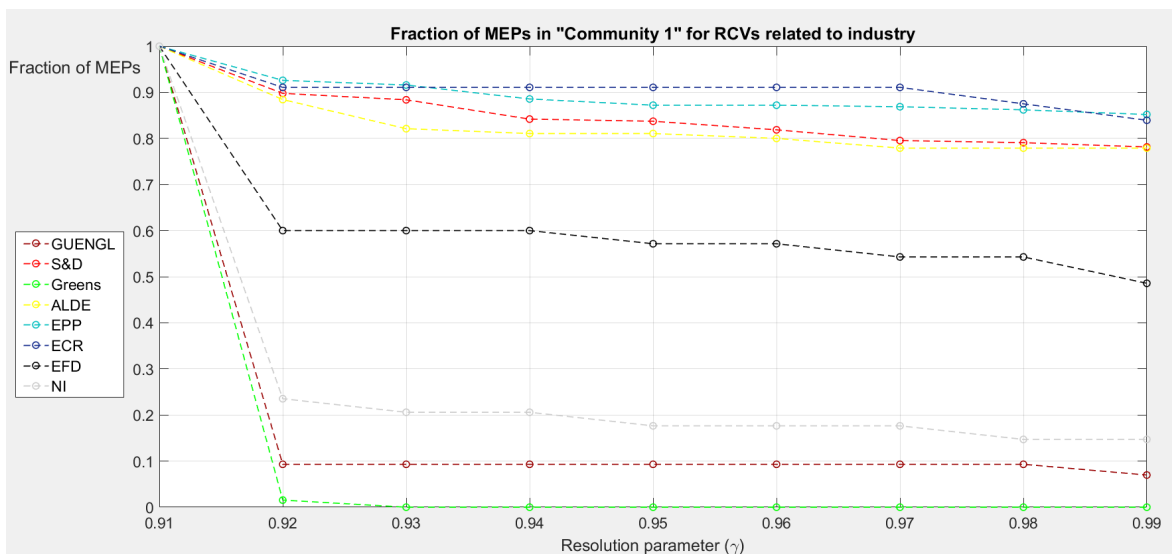
Figure 8.21: Fraction of MEPs in "Community 1" for RCVs related to industry

The EFD party is split between those two communities. We find that the "centrist" part of the EFD (20 MEPs) consists of 8 Italian MEPs and 12 MEPs from 9 other countries. The "extreme" part of the EFD (15 MEPs) consists of 10 British MEPs and 5 MEPs from 4 other countries. These British MEPs are all members of the UK Independence Party (UKIP).

8.3.4 Summary and comparison of results

Based on our modularity calculations, some policy areas seem to be more polarized than others. We find that the EP is relatively polarized for bills related to "budgetary control", "gender equality", and "internal EP regulations". Policy areas that are relatively less polarized include "fisheries", "internal market", and "legal affairs". Party modularity is relatively high for bills related to "civil liberties", "employment", and "environment".

The most robust network divisions split the European Parliament into two communities. Whether we can assign a clear label to each community depends on the policy area. When we focus on environment-related bills and gender-related bills, there seems to be a "left-wing" community and a "right-wing" community, which is in line with the result in [73] that there is a bloc around the Socialists & Democrats (S&D) and a bloc around the European People's Party (EPP). However, for agriculture-related bills and industry-related bills, it seems as if there is an "extreme" community and a "centrist" community (which includes both the S&D and the EPP).

We note that the observation by Mendonca et al. (2015) (that the EP is split in two major political coalitions, one including the EPP, and one including the S&D [73]) does not apply to all policy areas in our research. We also find different results from Cherepnalkoski and Mozetič (2016), who analyzed Twitter data and reported that most MEPs of each political group are contained in a single community [21]. Our results show some overlap with the results by Cherepnalkoski et al. (2016), who observed four blocs of coalitions in the EP (a "left-wing" bloc, "centre" bloc, "centre-right" bloc, and a "far-right" bloc) [20]. However, the

relatively robust partitions in our analysis divide the network into two communities instead of four.

Consequently, we confirm the importance of studying the European Parliament from different points of view, as the corresponding results may differ. It is useful to study different types of data, to have multiple approaches, and to consider different resolution parameters. Lastly, our results suggest that one should take into account the differences between policy areas when one wants to study community structure in the European Parliament.

Chapter 9

Discussion

In this chapter, we summarize our results and present some ideas for further research.

9.1 Conclusion

We started this dissertation by introducing notions from network science. Then, we discussed several methods to study roll-call voting. We focused on one particular method; the construction of a unipartite weighted similarity matrix. We listed the choices one can make when constructing a similarity matrix and we mentioned which options we chose.

In the chapter about community detection, we discussed what communities are, why it is useful to study them, and how they can be detected. We focused on one particular community-detection method; the optimization of modularity. We described strengths and weaknesses of this method, and we introduced one particular heuristic for optimizing modularity; the Louvain method.

In the chapter about the European Parliament, we introduced the political groups and discussed previous network analyses of the European Parliament. Then, we described our data, methods, and results.

We noted that some policy areas seem to be more polarized than others. The most robust partitions split the European Parliament into two communities. It depends on the policy area whether we can assign a label to each of the two communities. For bills related to the "environment" or "gender equality", there seems to be a "left wing" community and a "right wing" community. For bills related to "agriculture" and "industry", there seems to be an "extreme" community and a "centrist" community in the European Parliament. Hence, we conclude that one should take into account the differences between policy areas when one wants to study community structure in the European Parliament.

9.2 Future work

In this section, we discuss more approaches that one could use to study roll-call voting networks (of the European Parliament).

9.2.1 Longitudinal data

We only considered RCVs between 2009 and 2014 (i.e. RCVs in the 7th Assembly of the EP). It would be interesting to consider assemblies of the European Parliament over time and examine whether we can detect different community structures for different assemblies. For instance, Macon et al. (2012) examined longitudinal data for the United Nations General Assembly (1946–2008), and observed a sharp increase in maximum modularity after the end

of the Cold War [68]. For us, it would be interesting to analyse modularity values before and after countries joined the European Union (or to analyze modularity values after a future "Brexit").

9.2.2   Measuring the role of an individual MEP

We considered groups of MEPs (such as political parties), but one could also measure statistics of individual MEPs [31]. For instance, Waugh et al. (2012) discussed individual modularity measures called "solidarity" and "divisiveness", which calculate, respectively, how closely aligned a MEP is to his/her political group and how polarizing a MEP is [124]. One could also measure how often a MEP changes community membership during observation time [64].

9.2.3   Multilayer networks

We independently examined one network for each of the policy areas. To consider different policy areas simultaneously, one could instead use a *multilayer* network [56] in which every layer represents a policy area [10, 64]. Consequently, one can give a weight to each policy area and learn more about robust community structure in the European Parliament. One could also use multilayer networks to study community structure across multiple times [25] or multiple values of resolution parameter γ [79]. Mucha et al. (2010) demonstrated the usefulness of multilayer networks by applying them to the voting record in the U.S. Senate [79].

9.2.4   Overlapping communities

We divided the nodes of a network over disjoint communities (so-called "hard clustering"), but one could also allow communities to overlap (so-called "soft clustering") [31, 33, 42]. In this way, one could examine whether MEPs are part of multiple communities. For instance, an MEP from Greens-EFA could simultaneously belong to a "left wing" community, an "extreme" community, and a "Greens-EFA" community. Hence, it may be more realistic to assume that MEPs belong to multiple communities.

9.2.5   Core–periphery structure

Instead of only studying community structure, one could also examine other types of mesoscale structures, such as core–periphery structure [106]. A "core" is a (central) group of densely connected nodes and a "periphery" is a group of sparsely connected nodes. A node belongs to a core if and only if it is well connected both to other core nodes and to peripheral nodes. One can try to discover core–periphery structure in an entire network [106], or one can try to detect cores inside communities [24, 27]. (Note that a community can have multiple cores and/or multiple peripheries [106]). De Leo et al. (2013) noted that if one removes a node from a core of a community, this will affect a network partition more than if one removes a node from a periphery [27].

Bibliography

[1] Alemán E. (2014). Coauthorship Ties in the Colombian Congress, 2002–2006, *Colombia Internacional*, Vol. **83**, pp. 23–42

[2] Alemán E, Calvo E. (2013). Explaining Policy Ties in Presidential Congresses: A Network Analysis of Bill Initiation Data, *Political Studies*, Vol. **61**, No. 2, pp. 356–377

[3] Alemán E, Calvo E, Jones MP, Kaplan N. (2009). Comparing Cosponsorship and Roll-call Ideal Points, *Legislative Studies Quarterly*, Vol. **34**, No. 1, pp. 87–116

[4] Amelio A, Pizzuti C. (2015). Mining and Analyzing the Italian Parliament: Party Structure and Evolution, *Recommendation and Search in Social Networks*, pp. 249–279

[5] Arenas A, Fernandez A, Gomez S. (2008). Analysis of the Structure of Complex Networks at Different Resolution Levels, *New Journal of Physics*, Vol. **10**, No. 5, p.053039

[6] Attiná F. (1990). The Voting Behaviour of the European Parliament Members and the Problem of the Europarties, *European Journal of Political Research*, Vol. **18**, No. 5, pp. 557–579

[7] Bagrow JP. (2012). Communities and Bottlenecks: Trees and Treelike Networks Have High Modularity, *Physical Review E*, Vol. **85**, No. 6, p.066118

[8] Bakker R, Poole KT. (2013). Bayesian Metric Multidimensional Scaling, *Political Analysis*, Vol. **21**, No. 1, pp. 125–140

[9] Baller I. (2015). Specialists, Party Members or National Representatives: Patterns in Cosponsorship of Amendments in the European Parliament, *22nd International Conference of Europeanists, Council for European Studies*, pp. 1–22

[10] Bassett DS, Porter MA, Wymbs NF, Grafton ST, Carlson JM, Mucha PJ. (2013). Robust Detection of Dynamic Community Structure in Networks, *Chaos: An Interdisciplinary Journal of Nonlinear Science*, Vol. **23**, No. 1, 013142

[11] Bazzi M, Porter MA, Williams S, McDonald M, Fenn DJ, Howison SD. (2016). Community Detection in Temporal Multilayer Networks, with an Application to Correlation Networks, *Multiscale Modeling & Simulation*, Vol. **14**, No. 1, pp. 1–41

[12] Bertozzi AL, Flenner A. (2016). Diffuse Interface Models on Graphs for Classification of High Dimensional Data, *SIAM Review*, Vol. **58**, No. 2, pp. 293–328

[13] Beyle HC. (1931). *Identification and Analysis of Attribute-Cluster-Blocs*, The University of Chicago Press, (discussed in [66])

[14] Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. (2008). Fast Unfolding of Communities in Large Networks, *Journal of Statistical Mechanics: Theory and Experiment*, No. 10, P10008

[15] Boettcher S, Percus AG. (2001). Extremal Optimization for Graph Partitioning, *Physical Review E*, Vol. **64**, No. 2, 026114

[16] Bounova G, de Weck OL. (2012). Overview of Metrics and their Correlation Patterns for Multiple-Metric Topology Analysis on Heterogeneous Graph Ensembles, *Physical Review E*, Vol. **85**, 017117

[17] Brain Connectivity Toolbox by [108]. (Available at `https://sites.google.com/site/bctnet/`), *accessed on 28 August 2016*

[18] Brandes U, Delling D, Gaertler M, Gorke R, Hoefer M, Nikoloski Z, Wagner D. (2008). On Modularity Clustering, *IEEE Transactions on Knowledge and Data Engineering*, Vol. **20**, No. 2, pp. 172–188

[19] Briatte F. (2016). Network Patterns of Legislative Collaboration in Twenty Parliaments, *Network Science*, Vol. **4**, No. 2, pp. 266–271

[20] Cherepnalkoski D, Karpf A, Mozetič I, Grčar M. (2016). Cohesion and Coalition Formation in the European Parliament: Roll-Call Votes and Twitter Activities, *arXiv:1608.04917v1*

[21] Cherepnalkoski D, Mozetič I. (2016). Retweet Networks of the European Parliament: Evaluation of the Community Structure, *Applied Network Science*, Vol. **1**, No. 1, pp. 1–20

[22] Chiru M, Neamtu S. (2011). Parliamentary Representation under Changing Electoral Rules: Co-sponsorship in the Romanian Parliament, *Inaugural General Conference of the ECPR Standing Group on Parliaments, Parliaments in Changing Times (Dublin)*, pp. 1–22

[23] Cho WKT, Fowler JH. (2010). Legislative Success in a Small World: Social Network Analysis and the Dynamics of Congressional Legislation, *The Journal of Politics*, Vol. **72**, No. 1, pp. 124–135

[24] Dal Maso C, Pompa G, Puliga M, Riotta G, Chessa A. (2014). Voting Behavior, Coalitions and Government Strength through a Complex Network Analysis, *PLOS ONE*, Vol. **9**, No. 12, e116046

[25] Danchev V, Porter MA. (2016). Heterogeneity of Global and Local Connectivity in Spatial Network Structures of World Migration, *arXiv:1603.09313*

[26] Danon L, Diaz-Guilera A, Duch J, Arenas A. (2005). Comparing Community Structure Identification, *Journal of Statistical Mechanics: Theory and Experiment*, No. 9, P09008

[27] De Leo V, Santoboni G, Cerina F, Mureddu M, Secchi L, Chessa A. (2013). Community Core Detection in Transportation Networks, *Physical Review E*, Vol. **88**, No. 4, 042810

[28] Doron KW, Bassett DS, Gazzaniga MS. (2012). Dynamic Network Structure of Interhemispheric Coordination, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **109**, No. 46, pp. 18627–18628

[29] Duch J, Arenas A. (2005). Community Detection in Complex Networks Using Extremal Optimization, *Physical Review E*, Vol. **72**, No. 2, 027104

[30] Flake GW, Lawrence S, Giles CL, Coetzee FM. (2002). Self-Organization and Identification of Web Communities, *Computer*, Vol. **35**, No. 3, pp. 66–70

[31] Fortunato S. (2010). Community Detection in Graphs, *Physics Reports*, Vol. **486**, No. 3, pp. 75–174

[32] Fortunato S, Barthélemy M. (2007). Resolution Limit in Community Detection, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **104**, No. 1, pp. 36–41

[33] Fortunato S, Hric D. (2016). Community Detection in Networks: A User Guide, *arXiv:1608.00163*

[34] Fowler JH. (2006). Connecting the Congress: A Study of Cosponsorship Networks, *Political Analysis*, Vol. **14**, No. 4, pp. 456–487

[35] Fowler JH. (2006). Legislative Cosponsorship Networks in the US House and Senate, *Social Networks*, Vol. **28**, No. 4, pp. 454–465

[36] Fruchterman TM, Reingold EM. (1991). Graph Drawing by Force-Directed Placement, *Software: Practice and Experience*, Vol. **21**, No. 11, pp. 1129-1164

[37] Furey JB. (1954). Voting Alignment in the General Assembly, *Doctoral Dissertation Series Columbia University*, (discussed in [66])

[38] Gfeller D, Chappelier JC, De Los Rios P. (2005). Finding Instabilities in the Community Structure of Complex Networks, *Physical Review E*, Vol. **72**, No. 5, 056135

[39] Good BH, de Montjoye YA, Clauset A. (2010). Performance of Modularity Maximization in Practical Contexts, *Physical Review E*, Vol. **81**, No. 4, 046106

[40] Gross JH, Kirkland JH, Shalizi CR. (2012). Cosponsorship in the U.S. Senate: A Multilevel Two-Mode Approach to Detecting Subtle Social Predictors of Legislative Support, *unpublished manuscript*. (Available at `http://www.latinodecisions.com/files/4013/3840/2978/Gross-Kirkland-Shalizi_Multilevel-Cosponsorship_PolAnlys-submission.pdf`), *accessed on 8 July 2016*

[41] Guimerà R, Sales-Pardo M, Amaral LAN. (2004). Modularity from Fluctuations in Random Graphs and Complex Networks, *Physical Review E*, Vol. **70**, No. 2, p.025101

[42] Harenberg S, Bello G, Gjeltema L, Ranshous S, Harlalka J, Seay R, Padmanabhan K, Samatova N. (2014). Community Detection in Large-Scale Networks: A Survey and Empirical Evaluation, *Wiley Interdisciplinary Reviews: Computational Statistics*, Vol. **6**, No. 6, pp. 426–439

[43] Hix S, Noury A, Roland G. (2005). Power to the Parties: Cohesion and Competition in the European Parliament, 1979–2001, *British Journal of Political Science*, Vol. **35**, No. 2, pp. 209–234

[44] Hix S, Noury A, Roland G. (2007). *Democratic Politics in the European Parliament*, Cambridge University Press, (discussed in [20])

[45] Hix S, Noury A, Roland G. (2009). Voting Patterns and Alliance Formation in the European Parliament, *Philosophical Transactions of the Royal Society B*, Vol. **364**, No. 1518, pp. 821–831

[46] Hovet T. (1960). *Bloc Politics in the United Nations*, Harvard University Press, (discussed in [66])

[47] Hric D, Peixoto TP, Fortunato S. (2016). Network Structure, Metadata and the Prediction of Missing Nodes, *arXiv:1604.00255*

[48] Jaccard P. (1912). The Distribution of the Flora in the Alpine Zone, *New Phytologist*, Vol. **11**, No. 2, pp. 37–50

[49] Jakulin A, Buntine W. (2004). Analyzing the US Senate in 2003: Similarities, Networks, Clusters and Blocs. (Available at `polmeth.wustl.edu/files/polmeth/senate4.pdf`), *accessed on 14 August 2016*

[50] Jeub LG, Balachandran P, Porter MA, Mucha PJ, Mahoney MW. (2015). Think Locally, Act Locally: Detection of Small, Medium-Sized and Large Communities in Large Networks, *Physical Review E*, Vol. **91**, No. 1, 012821

[51] Jutla IS, Jeub LGS, Mucha PJ. (2014). A Generalized Louvain Method (version 2.0) for Community Detection Implemented in MATLAB. (Available at `netwiki.amath.unc.edu/GenLouvain`), (*accessed on 24 August 2016*)

[52] Kamada T, Kawai S. (1989). An Algorithm for Drawing General Undirected Graphs, *Information Processing Letters*, Vol. **31**, No. 1, pp. 7–15

[53] Karrer B, Levina E, Newman MEJ. (2008). Robustness of Community Structure in Networks, *Physical Review E*, Vol. **77**, No. 4, 046119

[54] Keller FB. (2015). Networks of Power: Using Social Network Analysis to understand who will rule and who is really in charge in the Chinese Communist Party, *European Consortium for Political Research.*(Available at `http://ecpr.eu/Filestore/PaperProposal/84876886-68d6-412d-995b-5fc5143de350.pdf`), *accessed on 30 April 2016*

[55] Kirkpatrick S, Gelatt CD, Vecchi MP. (1983). Optimization by Simulated Annealing, *Science*, Vol. **220**, No. 4598, pp. 671–680

[56] Kivelä M, Arenas A, Barthélemy M, Gleeson JP, Moreno Y, Porter MA. (2014). Multilayer Networks, *Journal of Complex Networks*, Vol. **2**, No. 3, pp. 203–271

[57] Kovács A. (2014). Political Networks in the European Parliament: Network Analysis of the 2013 Common Agricultural Policy Reform, *Proceedings of FIKUSZ 14 Symposium for Young Researchers*, pp. 107–122

[58] Kraskov A, Stögbauer H, Andrzejak RG, Grassberger P. (2005). Hierarchical Clustering Using Mutual Information, *EPL (Europhysics Letters)*, Vol. **70**, No. 2, pp. 278–282

[59] Krippendorff K. (2012). *Content Analysis: An Introduction to its Methodology*. Sage, (discussed in [20])

[60] Lambiotte R. (2010). Multi-Scale Modularity in Complex Networks, *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), Proceedings of the 8th International Symposium on IEEE*, pp. 546–553

[61] Lancichinetti A, Fortunato S. (2009). Community Detection Algorithms: A Comparative Analysis, *Physical Review E*, Vol. **80**, No. 5, 056117

[62] Lancichinetti A, Fortunato S, Kertész J. (2009). Detecting the Overlapping and Hierarchical Community Structure in Complex Networks, *New Journal of Physics*, Vol. **11**, No. 3, 033015

[63] Lazer D. (2011). Networks in Political Science: Back to the Future, *Political Science & Politics*, Vol. **44**, No. 1, pp. 61–68

[64] Lee SH, Magallanes JM, Porter MA. (2016). Time-dependent Community Structure in Legislation Cosponsorship Networks in the Congress of the Republic of Peru, *Journal of Complex Networks*, doi: 10.1093/comnet/cnw004

[65] Levandowsky M, Winter D. (1971). Distance between Sets, *Nature*, Vol. **234**, No. 5323, pp. 34–35

[66] Lijphart A. (1963). The Analysis of Bloc Voting in the General Assembly: A Critique and a Proposal, *American Political Science Review*, Vol. **57**, No. 4, pp. 902–917

[67] Lyons P, Lacina T. (2009). An Examination of Legislative Roll-call Voting in the Czech Republic Using Spatial Models, *Czech Sociological Review*, Vol. **45**, No. 6, pp. 1155–1190

[68] Macon KT, Mucha PJ, Porter MA. (2012). Community Structure in the United Nations General Assembly, *Physica A: Statistical Mechanics and its Applications*, Vol. **391**, No. 1, pp. 343–361

[69] McCarty N, Poole KT, Rosenthal H. (2009). Does Gerrymandering Cause Polarization?, *American Journal of Political Science*, Vol. **53**, No. 3, pp. 666–680

[70] McClurg SD, Lazer D. (2014). Political Networks, *Social Networks*, Vol. **36**, pp. 1–4

[71] McElroy G, Benoit K. (2007). Party Groups and Policy Positions in the European Parliament, *Party Politics*, Vol. **13**, No. 1, pp. 5–28

[72] Meilă M. (2007). Comparing Clusterings — An Information Based Distance, *Journal of Multivariate Analysis*, Vol. **98**, No. 5, pp. 873-895

[73] Mendonca I, Figueiredo R, Labatut V, Michelon P. (2015). Relevance of Negative Links in Graph Partitioning: A Case Study Using Votes from the European Parliament, *The Second European Network Intelligence Conference (ENIC)*, pp. 122–129

[74] Metz T, Jäckle S. (2016). Hierarchical, Decentralized, or Something Else? Opposition Networks in the German Bundestag, *Legislative Studies Quarterly*, Vol. **41**, No. 2, pp. 501–542

[75] Micozzi JP. (2014). Alliance for Progress? Multilevel Ambition and Patterns of Cosponsorship in the Argentine House, *Comparative Political Studies*, Vol. **47**, No. 8, pp. 1186–1208

[76] Mirshahvalad A, Lindholm J, Derlen M, Rosvall M. (2012). Significant Communities in Large Sparse Networks, *PLOS ONE*, Vol. **7**, No. 3, e33721

[77] MIT Strategic Engineering - Matlab Tools for Network Analysis (2006–2011). (Available at `http://strategic.mit.edu/downloads.php?page=matlab_networks`), *accessed on 26 August 2016*

[78] Moody J, Mucha PJ. (2013). Portrait of Political Party Polarization, *Network Science*, Vol. **1**, No. 1, pp. 119–121

[79] Mucha PJ, Richardson T, Macon K, Porter MA, Onnela JP. (2010). Community Structure in Time-Dependent, Multiscale, and Multiplex Networks, *Science*, Vol. **328**, No. 5980, pp. 876–878

[80] NetVotes by [73]. (Available at `github.com/CompNet/NetVotes`), *accessed on 22 August 2016*

[81] Network Community Toolbox. (Available at `http://commdetect.weebly.com/`), *accessed on 26 August 2016*

[82] Newman MEJ. (2004). Analysis of Weighted Networks, *Physical Review E*, Vol. **70**, 056131

[83] Newman MEJ. (2004). Fast Algorithm for Detecting Community Structure in Networks, *Physical Review E*, Vol. **69**, No. 6, 066133

[84] Newman MEJ. (2006). Finding Community Structure in Networks Using the Eigenvectors of Matrices, *Physical Review E*, Vol. **74**, No. 3, 036104

[85] Newman MEJ. (2006). Modularity and Community Structure in Networks, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **103**, No. 23, pp. 8577–8582

[86] Newman MEJ. (2010). *Networks: An Introduction*, Oxford University Press

[87] Newman MEJ. (2016). Community Detection in Networks: Modularity Optimization and Maximum Likelihood are Equivalent, *arXiv:1606.02319*

[88] Newman MEJ, Clauset A. (2016). Structure and Inference in Annotated Networks, *Nature Communications*, Vol. **7**, No. 11863, pp. 1–11

[89] Newman MEJ, Girvan M. (2004). Finding and Evaluating Community Structure in Networks, *Physical Review E*, Vol. **69**, 026113

[90] Newman MEJ, Reinert G. (2016). Estimating the Number of Communities in a Network, *Physical Review Letters*, Vol. **117**, 078301

[91] Ozer M, Kim N, Davulcu H. (2016). Community Detection in Political Twitter Networks Using Nonnegative Matrix Factorization Methods, *arXiv:1608.01771*

[92] Pajala A, Jakulin A, Buntine W. (2004). Parliamentary Group and Individual Voting Behaviour in the Finnish Parliament in Year 2003: A Group Cohesion and Voting Similarity Analysis. (Available at `citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.103.2295&rep=rep1&type=pdf`), *accessed on 15 August 2016*

[93] Poole KT, Rosenthall H. Voteview Website. (Available at `www.voteview.com`), *accessed on 6 June 2016*

[94] Poole KT, Rosenthall H. (1995). Congress: *A Political–Economic History of Roll Call Voting*. Oxford University Press.

[95] Porter MA, Friend AJ, Mucha PJ, Newman MEJ. (2006). Community Structure in the U.S. House of Representatives, *Chaos*, Vol. **16**, 041106

[96] Porter MA, Mucha PJ, Newman MEJ, Friend AJ. (2007). Community Structure in the United States House of Representatives, *Physica A: Statistical Mechanics and its Applications*, Vol. **386**, No. 1, pp. 414–438

[97] Porter MA, Mucha PJ, Newman, MEJ, Warmbrand CM. (2005). A Network Analysis of Committees in the U.S. House of Representatives, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **102**, No. 20, pp. 7057–7062

[98] Porter MA, Onnela J-P, Mucha PJ. (2009). Communities in Networks, *Notices of the AMS*, Vol. **56**, No. 9, pp. 1082–1097, 1164–1166

[99] Puccio E, Pajala A, Piilo J, Tumminello M. (2016). Structure and Evolution of a European Parliament via a Network and Correlation Analysis, *arXiv:1603.09319*

[100] Rand WM. (1971). Objective Criteria for the Evaluation of Clustering Methods, *Journal of the American Statistical Association* Vol. **66**, No. 336, pp. 846–850

[101] Reichardt J, Bornholdt S. (2006). Statistical Mechanics of Community Detection, *Physical Review E*, Vol. **74**, No. 1, 016110

[102] Rice SA. (1927). The Identification of Blocs in Small Political Bodies, *American Political Science Review*, Vol. **21**, No. 3, pp. 619–627

[103] Rice SA. (1928). *Quantitative Methods in Politics* (New York, 1928), (discussed in [66])

[104] Richardson T, Mucha PJ, Porter MA. (2009). Spectral Tripartitioning of Networks, *Physical Review E*, Vol. **80**, No. 3, 036111

[105] Riggs RE. (1958). Politics in the United Nations, A Study of United States Influence in the General Assembly, *Illinois Studies in the Social Sciences*, Vol. **41**, (discussed in [66])

[106] Rombach MP, Porter MA, Fowler JH, Mucha PJ. (2014). Core–Periphery Structure in Networks, *SIAM Journal on Applied Mathematics*, Vol. **74**, No. 1, pp. 167–190

[107] Rosvall M, Bergstrom CT. (2008). Maps of Random Walks on Complex Networks Reveal Community Structure, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **105**, No. 4, pp. 1118–1123

[108] Rubinov M, Sporns O. (2010). Complex Network Measures of Brain Connectivity: Uses and Interpretations, *NeuroImage*, Vol. **52**, pp. 1059–1069

[109] Schneider V, Lang A, Leifeld P, Gundelach B. (2007). Political Networks – A Structured Bibliography, *Culture*, Vol. **60**, pp. 1–84

[110] Shannon CE. (1948). A Mathematical Theory of Communication, *The Bell System Technical Journal*, Vol. **27**, pp. 379–423, 623–656, (discussed in [49])

[111] Signorelli M, Wit EC. (2016). A Penalized Inference Approach to Stochastic Blockmodelling of Community Structure in the Italian Parliament, *arXiv:1607.08743*

[112] Sirovich L. (2003). A Pattern Analysis of the Second Rehnquist U.S. Supreme Court, *Proceedings of the National Academy of Sciences of the United States of America*, Vol. **100**, No. 13, pp. 7432–7437

[113] Squartini T, Picciolo F, Ruzzenenti F, Garlaschelli D. (2013). Reciprocity of Weighted Networks, *Scientific Reports*, Vol. **3**, 2729

[114] Strehl A, Ghosh J. (2002). Cluster Ensembles — A Knowledge Reuse Framework for Combining Multiple Partitions, *Journal of Machine Learning Research*, Vol. **3**, pp. 583–617

[115] Traud AL, Kelsic ED, Mucha PJ, Porter MA. (2011). Comparing Community Structure to Characteristics in Online Collegiate Social Networks, *SIAM Review*, Vol. **53**, No. 3, pp. 526–543

[116] Traud AL, Mucha PJ, Porter MA. (2012). Social Structure of Facebook Networks, *Physica A: Statistical Mechanics and its Applications*, Vol. **391**, No. 16, pp. 4165–4180

[117] Victor JN. *Handbook of Political Networks*, Oxford University Press

[118] Vinh NX, Epps J. Bailey J. (2010). Information Theoretic Measures for Clusterings Comparison: Variants, Properties, Normalization and Correction for Chance, *Journal of Machine Learning Research*, Vol. **11**, pp. 2837–2854

[119] Visualization of Communities in MATLAB. (Available at `netwiki.amath.unc.edu/VisComms/VisComms`), *accessed on 4 September 2016*

[120] Voeten E. (2000). Clashes in the Assembly, *International Organization*, Vol. **54**, No. 2, pp. 185–215

[121] VoteWatch Europe. (Available at `www.votewatch.eu`), *accessed on 22 August 2016*

[122] Third Report VoteWatch.Eu - Voting in the 2009–2014 European Parliament How do MEPs Vote after Lisbon? (Available at `www.votewatch.eu/blog/wp-content/uploads/2011/01/votewatch_report_voting_behavior_26_january_beta.pdf`), *accessed on 6 June 2016*

[123] Ward MD, Stovel K, Sacks A. (2011). Network Analysis and Political Science, *Annual Review of Political Science*, Vol. **14**, pp. 245–264

[124] Waugh AS, Pei L, Fowler JH, Mucha PJ, Porter MA. (2012). Party Polarization in Congress: A Network Science Approach, *arXiv:0907.3509v3*

[125] van Wijk BCM, Stam CJ, Daffertshofer A. (2010). Comparing Brain Networks of Different Size and Connectivity Density Using Graph Theory, *PLOS ONE*, Vol. **5**, No. 10, e13701

[126] Yang J, Leskovec J. (2015). Defining and Evaluating Network Communities Based on Ground-Truth, *Knowledge and Information Systems*, Vol. **42**, No. 1, pp. 181–213

[127] Zanin M, Papo D, Sousa PA, Menasalvas E, Nicchi A, Kubik E, Boccaletti S. (2016). Combining Complex Networks and Data Mining: Why and How, *Physics Reports*, Vol. **635**, pp. 1–44

[128] Zhang Y, Friend AJ, Traud AL, Porter MA, Fowler JH, Mucha PJ. (2008). Community Structure in Congressional Cosponsorship Networks, *Physica A: Statistical Mechanics and its Applications*, Vol. **387**, No. 7, pp. 1705–1712