Computational Social Networks
a SpringerOpen Journal

RESEARCH   **Open Access**

Polarization of coalitions in an agent-based model of political discourse

Philip Leifeld

Correspondence:
philip.leifeld@uni-konstanz.de
University of Konstanz, Box 216,
78457 Konstanz, Germany

Abstract

Political discourse is the verbal interaction between political actors in a policy domain. This article explains the formation of polarized advocacy or discourse coalitions in this complex phenomenon by presenting a dynamic, stochastic, and discrete agent-based model based on graph theory and local optimization. In a series of thought experiments, actors compute their utility of contributing a specific statement to the discourse by following ideological criteria, preferential attachment, agenda-setting strategies, governmental coherence, or other mechanisms. The evolving macro-level discourse is represented as a dynamic network and evaluated against arguments from the literature on the policy process. A simple combination of four theoretical mechanisms is already able to produce artificial policy debates with theoretically plausible properties. Any sufficiently realistic configuration must entail innovative and path-dependent elements as well as a blend of exogenous preferences and endogenous opinion formation mechanisms.

Keywords: Political discourse; Policy debates; Discourse coalitions; Advocacy coalitions; Polarization; Social balance

Background

Political discourse is a complex phenomenon. Despite its intriguing prevalence in everyday politics, there have been only few attempts at developing explanations of how discourse works, and even fewer attempts at modeling this apparently ill-defined phenomenon in a formal way. Political discourse, as it shall be analyzed here, is the verbal interaction between political actors in a policy domain. For example, in the debate on nuclear energy policy, a number of state and non-state actors speak up in the media and call for specific policy instruments or make specific claims on the basis of their causal beliefs. This phenomenon is based on a complex set of properties:

First, political discourse is *dynamic* because political actors repeatedly participate in a policy debate. Consecutive statements of actors rather than simultaneous moves constitute the essence of a discourse.

Second, political discourse is *relational*. Neither are actors insulated when they make their claims in a debate nor are their statements randomly distributed across actors in the policy domain or over time [1,2]. There is rather an interaction in which actors frequently refer to what others said before. For this reason, social network analysis [3,4] is an appropriate tool for the measurement of empirical discourses. Ward et al. [5] conclude in their

review article that network analysis 'offers a means of addressing one of the holy grails of the social sciences: effectively analyzing the interdependence and flows of influence among individuals, groups, and institutions' [5].

Third, the motivation of actors to participate in the discourse is based on the *uncertainty* of other actors. Political discourse involves both 'power' and 'puzzling' over optimal policy design [6]. The primary reason for the existence of political discourse is (Knightian) uncertainty over causal relationships ('does policy *x* lead to outcome *y*?') and future constraints ('if we implement policy *x* to solve problem *y*, will this relationship hold true in the future?'). On the other hand, this uncertainty is exploited by those who have vested interests. By making certain kinds of claims in a discourse, organizations try to 'convince' the undecided to adopt their claims; hence, discourse is also *goal-driven* and can be understood as an exercise of power [7].

Fourth, as discourse is based on a dynamic, non-random interaction between actors, it tends to be *path-dependent* [8] because actors rarely come up with completely new ideas in a discourse. If all statements were based on completely new thoughts, the incentive to make a statement in the discourse in the first place would vanish because it would no longer be possible to achieve any goals by participating in the discourse.

Fifth, discourse is *fragmented* or *polarized*. Public policy research has shown that policy domains and hence discourses are composed of several competing coalitions of actors [2,9-11]. These advocacy coalitions or discourse coalitions tend to reiterate cohesive sets of policy beliefs or claims, but there is little belief overlap across coalitions. This reflects the notion of social balance as found in other mathematical and computational models of cultural segregation and polarization [12,13].

Sixth, however, the literature on advocacy coalitions does argue that *learning* takes places within and across coalitions [9]. In other words, coalitions in a discourse are fragmented or polarized but not completely insulated.

Seventh, political discourse is an out-of-equilibrium phenomenon. Given the assumption that participation in a discourse is instrumental for achieving goals like policy platforms or merely the reduction of uncertainty about the problem ahead, a discourse keeps moving forward along the time axis as long as the goals of almost any participating individual have not been fully met. As such an equilibrium rarely ever occurs in politics, political discourse as an out-of-equilibrium state has a possibly infinite time horizon. If it were in equilibrium, either actors would eiterate the same sets of ideas in a certain order over and over again or the discourse would end because previous contradictions between antagonists would have been resolved (as argued by deliberative democracy scholars and strongly opposed by liberal democracy scholars; for a discussion, see [14]).

The question then becomes how individuals select the normative concepts they advocate in the public such that the discussion remains ongoing and that the above requirements are met. The professed goal of the models presented below is therefore to ascertain under what theoretical micro-level conditions a political discourse as a meso or macro phenomenon remains out of equilibrium for a substantial amount of time and shows the properties outlined above.

Methods

I will present a dynamic, stochastic, and discrete agent-based model which rests mainly on graph theory and local optimization[a]. As Epstein argues, agent-based models are an

'analytical-computational approach to non-equilibrium social systems' [15]. For this reason, they are likely to be the most powerful tool available for theorizing about opinion dynamics, both in political and other social contexts. And indeed, agent-based models have been relatively often used to study opinion formation [16-20] while their prevalence is still comparatively low in political science in general [21,22].

Agent-based models follow a 'generative paradigm' [15] because they serve to model agents' individual behavior in a bottom-up fashion and then evaluate the resulting macro structure against the benchmark of empirical observations or theoretical expectations of this macro structure. For this reason, agent-based models have been called 'thought experiments' in the literature [23].

This is precisely the research design employed here. Some simple behavioral rules are designed about how agents decide which claims they want to contribute to the discourse. Given each mechanism, how will the behavior of the agents affect the macro structure of the policy-domain-wide political discourse? Moreover, how do these separately plausible assumptions have to be combined to yield artificial discourses which are compatible with the theoretical properties of discourses outlined above? Social network analysis [5] is employed as a measurement model of the resulting discourse.

The model presented hereafter is modular insofar as various kinds of structural or individual effects can be easily plugged into it or omitted from it. This design principle allows to start off with a baseline model of completely random behavior, then testing of effects of any single behavioral mechanism, and later on combining these mechanisms by attaching weights to each of them.

Definitions and basic setup

In each round of the discourse, an actor makes a statement by publicly choosing and announcing a concept. Let $A = \{a_1, \ldots, a_i, \ldots, a_m\}$ be the set of actors in the discourse, with i being the index of the evaluating actor (*ego*) and i' being the index of an arbitrary *alter*. Every single actor is associated with a randomly assigned ideal point on a single ideological dimension (e.g., leftist versus rightist ideology). While drawing the ideological ideal points of actors from a continuous, bimodal random distribution would be more compatible with conventional political science models, the model described here rather fixes the ideal points of actors at the extreme points of the ideological dimension, that is, $P(\Phi(a) = \phi(a)) = \frac{1}{2}$ with $\Phi = \{0; 1\}$, in order to yield a clean bipolar ideology effect. This dichotomous ideology is exogenously given and fixed over time. In theoretical terms, it corresponds to what Sabatier calls 'deep core beliefs' [9] - unalterable normative ideal points which are the source of much of the ideological conflict that is observable in political discourses. These ideology constants work like in spatial models of politics where the ideal points of agents are exogenously given[b]. Actors are also exactly one out of two types: an interest group or a governmental actor, irrespective of their ideology score. The actor type is captured by a dummy variable $\theta(a_i) = \{0; 1\}$, which equals 1 iff a_i is a governmental actor.

Furthermore, let $C = \{c_1, \ldots, c_j, \ldots, c_n\}$ be the set of concepts, or claims, in the discourse. A random ideological score $\phi(c) \sim U(0, 1)$ is associated with each c, with one extreme of this dimension representing leftist values and the other end representing rightist values. This ideological dimension is the same as for the actors, but the ideological scores of the concepts are drawn from a continuous uniform distribution.

Each actor shall have a publicly visible history of the o last statements a_i has made in the media. Let $H_i = \{h_{i,1}, \ldots, h_{i,t}, \ldots, h_{i,o}\}$ be the history of actor i with concepts at time t. By default, the history has $o = 5$ rounds with o being the most recent round, $o - 1$ the second-most recent round, etc. The current decision about which concept to choose is being made at $o + 1$ and thus does not yet belong to the history.

In each round, exactly one actor makes a statement. The actor is chosen randomly with a probability of κ that a governmental actor is chosen. By default, there are five leftist governmental actors, five rightist governmental actors, 15 leftist interest groups, and 15 rightist interest groups, and $\kappa = 0.6$. These proportions roughly correspond to empirical discourses in various policy domains [1,2,10]. By default, C comprises eight different normative concepts. This number is sufficiently large to yield meaningful results, but it is low enough to ensure a reasonable computing time.

Actors make their concept choices based on a behavioral mechanism or based on a custom blend of eight possible mechanisms. In each round, actor i computes an aggregated score for each concept, which is based on various criteria and then adopts the concept with the highest score. The details of these criteria or mechanisms and the aggregation via a utility function are described in the following paragraphs.

Exogenous ideology

Ego assigns a score to c_j which corresponds to the similarity between his or her ideology and the ideology of c_j:

$$\text{XI}_i\left(c_j\right) := 1 - \left|\phi(a_i) - \phi\left(c_j\right)\right| \tag{1}$$

This mechanism represents the classic view that interest groups have a rather fixed set of preferences and try to pull the outcome of policy-making into their ideological directions. Actors make normative statements in the discourse which are close to their own ideal point [24]. It should be straightforward to assume that interest groups will rather express statements that support their ideology than statements running against their own ideology.

Endogenous ideology

Ideology does not have to be exogenously given. It can be conceptualized as being path-dependent such that ego assesses whether a concept is ideologically compatible with the concepts ego has named before. Ego computes the average (or total) similarity between the concepts in his or her history and c_j:

$$\text{NI}_i\left(c_j\right) := \sum_{t=1}^{t=o} \left(1 - \left|\phi\left(h_{i,t}\right) - \phi\left(c_j\right)\right|\right) \tag{2}$$

Concept popularity or preferential attachment to concepts

For a given actor i and a given concept j, let M be the set of alters i' who have announced concept j in the latest round, o.

$$M_{i,j} := \left\{i' \mid i' \neq i \wedge a_{i'} \in A \wedge h_{i',t=o} = c_j\right\} \tag{3}$$

Ego assigns the number of alters in the set as the score:

$$\text{CP}_i\left(c_j\right) = \left|M_{i,j}\right| \tag{4}$$

In other words, ego chooses a concept that is currently popular. This may be one of the drivers of 'political waves' [25] or 'issue attention cycles' [26]; topics or claims within a debate come and go in 'waves' because actors jump on the bandwagon once another actor has started to use a concept in the discourse.

A politician or a party that just ignores claims brought up by others will not be re-elected, and similarly a government department will be criticized if it ignores important facts of the problem it is in charge of.

In the literature on social networks, the choice of already popular nodes is known as 'preferential attachment' and has been described as one of the building blocks of network formation [27].

Actor similarity or coalition formation

The actor similarity score follows from the prediction of the Advocacy Coalition Framework [9] and other interest group theories [28] that interest groups are organized in coalitions. According to the former theory, actors learn more easily from other actors if they are in the same coalition. A coalition, however, is not a formal arrangement. It is rather constituted by the similarity of 'belief systems' between actors [9]. Consequently, actors identify other actors who are in the same coalition by measuring their overall belief similarity and then adopt their beliefs.

More precisely, for any given concept j, an actor first considers the histories of all other actors (that is, their last five concepts) who recently chose j and computes their similarities to his or her own history list by counting the co-occurrences of concepts. Actor i chooses the concept which maximizes the fit between actor i and actors i' who recently chose j. This method directly implements the idea of adaptation from actors who have similar belief systems and is thus a device of discursive coalition formation. Given Equation 3, ego assigns the following score to concept j:

$$\mathrm{AS}_i\left(c_j\right) = \begin{cases} 0 & \text{if } M_{i,j} = \emptyset \\ \frac{\sum_{i' \in M} |H_i \cap H_{i'}|}{|M_{i,j}|} & \text{if } |M_{i,j}| \geq 1, \end{cases} \tag{5}$$

where $|H_i \cap H_{i'}|$ is the size of the intersection of the histories of a_i and $a_{i'}$ and where the whole numerator captures the total similarity of a_i to all $a_{i'}$ who announced c_j in the last round. The function therefore assigns a value to c_j which is equivalent to the average similarity of the history of a_i to other actors' histories who recently announced c_j and thus captures the degree to which c_j was recently chosen by similar actors (i.e., actors from the same coalition) and operationalizes coalition formation.

Concept similarity or following the wisdom of the crowds

Define G_j as the set of alters who contain concept c_j in their history:

$$G_{i,j} := \left\{ i' \mid i' \neq i \wedge a_{i'} \in A \wedge c_j \in H_{i'} \right\} \tag{6}$$

Let $s(c_j, c_{j'})$ be the similarity between two concepts as measured by the number of actors who have both concepts in their histories:

$$s(c_j, c_{j'}) = \left| G_{i,j} \cap G_{i,j'} \right| \tag{7}$$

Ego sums up the similarities between all of his or her past concepts and concept *j*. The result is the concept similarity score of c_j:

$$\text{CS}_i\left(c_j\right) = \sum_{t=1}^{t=o} s\left(c_j, h_{i,t}\right) \tag{8}$$

In other words, a co-occurrence matrix is created which contains the similarities between all concepts. Two concepts are more similar, the more often an actor refers to both of them in his or her history. For every concept in the discourse, the method then adds up the similarities between this concept and each item in the history list of the current actor. This results in the total similarity between a concept and the current actor, based on both his or her history and the evaluations of all other actors.

One rationale for this method is the tendency of subjects to deny information that does not conform to their own beliefs [29]. Concepts should be as compatible to one's own prior statements as possible, and other actors' use of concepts may be an important data source for evaluating how similar two concepts actually are.

At the same time, political actors strive for new arguments and concepts in support of their claims ([30], pp. 290). Actors try to position themselves in the discourse, avoid being isolated with their claims, and they closely observe which other concepts are both popular and close to their own position. If they adopt such a claim, it is more likely that other actors will refer to them and in turn adopt their other concepts.

Actor's history or self-consistency

Ego assigns a positive score to concept *j* if ego named the concept in his or her history, otherwise 0. The more recent the statement, the higher the score:

$$\text{AH}_i\left(c_j\right) = \begin{cases} \max\left\{t \mid c_j = h_{i,t}\right\} & \text{if } c_j \in H_i \\ 0 & \text{if } c_j \notin H_i \end{cases} \tag{9}$$

This mechanism stems from social psychology. Individuals (and organizations as collective actors) try to maximize consistency with their previous statements. This results in an actor's own latest statement scoring highest, the one before second-highest, etc. Concepts which are not in the actor's history receive the minimum score.

This method is based on cognitive dissonance theory and its descendants [31,32]: Aronson's self-consistency theory implies that actors strive for consistent views of themselves in order to avoid dissonance [29]. Cialdini et al. define consistency as the 'tendency to base one's responses to incoming stimuli on the implications of existing (prior entry) variables, such as previous expectancies, commitments and choices' [33]. Actors weight their own previous statements higher than other concepts because they want to maintain and communicate a coherent image of themselves, leading to path dependence of an actor's statement choice.

Rare concepts or agenda setting

Particularly interest groups find it worthwhile to revive a sleeping discourse. If there is a concept that has not been actively discussed for quite some time, it proves to be a good opportunity to pick up this concept and re-introduce it into the discourse in order to pull the debate into a certain direction. Actors may thus have a tendency to make claims which

are not on many actors' agenda in order to act as an agenda setter. Especially those rare concepts are attractive which are close to the actor's own ideal point. Therefore, it makes sense to combine the mechanism of choosing rare concepts with the ideology mechanism described above.

The agenda-setting score is the number of alters who have concept c_j in their history divided by the number of alters in the discourse, subtracted from 1.

$$\text{RC}_i\left(c_j\right) = 1 - \frac{\left|G_{i,j}\right|}{m-1} \tag{10}$$

Based on Equation 6, this strategy resembles inverse concept popularity (but concerning all rounds in the history, not just the latest round).

Government coherence

As an extension of Equation 6, let $B_{i,j}$ be the set of alters who hold concept c_j in their history and who are governmental actors.

$$B_{i,j} := \left\{i' \mid i' \neq i \wedge a_{i'} \in A \wedge c_j \in H_{i'} \wedge \theta(a_i) = 1\right\} \tag{11}$$

Government coherence is the number of alters who are governmental actors and who have c_j in their history:

$$\text{GC}_i\left(c_j\right) = \left|B_{i,j}\right| \tag{12}$$

For every concept, the government coherence method counts by how many (other) governmental actors the concept was previously chosen, i.e., in how many (other) governmental actors' history lists it is present. The method serves two different purposes for interest groups and governmental actors.

For governmental actors, this in an important procedure because they usually share common goals, despite potential ideological differences due to coalition governments. In most political systems, the various government departments are tied together by common objectives, parties, coalition contracts, or presidents. The aim of the method is to unify governmental actors by aligning them with each other.

For interest groups, this method defines the degree to which they adhere to the collective ideal point of the government. In consensual or corporatist political systems, the importance of this mechanism for interest groups should be high, while it should be low in majoritarian or pressure-pluralist systems [34].

Normalization

Each of these variables, or *evaluation functions*, yields a score that an actor assigns to a concept. Their ranges (in the statistical sense) differ substantially. However, they should be on the same scale in order to be comparable in a meaningful way. The following procedure thus converts the list of scores into a ranking list with 1 being the lowest rank and n (the number of concepts) being the highest or best rank. Two items with the same frequency (that is, ties) should be assigned the lower of the two rank positions. For example, if there are three concepts with the same score, and there are two items with a lower and three items with a higher score than the three items under consideration, the three items should all be assigned rank 3.

Let $f(c_j)$ be an evaluation function that assigns a score to concept j. Furthermore, let

$$D = \left(f(c_1), \ldots f(c_j), \ldots, f(c_n) \right) = (d_1, \ldots, d_l, \ldots, d_n) \tag{13}$$

be an n-tuple of assigned scores such that $\forall l < n : d_l \leq d_{l+1}$. Then, let $g : D \twoheadrightarrow S$ be a surjection with

$$g(d_l) = \begin{cases} 1 & \text{if } l = 1 \\ l & \text{if } l \geq 2 \wedge d_l > d_{l-1} \\ g\left(d_{l-1}\right) & \text{if } l \geq 2 \wedge d_l = d_{l-1} \end{cases} \tag{14}$$

Finally, $(g \circ f)(c_j)$ is the normalized value of $f(c_j)$.

A fictitious example: Assume there are four concepts with different scores resulting from the concept popularity evaluation function, $\mathrm{CP}_i(c_1) = 14$, $\mathrm{CP}_i(c_2) = 3$, $\mathrm{CP}_i(c_3) = 14$, and $\mathrm{CP}_i(c_4) = 19$. Then, D is the ordered set $\{3, 14, 14, 19\}$. The normalization function assigns the values $\{1, 2, 2, 4\}$ because the first item matches the first rule of the surjection, the second and fourth items the second rule, and the third item the third rule of the g function.

Utility functions

Assuming that there are $p = 8$ evaluation functions, the utility of actor i for choosing concept j is defined as

$$\begin{aligned} U_i^{\mathrm{ig}}(c_j) = &\ \beta_1 g\left(\mathrm{XI}_i(c_j)\right) + \beta_2 g\left(\mathrm{NI}_i(c_j)\right) + \beta_3 g\left(\mathrm{AS}_i(c_j)\right) + \beta_4 g\left(\mathrm{CS}_i(c_j)\right) \\ &+ \beta_5 g\left(\mathrm{AH}_i(c_j)\right) + \beta_6 g\left(\mathrm{CP}_i(c_j)\right) + \beta_7 g\left(\mathrm{RC}_i(c_j)\right) + \beta_8 g\left(\mathrm{GC}_i(c_j)\right) \end{aligned} \tag{15}$$

with $\beta \in \mathbb{R}$ being an arbitrary weight. As mentioned above, there are governmental actors and interest groups. The utility function with the β weights introduced in Equation 15 is only applicable to interest groups. A second utility function for governmental actors exists where the β weights are replaced by γ weights. In all other regards, the functions are identical. This allows for a subclass of model specifications where there is in fact only a single actor type, that is, in the case where $\forall k : \beta_k = \gamma_k$. Governmental actors thus maximize the following utility function:

$$\begin{aligned} U_i^{\mathrm{gov}}(c_j) = &\ \gamma_1 g\left(\mathrm{XI}_i(c_j)\right) + \gamma_2 g\left(\mathrm{NI}_i(c_j)\right) + \gamma_3 g\left(\mathrm{AS}_i(c_j)\right) + \gamma_4 g\left(\mathrm{CS}_i(c_j)\right) \\ &+ \gamma_5 g\left(\mathrm{AH}_i(c_j)\right) + \gamma_6 g\left(\mathrm{CP}_i(c_j)\right) + \gamma_7 g\left(\mathrm{RC}_i(c_j)\right) + \gamma_8 g\left(\mathrm{GC}_i(c_j)\right) \end{aligned} \tag{16}$$

In either case, actors choose an optimal concept which maximizes their utility:

$$c^* = \arg\max_{c_j \in C} U_i\left(c_j\right) = \arg\max_{c_j \in C} \sum_{p=1}^{8} \beta_p \left(\left(g \circ f_p\right)(c_j) \right) \tag{17}$$

If there are several optimal concepts, the actor chooses from a discrete uniform distribution:

$$\left| C^* \right| > 1 \Rightarrow P(C = c) = \frac{1}{|C^*|} \tag{18}$$

This step ensures that there is always exactly one concept per actor and round.

The discourse model is run for several thousand time steps. At each step, an actor is selected according to the probability rule described above, and this actor makes a publicly visible statement according to the outcome of the utility function, thereby updating his or her history of concepts.

Measurement

As shown in the literature on 'discourse networks' [1,2,35,36], network analysis can be employed to study empirical aspects of political discourse like the shape, stability, and coherence of discourse or advocacy coalitions, cleavage lines in a policy domain, diversity of arguments, and the degree of polarization of a discourse. A discourse can be operationalized as follows. If A is the set of actors in a discourse and C denotes the set of concepts in the discourse, then a bipartite graph $G^{\text{aff}} = \left(A, C, E^{\text{aff}}\right)$ with edges $e^{\text{aff}}(a, c) \in E^{\text{aff}}$ can be constructed. The *aff* superscript indicates that this is an affiliation network, or a bipartite graph. The bipartite graph can be converted into a one-mode projection, or 'co-occurrence' network, $G^{\text{a}} = (A, E^a)$, where the superscript a denotes a network composed only of actors. Formally, this can be achieved by considering neighbors of actors in the bipartite graph. The set of neighbors is defined as the collection of concepts which are incident to an actor, that is, $N_{G^{\text{aff}}}(a) := C \in E^{\text{aff}}(a)$. Accordingly, edge weights between actors in G^{a} are computed as $w(a, a') = |N_{G^{\text{aff}}}(a) \cap N_{G^{\text{aff}}}(a')|$, which amounts to the number of concepts two actors share. The resulting network provides a cross-sectional map of the discursive landscape of political actors. Ties between actors show their discursive similarity; the absence of ties between actors or groups of actors represents discursive dissimilarity. The full array of network-analytic methods can be used to describe the discourse in a precise way, e.g., the degree of polarization between groups of actors, the number of components or clusters as instances of discourse coalitions or advocacy coalitions, or the change of the discursive network structure over time (and thus discursive equilibria).

This measurement model is used to analyze the simulation outcomes as follows. After each new statement, the collective histories of all actors are visualized as an actor × actor co-occurrence network over all concept histories of actors. Beyond visual inspection, five particular statistics are used to analyze the resulting networks: a new measure of ideological polarization, betweenness centralization, the number of components as a share of the initial number of concepts, the proportion of concepts still alive, and the concept replacement rate in the histories of all agents. These measures are discussed below.

Each of the five indices is calculated after a single new step in the discourse. Every particular configuration of the utility functions (in terms of the β and γ weights) of interest is simulated 100 times in order to guarantee that the results are reliable. Since there are 100 simulation runs, there is a random sample of 100 observations for each measure per time step. The mean values of the 100 simulations are plotted as a time series, and this is done for each of the five indices (see Figure 1). Dashed lines around the time series lines represent the 95% confidence interval. Ninety-five out of the 100 simulations lie within these boundaries. The procedure is repeated for several different configurations of the utility functions.

While the agent-based model was implemented in the programming language Java, additional capabilities from the JGraphT[c] and RepastJ[d] add-on libraries are used for the analysis.
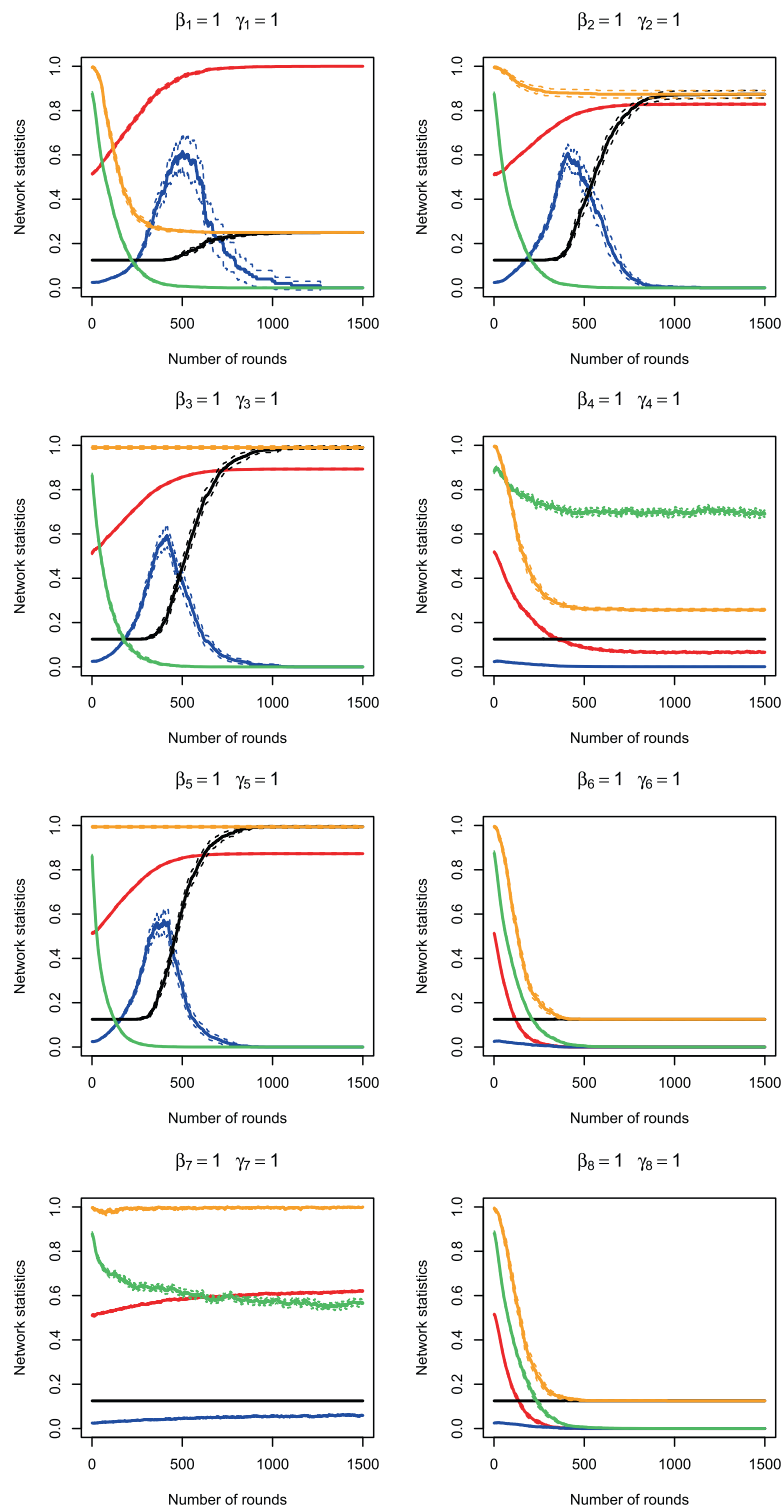
Figure 1 Simple configurations. Results of 100 simulation runs. 95% confidence intervals as dashed lines. Ideological polarization (red), share of concepts still alive (orange), share of recent concept changes (green), betweenness centralization (blue), and number of components as a share of the initial number of concepts (black). β = interest group and γ = governmental actor. β_1 and γ_1 = exogenous ideology (XI), β_2 and γ_2 = endogenous ideology (NI), β_3 and γ_3 = actor similarity/coalition formation (AS), β_4 and γ_4 = concept similarity/wisdom of the crowds (CS), β_5 and γ_5 = actor's history/self-consistency (AH), β_6 and γ_6 = concept popularity (CP), β_7 and γ_7 = rare concepts/agenda setting (RC), and β_8 and γ_8 = government coherence (GC).

Betweenness centralization

The primary structural feature of interest is whether everybody aligns with everybody else, or whether the network tends to be composed of factions or 'coalitions' [9,11]. Several studies have shown that empirical discourses tend to be composed of distinct coalitions with few bridging ties or nodes [1,2,35]. Betweenness centralization [37] captures this aspect by measuring the tendency of a vertex to act as a bridge between many other vertices. The notion of betweenness is operationalized by the number of shortest paths (geodesics, g_{jk}) the vertex is situated on ($g_{jk}(n_i)$), standardized by the number of dyads not involving the vertex for which betweenness centrality is being calculated ([4], pp. 190):

$$C'_B(n_i) = \frac{\sum_{j<k} \frac{g_{jk}(n_i)}{g_{jk}}}{\frac{(m-1)(m-2)}{2}} \qquad (19)$$

Centralization (in contrast to centrality) is a network-level index which sums up the differences between the highest centrality value found in the network, denoted as $C'_B(n^*)$, and the centrality values of all other nodes, $C'_B(n_i)$, and divides this sum by the maximum sum that is theoretically possible:

$$\text{Betweenness centralization} = \frac{\sum_{i=1}^{g} \left[C'_B(n^*) - C'_B(n_i) \right]}{\max \sum_{i=1}^{g} \left[C'_B(n^*) - C'_B(n_i) \right]} \qquad (20)$$

Centralization thus captures the tendency of a network to have few (in the extreme case: one) very central and many peripheral actors ([4], pp. 177). Applied to the problem at hand, betweenness centralization measures the tendency of a network to have very few vertices that interconnect distinct factions in the network. The drawbacks of this measure are that betweenness centralization is zero if the factions lose their interconnection completely and that the factions do not necessarily correspond to ideological cliques. Therefore, another measure of ideological polarization is introduced below.

Ideological polarization

Ideological polarization is a variant of assortative mixing by scalar properties [38]. It measures the degree to which vertices with a similar ideology score are connected and dissimilar actors are disconnected and thus to what extent the whole network is polarized with regard to the nodal attribute 'ideology'. Polarized networks exhibit two ideological clusters.

To compute nodal attribute polarization, three equations are necessary. First, the sum of absolute ideological differences between *non-connected* (separated) actors has to be computed:

$$\forall i: \ i \neq i' \wedge H_i \cap H_{i'} = \emptyset \text{ and } \forall i': \ i' \neq i \wedge H_i \cap H_{i'} = \emptyset:$$

$$\text{ID}_{\text{sep}} = \sum_{i=1}^{m} \sum_{i'=1}^{m} |\phi_i - \phi_{i'}| \qquad (21)$$

Second, calculate the sum of absolute ideological differences between *connected* actors, but this time, the difference has to be multiplied by the edge weight each time:

$$\forall i: \ i \neq i' \wedge H_i \cap H_{i'} \neq \emptyset \text{ and } \forall i': \ i' \neq i \wedge H_i \cap H_{i'} \neq \emptyset:$$

$$\text{ID}_{\text{con}} = \sum_{i=1}^{m} \sum_{i'=1}^{m} (|\phi_i - \phi_{i'}| \cdot |H_i \cap H_{i'}|) \qquad (22)$$

Note that the vertical bars denote absolute values in the first case and cardinality of the set in the second case. Finally, ideological polarization can be measured as

$$\text{Polarization} = \frac{\text{ID}_{\text{sep}}}{\text{ID}_{\text{sep}} + \text{ID}_{\text{con}}} \tag{23}$$

These equations define the polarization measure between 0 and 1. Values of 0.5 do not show any association, values close to 1.0 a strong positive association, and values close to 0.0 a strong negative association.

In contrast to betweenness centralization, the polarization score remains high if two distinct components are completely separated, and the measure captures only ideological polarization, not any other polarization tendencies such as endogenous coalition formation. On the other hand, betweenness centralization does not become obsolete because it is still a useful measure in situations where polarization occurs along other dimensions.

In comparison to assortative mixing by scalar properties [38], this ideological polarization measure is compatible with valued graphs as employed here, it is applicable in cases where a maximum of two coalitions or communities is possible, it scales between 0 and 1 (rather than -1 and $+1$), and it has a simpler formulation.

Number of components

A component is a subgraph that is not connected to the remaining network. If the discourse gets so bi- or multipolar that two or more separate components exist, there is no more common ground between different factions. This may occasionally happen in real-world policy debates, but the situation should be reversible - usually after a couple of rounds.

To obtain a standardized index between zero and one, the number of components is divided by the initial number of concepts in the discourse because this corresponds to the maximum number of components possible.

Proportion of concepts still alive

The fourth index measures the integrity of the discourse by counting how many concepts are still alive (that is, mentioned by at least one actor in his or her latest round). In a healthy and ongoing debate, at least two thirds of the initial ideas should still be present after a substantial amount of time. This integrity measure is defined as follows:

$$\text{Concepts still alive} = \frac{\left|\left\{c_j \mid \exists a_i \in A : c_j = h_{i,t=o}\right\}\right|}{n} \tag{24}$$

Concept attrition is high if this integrity index is low.

Number of recent concept changes

The following replacement index captures the number of recent concept changes by counting the number of actors whose latest concept differs from his or her concept in the previous round, divided by the number of actors in total:

$$\text{Recent concept changes} = \frac{\left|\left\{a_i \mid h_{i,t=o} \neq h_{i,t=o-1}\right\}\right|}{m} \tag{25}$$

The replacement index thus measures in how far the discourse is in motion and can serve as an early indicator of whether the discourse is in equilibrium or not. The goal is to obtain a replacement level that is neither close to 0 nor close to 1 in order to yield a steady-state but out-of-equilibrium discourse as described at the beginning of the article.

Results and discussion

The first question that is going to be answered in this section is assuming that one of the eight mechanisms elaborated above is exclusively at work in a discourse, then how does this affect the overall properties of the discourse over time? After answering this question for each of the eight mechanisms separately, interactions of several mechanisms are evaluated in terms of their effects on discourse evolution.

As a first validity check, a null model without any mechanism in the utility functions is run, that is, $\forall k : \beta_k = \gamma_k = 0$. As expected, the result is an Erdős-Rényi random graph.
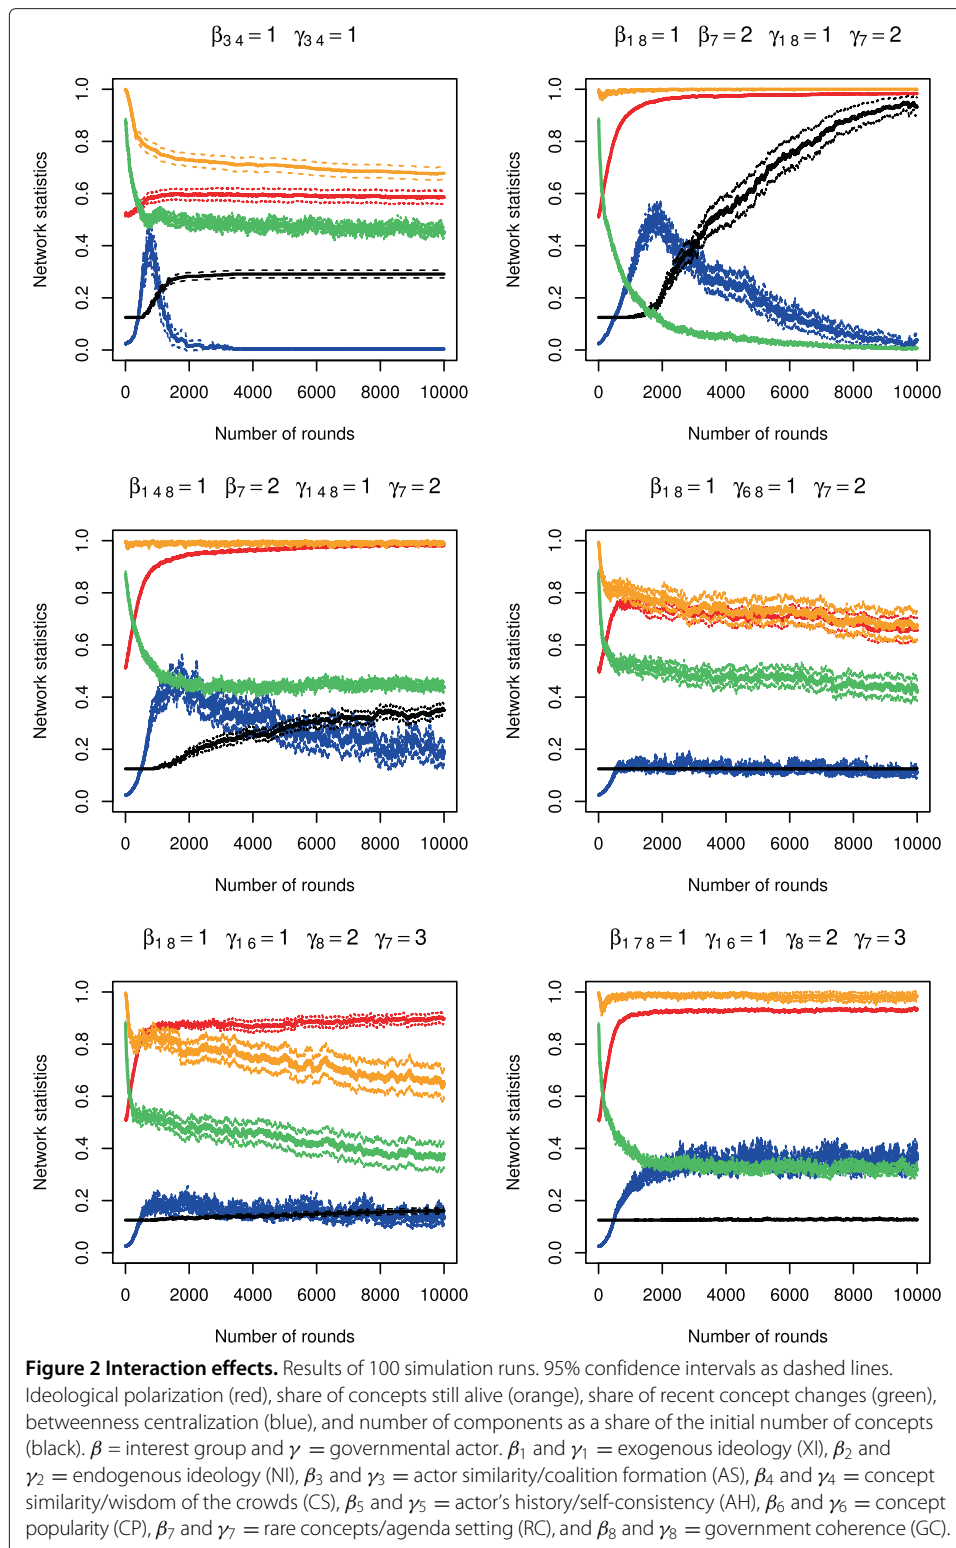
Figures 1 and 2 show the results of 14 different configurations of the objective function. The first eight configurations (Figure 1) test the single effects of each of the eight components of the utility functions. In each diagram, there is an initial rewiring process, or a phase transition, from an Erdős-Rényi random graph to an equilibrium based on the specific parameter settings [39].

Analysis of each of the eight mechanisms

If ideological fit is the only criterion actors maximize (Figure 1, first diagram), the discursive space ends up with two distinct components after about 1,000 time steps. They correspond to two distinct ideological camps, which are internally perfectly connected. As shown by the blue curve, betweenness centralization increases in the initial rounds of the rewiring process until the discourse is so bipolar that the two camps become disconnected. When this happens, the black curve shows that the number of components is increased from one to two (shown here is a fraction of the number of concepts). The red curve reveals that ideological polarization is maximal, so the two components are each composed of one ideology. The number of concept changes (the green curve) goes down to zero, which means that the network does not change anymore after this stage. As indicated by the orange curve, there are only two concepts left in the discourse (one in each coalition), and all others become extinct. Purely ideology-driven actors are thus fairly unrealistic.

What happens if one replaces exogenous ideology with an endogenous specification of ideology, as plotted in the second diagram? If actors base their decision on the ideological distance between their own history and a new concept under consideration, they always opt for concepts they have named before, leading to path dependency of concept choice. At the aggregate level, this leads to an equilibrium where each actor chooses one single concept for the remaining time, and thus there are almost as many distinct components as there are concepts in the discourse. One concept becomes extinct, all others remain alive. Ideological polarization is fairly high, but the number of distinct camps is unrealistic if compared to real-world discourses and theories of the policy process [9,11].

If all actors pursue a coalition-formation strategy by adopting concepts from other actors with a similar concept profile (third diagram), the result is very similar. In this case, all concepts remain alive, and there are as many components in the discourse after about

Figure 2 Interaction effects. Results of 100 simulation runs. 95% confidence intervals as dashed lines. Ideological polarization (red), share of concepts still alive (orange), share of recent concept changes (green), betweenness centralization (blue), and number of components as a share of the initial number of concepts (black). β = interest group and γ = governmental actor. β_1 and γ_1 = exogenous ideology (XI), β_2 and γ_2 = endogenous ideology (NI), β_3 and γ_3 = actor similarity/coalition formation (AS), β_4 and γ_4 = concept similarity/wisdom of the crowds (CS), β_5 and γ_5 = actor's history/self-consistency (AH), β_6 and γ_6 = concept popularity (CP), β_7 and γ_7 = rare concepts/agenda setting (RC), and β_8 and γ_8 = government coherence (GC).

1,000 time steps. The degree of ideological polarization is realistic, but the component structure is not.

The pattern is very different if actors choose their concepts by drawing on the collective judgment of other actors, as in the concept similarity mechanism plotted in the fourth

diagram. Actors first assess the similarity of any pair of concepts via common referrals by other actors and then select the concept that makes the best match with the actor's own previously named concepts. Such a behavior permits the discourse to be constantly in flux, with about 70% of all new statements deviating from the previously named concept by the same actor (the green curve). The number of concepts between which the actors then choose, however, is extraordinarily small (the orange curve), and ideological polarization is significantly less present than in an Erdős-Rényi random graph, that is, there appears to be a substantial amount of cross-ideology mixing. At all time steps, there is only one component with extremely low polarization.

The fifth diagram shows what happens if actors are consistency-maximizing such that they give highest priority to their own previous concepts. As in the other mechanism leading to path-dependency (the actor similarity mechanism), the number of components becomes maximal because every actor chooses not to deviate from their previous concept anymore.

The sixth diagram shows the effects of the concept popularity mechanism. If actors judge concepts only by their prevalence in the latest round of all actors' histories, that is, recent media coverage, all actors will eventually agree on a single concept, and all other concepts will become extinct after about 500 time steps.

The seventh diagram demonstrates the effect of purely agenda-setting actors. They tend to choose concepts which have been named very infrequently. This is basically the opposite of the concept popularity effect described before. In this case, the degree of innovation is high, and all concepts stay alive. Ideological polarization is not present, and there is only one component at all time steps. The aggregate discourse is largely chaotic.

The eighth diagram shows what happens if all actors maximize coherence between their concept choice and the majority of governmental actors. As with the concept popularity mechanism, all actors end up agreeing on one single concept after few time steps.

All of these situations are clearly unrealistic. They yield equilibria after at most 1,000 time steps. Moreover, in all cases, at least one of the indicators deviates significantly from what one would expect given theoretical characterizations of political discourses. In real-world discourses, there should be a fairly high ideological polarization between approximately two camps. These camps should be loosely connected, and they should exhibit a higher density internally. The population of concepts should not become too small because real-world discourses usually include many different aspects as well. None of the configurations presented so far meets these requirements even approximately. Figure 2 therefore shows some more complex configurations of the utility functions. The phase transition takes considerably more time steps to complete than in the case of single mechanisms. The simulations are run for 10,000 time steps in order to assess the equilibrium behavior of the models.

Analysis of simple interaction effects

An interesting finding from the previous configurations of the utility functions is that actor similarity (β_3 and γ_3) and concept similarity (β_4 and γ_4) have complementary strengths. Each of the two models separately is unrealistic but has certain strengths. The actor similarity model shows a high degree of ideological polarization, and the diversity of concepts remains high. The concept similarity model, on the other hand, shows a high

number of concept changes and a low number of components. It might therefore be a good idea to mix the two mechanisms in order to generate a more realistic aggregate picture of a political discourse.

In the first case, therefore, all actors apply the actor similarity (coalition formation) mechanism and the concept similarity (wisdom of the crowds) mechanism with equal strength. That is, the β and γ weights of the third and fourth components of the utility function are set to 1. The aggregate pattern is already more realistic than any single, separate mechanism (Figure 2, first diagram): ideological polarization is slightly higher than in a random graph, the number of concepts remains fairly high, and the innovative capacity of the discourse is moderate. These are fairly realistic patterns if compared to theoretical expectations of real-world discourses. However, the discourse falls into two separate components or camps, which stay separated for the remaining time steps. This is clearly unrealistic, given the large amount of overlap between the advocacy coalitions observed in empirical case studies [2,40]. Moreover, ideological polarization is only slightly higher than pure chance would predict, so some kind of exogenous ideology seems to be necessary in the model.

The second model assumes that all actors maximize ideological fit as well as coherence with the line of the government, and exhibit agenda-setting behavior. Agenda setting is weighted twice as high as each of the other two components of the utility function. In other words, the innovative potential of each actor is exactly as high as governmental coherence and exogenous ideology together. The scenario deviates from the expectations insofar as a very high number of distinct ideological camps emerges.

Beside government coherence, there is no mechanism in the model that would provide incentives for actors to adopt concepts from other coalitions ('learning across coalitions'). The third diagram therefore incorporates such an element into the utility function. The configuration is extended by adding the concept similarity mechanism to the equation. The result is as expected: the number of distinct components is significantly lower than before, and the number of concept changes is much higher. On the other hand, there are approximately two components at the end of the 10,000 time steps (showing a relatively high variance), with a gradual upward trend. It would be better if the model generated one single component that sometimes fell into two separate components and then merged them again, rather than a situation where there are usually two separate components which are sometimes merged and then dissolve again, given the literature on policy brokers that tie together competing factions in policy domains [40]. Moreover, all models presented so far have not distinguished interest groups from governmental actors.

The fourth model therefore makes some simple and presumably fairly straightforward assumptions about the political process by giving governmental actors and interest groups two different utility functions. In this model, interest groups maximize ideological fit and governmental coherence. This is presumably close to empirically observable discourses for two reasons: first, because interest groups have different interests - modeled here as two extremes on a one-dimensional scale; second, because interest groups always try to convince decision-makers of their favored policy image [28], which implies that they reiterate the ideas of governmental actors if they match their ideology. Repeating favorable ideas of governmental actors serves the purpose of supporting decision-makers when they happen to call for the 'right' policy instruments in their perspective. For these reasons, combining ideology and governmental coherence in the utility function of interest

groups is a plausible and simple assumption. Governmental actors, on the other hand, pursue a threefold strategy. Of course, they opt for governmental coherence (γ_8). As Heclo asserts, politicians also 'puzzle' about what solution concepts are desirable because their environment is complex and they are no experts [6]. For this reason, they show a tendency to adopt popular claims in the debate (γ_6). Finally, politicians have an interest in determining the politics of the day and presenting innovative solution concepts that are currently not on the political agenda [30]. They adopt rare concepts with the same strength as they go for government coherence and concept popularity together (γ_7). Such a configuration of the utility function yields exactly one component at all time steps, with moderate polarization and a slowly decreasing number of concepts. This configuration of the discourse is already close to the theoretical expectations and empirical observations generated by the literature on advocacy coalitions [9,10], but the separation of the coalitions should be clearer than implied by the model presented here. Especially governmental actors are modeled as largely chaotic beings who do not have any preconceptions about what is right or wrong in the first place. This erratic trend-hunting behavior is apparently the main reason why the separation between the ideological camps is so low.

The fifth model therefore changes the utility function of governmental actors and adds exogenous ideology to the calculus of politicians (γ_1). To keep up a good balance between continuity and change, it is always necessary to balance path-dependent mechanisms like ideology or self-consistency maximization, and innovative mechanisms like agenda-setting behavior, so a weight of 3 is attached to agenda setting (γ_7). In this configuration, ideological polarization is indeed stronger and more stable over time. However, the number of concepts is still slowly decreasing over time, and betweenness centralization is rather low. The reason may be that interest groups do not have any innovative potential at all in this model. They merely follow the majority of governmental actors and their own convictions.
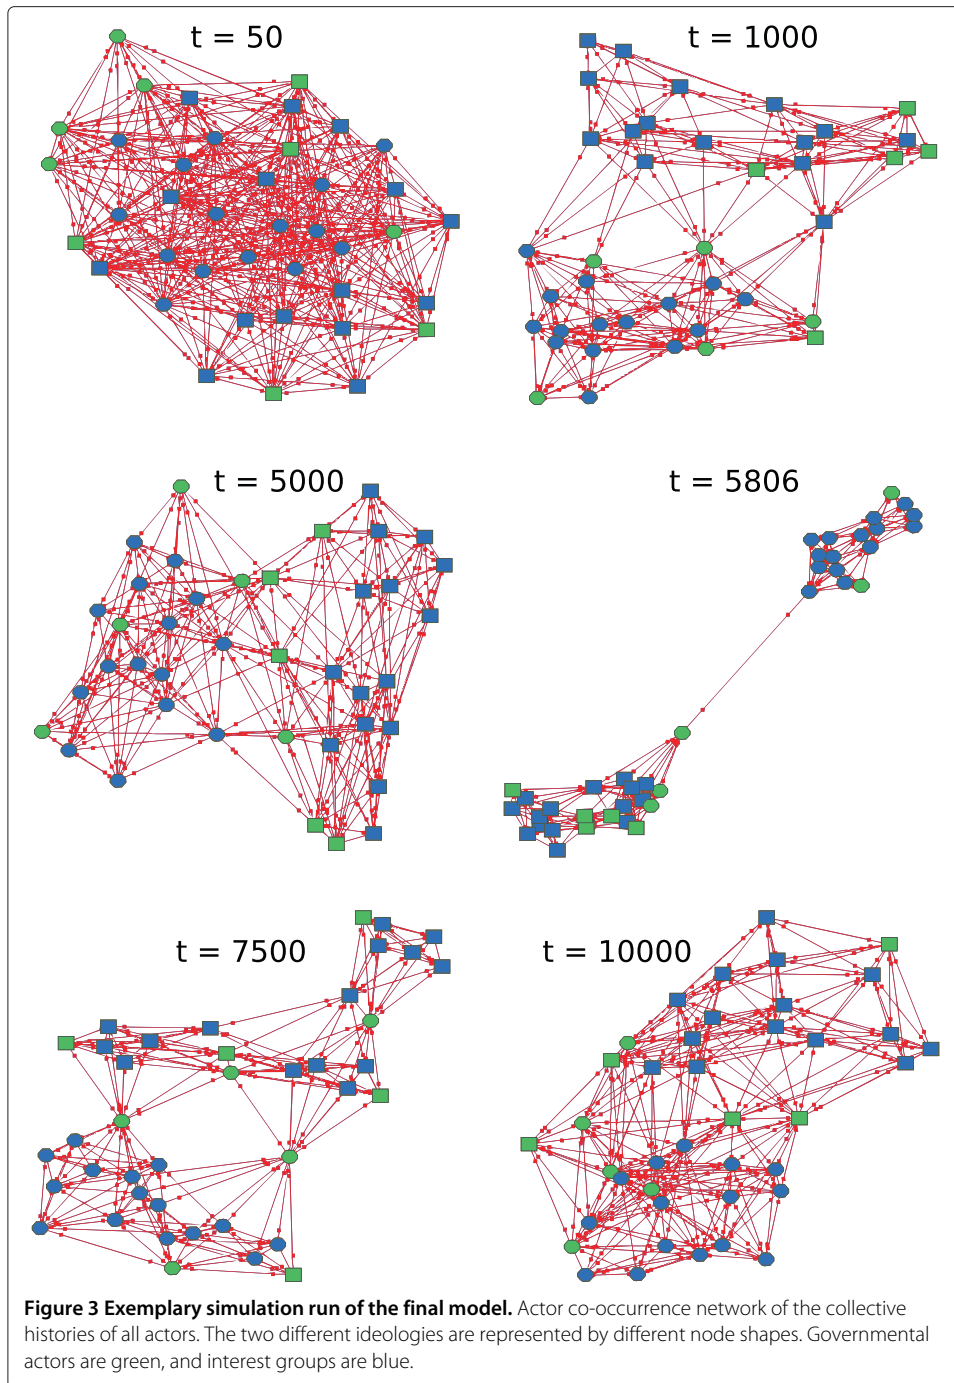
The final model therefore takes into account that interest groups may have higher payoffs from selecting concepts with lower prevalence in the discourse. The agenda-setting behavior of interest groups makes up one third of their utility (β_7). The result is strikingly realistic when being compared to the theoretical expectations raised above as well as recent empirical evidence on real-world policy discourses [1,2,35]: there is usually one component, which shows substantial ideological polarization. At the same time, ideological polarization is not maximal; there are still many ties in the network that allow for cross-coalition mixing, effectively binding the two ideological camps together in one single component. Every now and then, the discourse becomes so polarized that the overlap between the coalitions disappears. During these stages, there is no common ground for discussion anymore. Soon after, however, the camps engage in mutual deliberation again and establish some common ground for potential consensus. This is shown by the black line, which is usually showing only one component, but which sometimes slightly deviates upwards. The upward deviations are apparently so small because the line shows the average share of components of all 100 simulation runs, but the deviations are still visible. Betweenness centralization is moderate at all time steps, with substantial variation in terms of the confidence intervals. On the one hand, this shows that two clusters can be indeed clearly identified within the single component. On the other hand, the variation shows that the discourse is still far from being in equilibrium in the

sense that every actor would reiterate the same concepts over and over again, as in the first eight simple configurations shown in the previous figures. This claim is supported by the green curve, which demonstrates that the number of recent concept changes is significantly above zero. The population of concepts is diverse and 'healthy' at all time steps. It may happen that one or two concepts disappear from the discourse, but the relatively strong agenda-setting component always reinvigorates them after a couple of time steps. Finally, even though the discourse is out of equilibrium in the sense that actors still come up with new concepts and the discussion is vivid and ongoing, the observable network statistics time series become stationary after the initial rewiring process. This means that the structure of the discourse is very likely to remain approximately stable as shown in the sixth diagram for an unlimited time. The discourse is constantly in flux, but the basic bipolar structure is stationary, with some random fluctuations over time. This final model therefore seems to capture the essence of political discourse relatively well.

Figure 3 validates the claims about the structure of the discourse in the final model by visualizing an exemplary simulation run. Different node shapes (circles and squares) reflect the two different ideologies. Governmental actors are represented by green nodes and interest groups by blue vertices. In the first panel, the actor co-occurrence network is in the early stages of the rewiring process from a random graph to a discourse network. There is no typical discourse structure yet. After 1,000 time steps, the rewiring process has been completed. There are two coalitions corresponding to the two ideologies, with several connections between the two factions. Governmental actors are particularly likely to cluster together and connect the two camps. This is also visible after 5,000 time steps. At time step 5,806, the two coalitions almost lose their last bit of overlap. Only one governmental actor serves as a bridge between the two coalitions. The polarization slightly decreases again soon after. After 7,500 time steps, one can see that discourse does not necessarily have to be one-dimensional. As can be observed in empirical discourse networks [2], coalitions can sometimes decompose into two or more subcoalitions with different aims and beliefs. The rectangles at time step 7,500 make up two or three subclusters within the rectangular advocacy coalition. Moreover, the circles are connected to the rectangles via two different brokers. It is also noteworthy that actors from one coalition may sometimes join the opponent and then come back - a feature of the discourse that is also occasionally observable in empirical discourse networks [2]. Finally, the situation goes back to a somewhat polarized normal state with two relatively homogenous coalitions again after 10,000 time steps.

Conclusions

This article has presented a family of agent-based computational discourse network models. The most striking finding is that simple interactions of some of the proposed mechanisms yield discursive structures that are in line with what is expected theoretically and empirically [9-11,40], while none of the basic mechanisms separately has implications that seem to correspond to the real world. In particular, ideology, concept popularity, agenda setting, and government coherence each produce unrealistic discourse networks. If combined in a single utility function of agents, and if one distinguishes between the roles of interest groups and governmental actors, however, these mechanisms lead to plausible discursive structures.

Figure 3 Exemplary simulation run of the final model. Actor co-occurrence network of the collective histories of all actors. The two different ideologies are represented by different node shapes. Governmental actors are green, and interest groups are blue.

Apparently, neither purely exogenous preferences nor purely endogenous opinion formation can explain the structure of real-world policy debates. Only a combination of the two paradigms is fruitful in modeling political discourse. This finding may help to build a bridge between rational choice and constructivism, which have been subject to extensive controversies.

Moreover, only a simultaneous presence of both innovative and path-dependent mechanisms ensures that a simulated discourse lives up to the implications of its real-world

counterparts. This finding suggests that even discourses which underlie 'normal' states of policy making are guided by a balance of stochastic elements and the prosecution of self-interest.

The parameter settings tested here have been arrived at in an explorative fashion. There may be other combinations of mechanisms that produce similar or even more realistic results. Future research could explore this possibility by employing a combinatorial optimization algorithm such as simulated annealing to find optimal combinations of parameter settings in order to find the simplest relatively accurate model. Employing a combinatorial search heuristic would require that all relevant metrics like ideological polarization etc. can be combined in a single goodness-of-fit statistic which can then be optimized by the algorithm, which is currently not the case. Beside these potential improvements in the analysis stage, there are several assumptions which could be modified in future work to get a more nuanced picture of specific aspects of political discourse[e]. Future research should also aim to test specific mechanisms or interactions outlined here in experimental settings or using empirical data to further our understanding of the social processes underlying political discourse. One promising avenue for this are the recent advances in relational event modeling [41].

Endnotes

[a] The replication source code of the model is provided as Additional File 1.

[b] If the ideological scores are not fixed at the extreme poles, one of the following consequences hold: the simulation runs exhibit the same results but take considerably longer to show a recurring pattern, or multiple 'coalitions' emerge instead of just two coalitions (hence the discourse becomes multipolar instead of bipolar). The parameter choice is made because the majority of real-world discourses appear to be bipolar rather than multipolar.

[c] http://www.jgrapht.org (as of October 9, 2013).

[d] http://repast.sourceforge.net/ (as of October 9, 2013).

[e] In particular, the following assumptions could be modified to make the model more realistic (but less parsimonious): other actors like scientists, opposition, voters, or the media do not play a role; actors may have different resource endowments and therefore different skill levels regarding the accuracy of their information; the number of actors stays constant over the whole discourse; a discourse does not interfere with other topical discourses; the number of concepts is constant; actors are not allowed to hold contradictory positions regarding the same normative concept; new concepts are never introduced to the discourse; external perturbations [9] do not exist; exactly one statement is made at every new time step; actors observe all other actors' statements, that is, there is complete and perfect information; actors' ideological positions are extreme (instead of being, for example, normally distributed with two modes); the eight mechanisms in the utility function are exhaustive; and the probability that a governmental actor makes a statement is 60%.

Additional file

Additional file 1: Compressed ZIP archive. This file contains the Java source code of the simulation model and R code for the replication of the statistical analysis.

References
1.   Leifeld, P, Haunss, S: Political discourse networks and the conflict over software patents in Europe. Eur. J. Pol. Res. **53**(3), 382–409 (2012). doi:10.1111/j.1475-6765.2011.02003.x
2.   Leifeld, P: Reconceptualizing major policy change in the advocacy coalition framework. a discourse network analysis of German pension politics. Policy Stud. J. **41**(1), 169–198 (2013). doi:10.1111/psj.12007
3.   Brandes, U, Erlebach, T(eds.): Network Analysis: Methodological Foundations. Springer, Berlin (2005)
4.   Wasserman, S, Faust, K: Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge (1994)
5.   Ward, M. D, Stovel, K, Sacks, A: Application of network analysis to political problems. Ann. Rev. Pol. Sci. **14**(1), 245–264 (2011)
6.   Heclo, H: Issue networks and the executive establishment. In: King, A (ed.) The New American Political System, pp. 87–124. American Enterprise Institute, Washington D.C. (1978)
7.   Fairclough, N: Language and Power. Longman, London (1989)
8.   Pierson, P: Increasing returns, path dependence, and the study of politics. Am. Pol. Sci. Rev. **94**(2), 251–267 (2000)
9.   Sabatier, P. A: An advocacy coalition framework of policy change and the role of policy-oriented learning therein. Policy Sciences **21**(2), 129–168 (1988)
10. Ingold, K. M: Network structures within policy processes: coalitions, power, and brokerage in Swiss climate policy. Policy Stud. J. **39**(3), 435–459 (2011)
11. Hajer, MA: The Politics of Environmental Discourse: Ecological Modernization and the Policy Process. Oxford University Press, Oxford (1995)
12. Marvel, SA, Kleinberg, J, Kleinberg, RD, Strogatz, SH: Continuous-time model of structural balance. Proc. Natl. Acad. Sci. **108**(5), 1771–1776 (2011)
13. Traag, VA, van Dooren, P, de Leenheer, P: Dynamical models explaining social balance and evolution of cooperation. PLOSone **8**(4), 60063 (2013). doi:10.1371/journal.pone.0060063
14. Miller, D: Deliberative democracy and social choice. Pol. Stud. **40**(s1), 54–67 (2007)
15. Epstein, JM: Agent-based computational models and generative social science. Complexity **4**(5), 41–60 (1999)
16. Baldassarri, D, Bearman, P: Dynamics of political polarization. Am. Sociol. Rev. **72**(5), 784–811 (2007)
17. Bhavnani, R, Findley, MG, Kuklinski, JH: Rumor dynamics in ethnic violence. J. Pol. **71**(03), 876–892 (2009)
18. Bray, D, Shackley, S: The social simulation of the public perception of weather events and their effect upon the development of belief in anthropogenic climate change. Tyndall Centre Working Paper 58, Tyndall Centre for Climate Change Research, Norwich (2004)
19. Henry, AD: Simulating the evolution of policy-relevant beliefs: can rational learning lead to advocacy coalitions?. In: Ostrom, E, Schlüter, A (eds.) The Challenge of Self-Governance in Complex, Globalizing Economies. Collection of revised papers of a PhD seminar; 17[th] to the 26[th] of April 2007 in Freiburg. Arbeitsbericht, pp. 135–160. Institut für Forstökonomie, Universität Freiburg, Freiburg, (2007)
20. Lustick, I, Miodownik, D: Deliberative democracy and public discourse: the agent-based argument repertoire model. Complexity **5**(4), 13–30 (2000)
21. Cederman, LE: Agent-based modeling in political science. Pol. Methodologist **10**(1), 16–22 (2001)
22. Johnson, PE: Simulation modeling in political science. Am. Behav. Scientist **42**(10), 1509–1530 (1999)
23. Krackhardt, D, Stern, RN: Informal networks and organizational crises: an experimental simulation. Soc. Psychol. Q. **51**(2), 123–140 (1988)
24. Hinich, MJ, Munger, MC: Analytical Politics. Cambridge University Press, Cambridge (1997)
25. Wolfsfeld, G, Sheafer, T: Competing actors and the construction of political news: the contest over waves in Israel. Pol. Commun. **23**(3), 333–354 (2006)
26. Downs, A: Up and down with ecology: the issue attention cycle. Public Interest **28**(1), 38–50 (1972)
27. Barabási, AL, Albert, R: Emergence of scaling in random networks. Science **286**, 509–512 (1999)
28. Baumgartner, FR, Jones, BD: Agenda dynamics and policy subsystems. J. Pol. **53**(4), 1044–1074 (1991)
29. Aronson, E: Dissonance theory: progress and problems. In: Abelson, RP, Aronson, E, McGuire, WJ, Newcomb, TM, Rosenberg, MJ, Tannenbaum, PH (eds.) Theories of Cognitive Consistency. A Sourcebook, pp. 5–27. Rand McNally, Chicago, (1968)
30. Hall, PA: Policy paradigms, social learning, and the state: the case of economic policymaking in Britain. Comparative Pol. **25**(3), 275–296 (1993)
31. Abelson, RP, Aronson, E, McGuire, WJ, Newcomb, TM, Rosenberg, MJ, Tannenbaum, PH: Theories of Cognitive Consistency. A Sourcebook. Rand McNally, Chicago (1968)
32. Festinger, L: A Theory of Cognitive Dissonance. Stanford University Press, Redwood City (1957)
33. Cialdini, RB, Trost, MR, Newsom, JT: Preference for consistency: the development of a valid measure and the discovery of surprising behavioral implications. J. Pers. Soc. Psychol. **69**(2), 318–328 (1995)
34. Lijphart, A: Patterns of Democracy. Government Forms and Performance in Thirty-Six Countries. Yale University Press, London (1999)
35. Fisher, DR, Leifeld, P, Iwaki, Y: Mapping the ideological networks of American climate politics. Climatic Change **116**(1), 523–545 (2013). doi:10.1007/s10584-012-0512-7
36. Fisher, DR, Waggle, J, Leifeld, P: Where does political polarization come from? locating polarization within the U.S. climate change debate. Am. Behav. Sci. **57**(1), 70–92 (2013). doi:10.1177/0002764212463360
37. Freeman, LC: Centrality in social networks: conceptual clarification. Soc. Netw. **1**(3), 215–239 (1979)

38. Newman, MEJ: Mixing patterns in networks. Phys. Rev. E **67**(2), 026126 (2003)
39. Watts, DJ, Strogatz, SH: Collective dynamics of "small-world" networks. Nature **393**(6684), 440–442 (1998)
40. Ingold, KM, Varone, F: Treating policy brokers seriously: evidence from the climate policy. J. Public Adm. Res. Theory **22**(2), 319–346 (2012)
41. Lerner, J, Bussmann, M, Snijders, TAB, Brandes, U: Modeling frequency and type of interaction in event networks. Corvinus J. Sociol. Soc. Policy. **4**(1), 3–32 (2013)