D4.3 Populism and Opinion Dynamics

WP4 – Causal, policy and futures analysis

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 822337.

Populism and Civic Engagement

About this document

Work Package 4 Lead: The Centre for Liberal Strategies Foundation (CLS)

Task Leads: Bruce Edmonds and Ruth Meyer (MMU)

Document history

Version	Date	Description	Reason for change	Distribution	Authors
V 1	May 2022	Draft of Populism and Opinion Dynamics			Bruce Edmonds (MMU)

Due date	April 30, 2022
Delivery date	Submitted May 2022.
Type	Research Results
Dissemination level	Report, Public

Disclaimer

The project PaCE has received funding from the European Union´s Horizon 2020 research and innovation programme under the grant agreement No 822337.

Contents

Figures

Introduction

One of the more innovative (and hence higher-risk) approaches to analysis undertaken in PaCE is the agent-based modelling (ABM) of the co-evolution of voter behaviour (social communication, adaption of views, voting choices) and party behaviour (in terms of its declared policy positions). ABM is unique in its ability to formally represent the relationship between micro- and macro-levels of evidence (as illustrated below in Figure 1) and thus can provide some unique insights into political phenomena. Apart from informal, discursive, accounts other approaches to political phenomena focus on either the micro or macro levels.

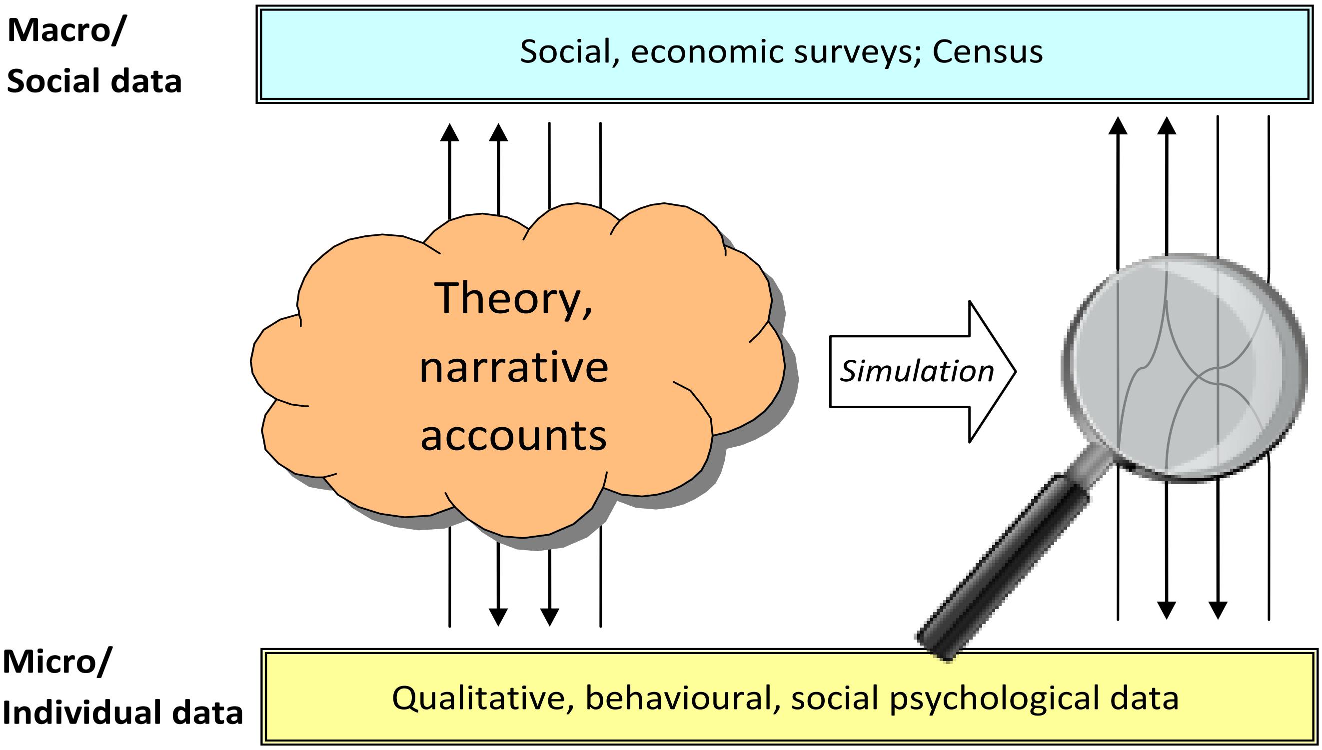

Figure 1. The role of agent-based simulation in making micro-macro linkages explicit

In this report we describe how an evidence-led ABM of political behaviour in Austria (2013-2017) can be used to explore how different micro-level process that contribute to the overall macro-outcomes, might combine – something very difficult to do in any other way. Thus the purpose of this exercise is to explore the effects of different assumptions/mixtures on the outcomes (what is called "theoretical exploration" in (Edmonds et al. 2019). The insights gained from these simulation experiments do not tell us directly about observed political phenomena (this is *in vitro* rather than *in vivo* research) but might suggest hypotheses that could later be empirically investigated.

In particular, to answer some of the following questions concerning this model:

- How important are the social influence processes in producing plausible outcomes?
- How important are the network dynamics in producing plausible outcomes?
- How important are party dynamic behaviour in producing plausible outcomes?
- How sensitive is the model to its various settings and parameters?
- What impact do the settings have on the convergence of voter attitudes and how 'effective' is a resulting (notional) government in reflecting these attitudes?
- Are there any insights from this analysis that relate to populism?
- Does this suggest any areas for further empirical research?

Brief Summary of the Simulation

Here we give a surface-level introduction to the simulation, for the convenience of readers. For the full details about the simulation see deliverable: ***D2.5 "Lessons learned from the simulation analysis"***.

An agent-based model (ABM) simulates the actions and interactions of autonomous agents, which can represent both individuals (e.g., voters) or collective entities (e.g., parties), in order to understand the behaviour of the system under investigation. Specific advantages of this approach are the capability to explicitly model individual behaviour, include heterogeneous decision-making processes, and integrate quantitative and qualitative data from a range of sources. It also allows for exploring what-if scenarios and counterfactual reasoning. The PaCE project (WP 2) uses real world data on voters' attitudes and party preference, on the position of political parties, and the external salience of issues in the mass public and combines this with theories on voters' decision-making (Lau et al. 2018) and parties' strategic moves in the political space (Muis and Scholte 2013, Laver and Sergenti 2012).

In this model, we simulate the development in Austrian party politics between the national elections of 2013 and 2017, a period that was affected by the refugee crisis of 2015/2016 and the above-mentioned leadership-change in and shift to the right by the conservative ÖVP (as described in PaCE reports D2.3 and D2.5). The best available model uses a mix of voter strategies and successfully reproduces the trends in opinion polls, namely the rise of the FPO. Subsequent simulations will not only use what-if scenarios based on different voter strategies (such as the relative importance of rational choice) but will focus on what-if scenarios concerning real-world developments such as the refugee crisis and the leadership and programmatic change of the ÖVP. This will be compared to the development in Germany, where mainstream parties are confronted with populism not only on the right (AfD) but alsso on the left of the political spectrum (Linke). The purpose was to better understand the dynamics of the rise of the AfD there, We will not give the full simulation details again here, but quickly summarise to provide some context.

The simulation modelled the interaction between voters and the parties both situated within a 7-diminesional space of attitudes to 7 key issues (economy, welfare state, budge, immigration, environment, society and law&order). Voters could change their opinion on these attitudes due to social influence, and could use a variety of strategies to decide which party to vote for (rational choice, confirmatory, fast&frugal, heuristics-based, go with gut and identity-based). Simultaneously the parties (SPÖ, ÖVP, FPÖ, Grüne, NEOS, BZÖ, Team Stronach) could use a variety of strategies ('Aggregator', 'Satisficer', 'Hunter' and 'Sticker') to use in order to decide their policies.

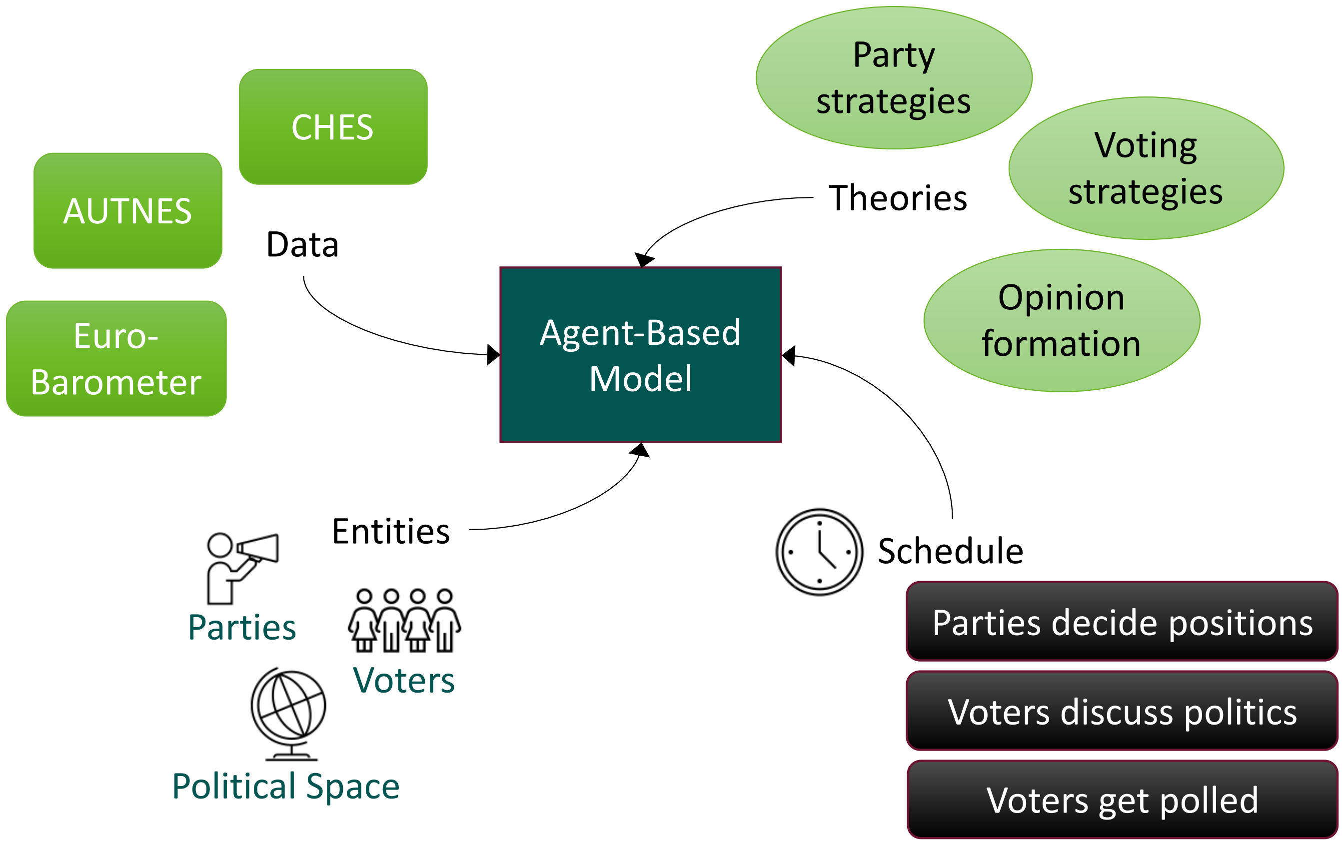

Figure 2. An illustration of the main simulation elements

As far as possible all these settings (and the others) were based upon available evidence. The voters were created using characteristics from the AUTNES data set, the parties using the CHES data set.

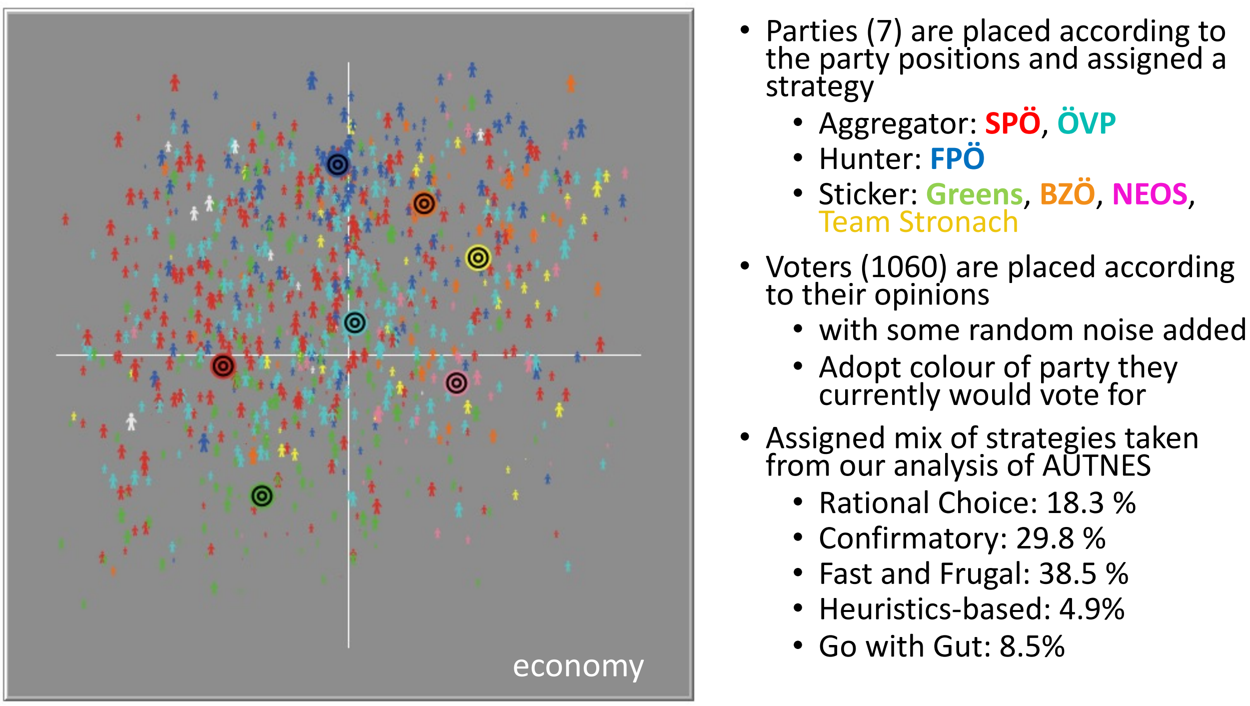

Figure 3. An illustration of two of the dimensions in the space of attitudes with parties and position of parties shown

Then the simulation was explored to see how close to the observed results it could get. These target results were the polling for these parties in the 2013-2017 period.

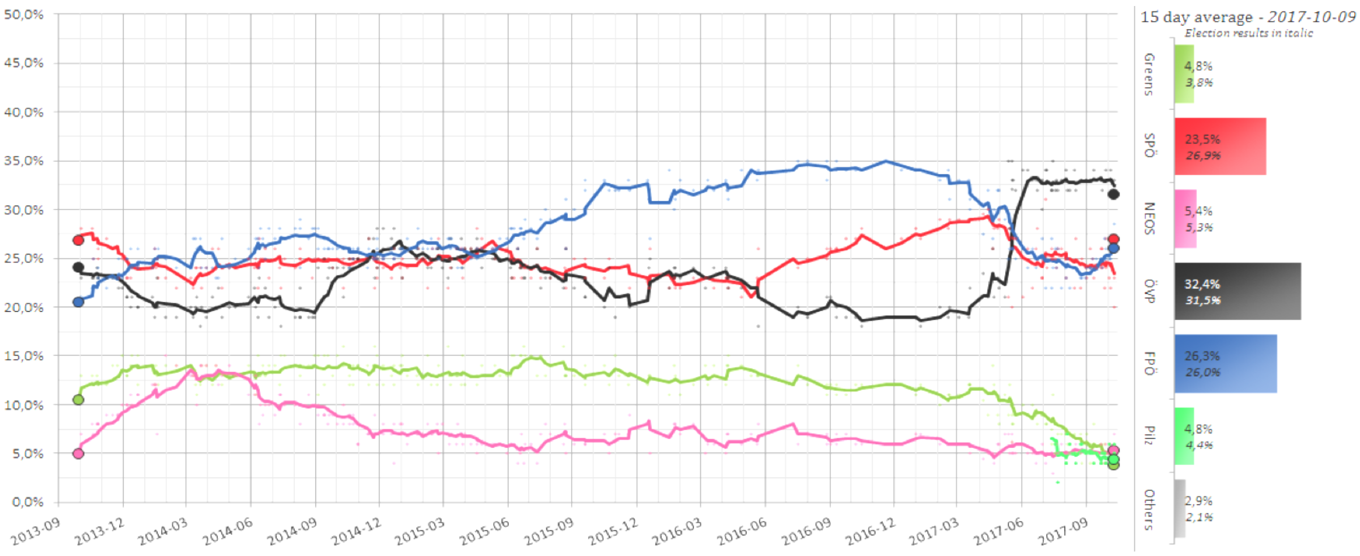

Figure 4. The target polling data, Austria 2013-2017

The best model results obtained looked like that below, which are quite a good match to those above. The essence of its plausibility is in the gradual decline in SPO and Greens, sharp competition between OVP and FPO and low levels for the other parties.

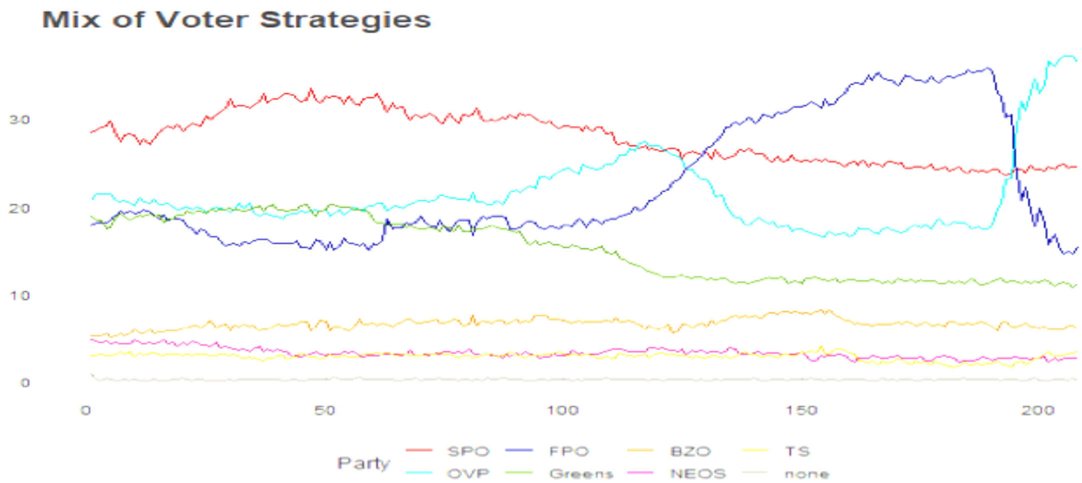

Figure 5. The simulation-generated polling proportions from the best results of the model.

However, when you run the simulation many times using the same parameters and data, you get a *variety* of different patterns, only sometimes matching the target polling data. The diagram below (Figure 6) gives an idea of the variety. As described elsewhere (D2.4) this seems to occur due to the FPO party agent's opportunism in terms of moderating its immigration policy to gain votes from the convergence of immigration-concerned but broadly centrist voters.

Figure 6. The polling proportions generated by the simulation in nine different simulation runs, illustrating the variation of outcomes with the same initial conditions.

In this report we start from this simulation, using these settings and then explore various of the options in terms of the interaction between voters and the parties – going further than the empirically-led case to experiment with the options and settings of the simulation. This allows us to better understand what is happening in the simulation, and hence the nature of any explanation provided by the empirical case.

What is essential to producing ‘plausible’ outcomes?

In this section we report on some simulation experiments that show which of the component processes are necessary in order to get the above ‘plausible’ outcomes. For this we need a base-line (set of settings) for comparison. Given the variation displayed in simulation runs with the same initial conditions (Figure 6) this can only be convincingly shown by showing aggregate results over a set of runs (which we look at in the next section). However, such aggregate summaries do not give one a ‘feel’ for the qualitative differences in each case, so in this section we give a single example run in each case, showing the whole simulation interface at the end of the run in each case.

The Reference Case

The base case is as illustrated below (Figure 7).

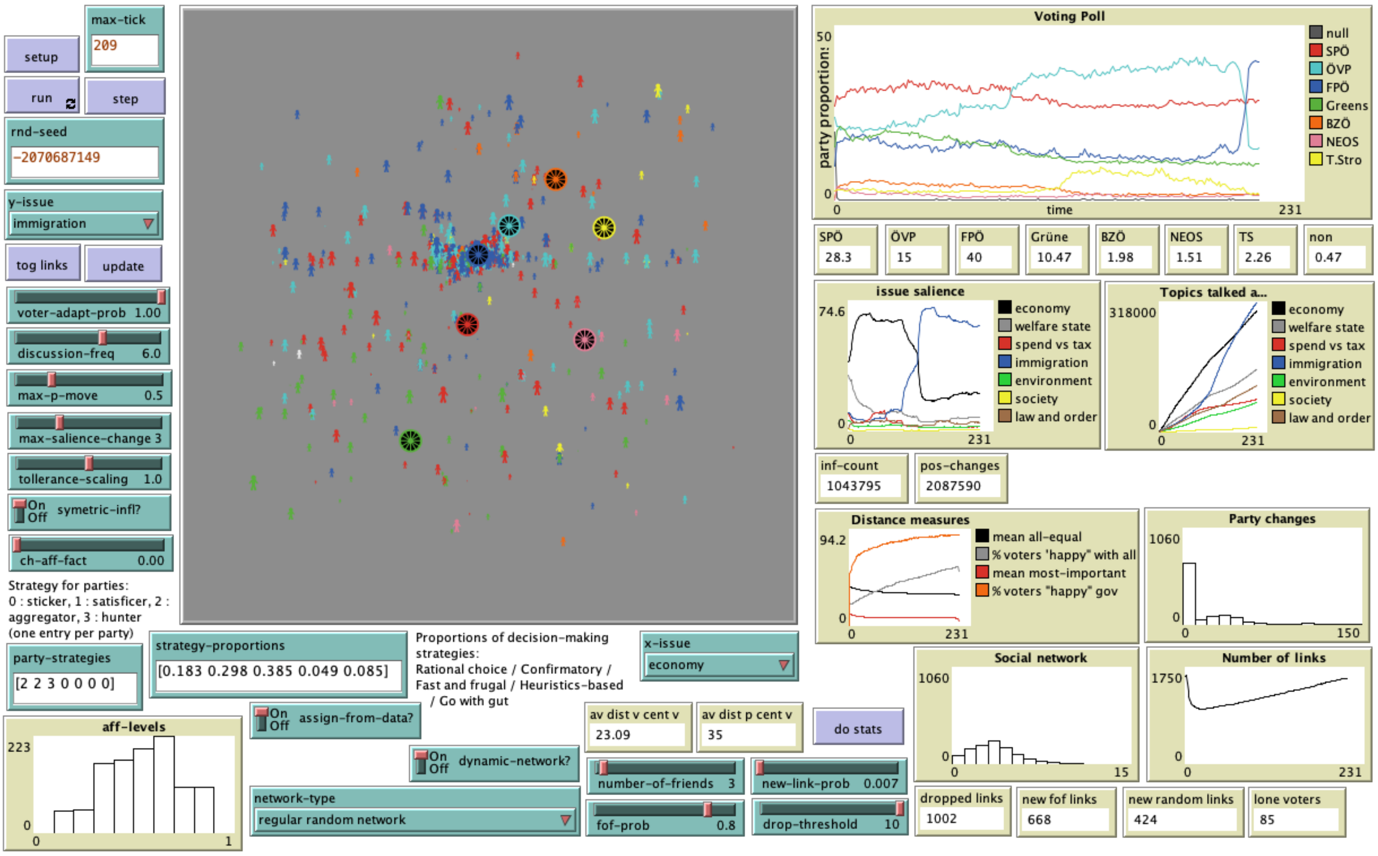

Figure 7. An illustration of a standard reference run

Some things to note about this.

The main world-view (the grey square) shows the positions of the voters (small people) and the parties (the wheels) at the end of the run (representing the 2017 election), with the axes showing the two key dimensions immigration (y-axis) and the economy (x-axis). The colours of the voter agents indicate which party they voted for.

The top-right graph shows the simulated polling data for the support for the parties, with the pattern commented on above.

The two graphs immediately below this show: (left) the voter agent salience of issues for the voters with immigration overtaking the economy as the main item of concern and (right) the cumulative number of topics ‘talked about’ between agents. Note how immigration becomes the most talked about and thus salient issue for voters – allowing the FPO to gain most votes if it moderates its position on immigration.

Below those is a graph showing the ‘satisfaction’ of voters measured in four different ways (the way we will concentrate on is the orange one – the proportion of voters that are near the policies of a notional elected government that might be elected on the basis of these polls). To its right is a histogram showing how frequently voters change who to vote for (most do not change).

Below that are two about the structure of the social network: (left) a histogram showing the how many links agents have and (right) how the number of links in the network changes over time.

Varying The Social Influence Process

When agents communicate with each other (over the social network) they may change each other's attitudes so their opinions are closer (the rate controlled by the parameter: *voter-adapt-prob*) as well as how important the issues are to them (the rate controlled by the parameter: *max-salience-change*) but only if all their opinions are sufficiently close. To see the effects of these we compare the cases when these two processes are switched off (in turn).

First a typical run with no social interaction between voter agents (below).

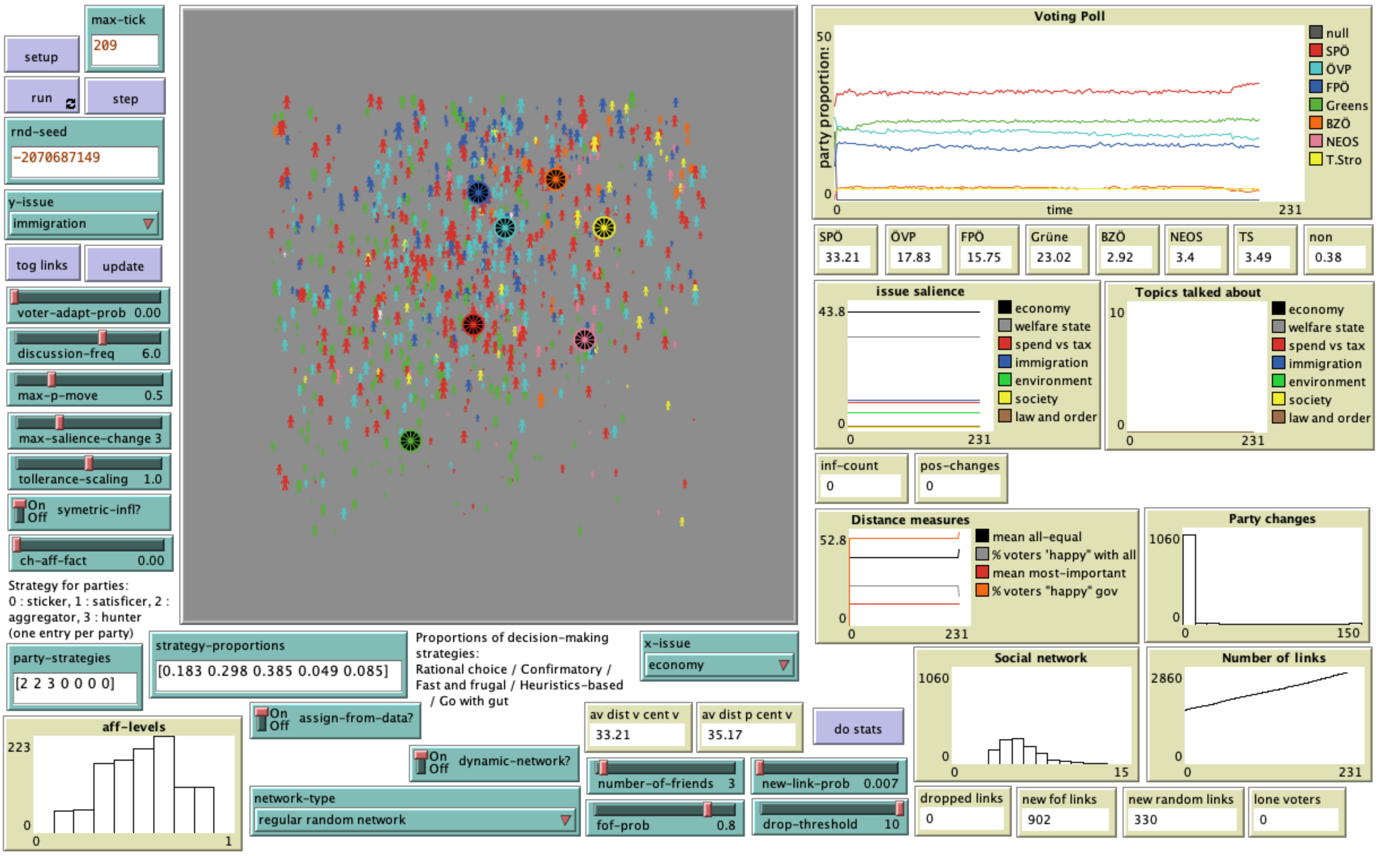

Figure 8. A 'no social interaction' example run

Unsurprisingly, there are not much dynamical behaviour displayed in the model, not at all like the reference case. Also the proportion of voter agents 'happy' with the final election is low (just over 50%). Due to a lack of interaction, no links are dropped (due to this only happening when a certain threshold of interactions that fail to influence have occurred).

Compare this to the case where voter agents do influence each other's attitudes but not the voter agent salience of issues for them.

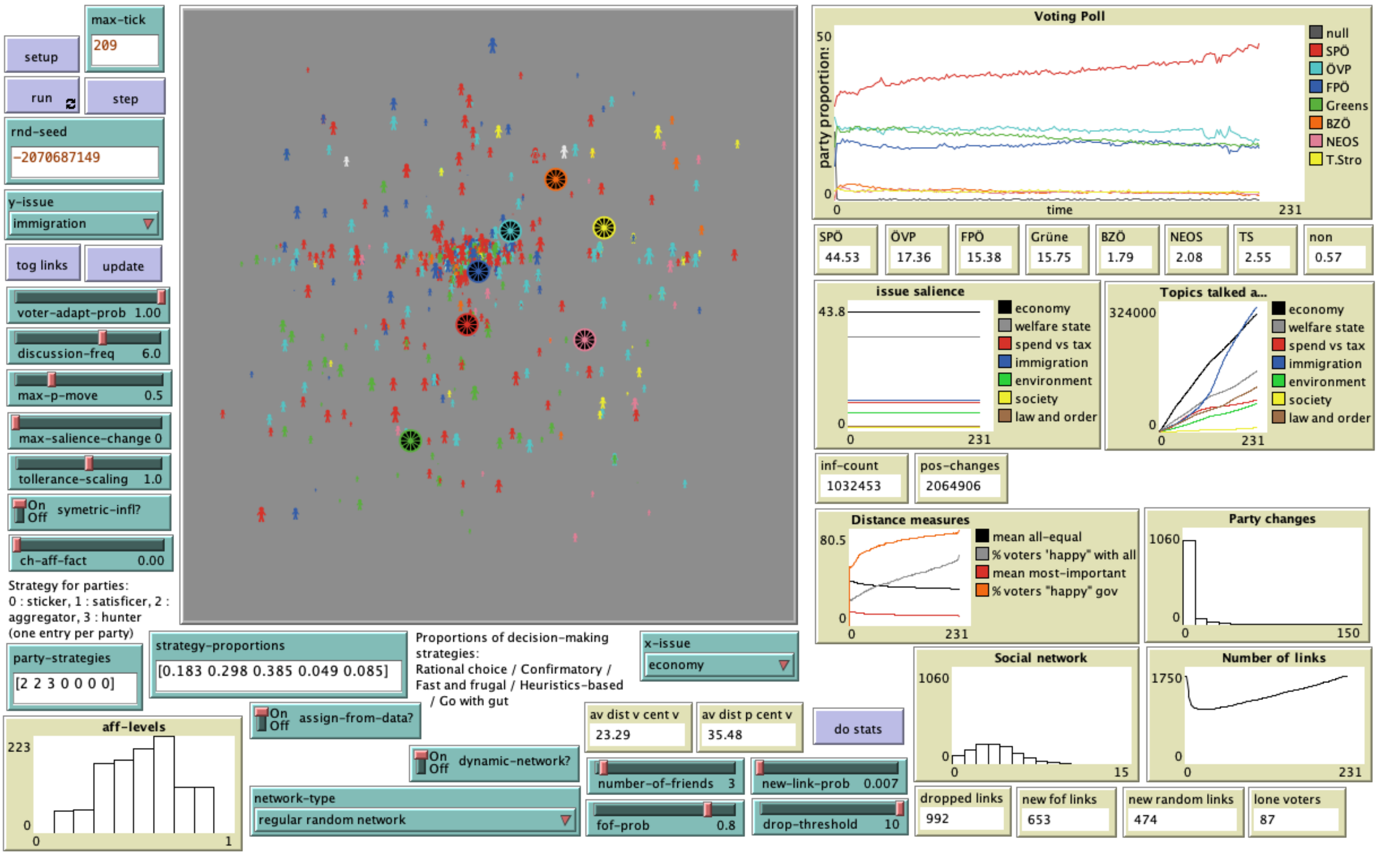

Figure 9. A no salience change example run

Here we do have a convergence of attitudes, but voter agent saliences do not change (so that the salience of the immigration issue remains constant) with the result that the SPO clearly gets most votes (it is central on the economy and many other issues).

We now look at some other ways we might change the social influence process. The first is by reducing the overall 'tolerance' of voter agents (by changing the scaling parameter, *tollerance-scaling*). This means that they only interact with others with a very similar view to their own. However a side effect of this is that there are significantly fewer interactions over all, so we also increase the frequency of discussion (*discussion-freq*) so that there are roughly the same total number of interactions as in the standard case.

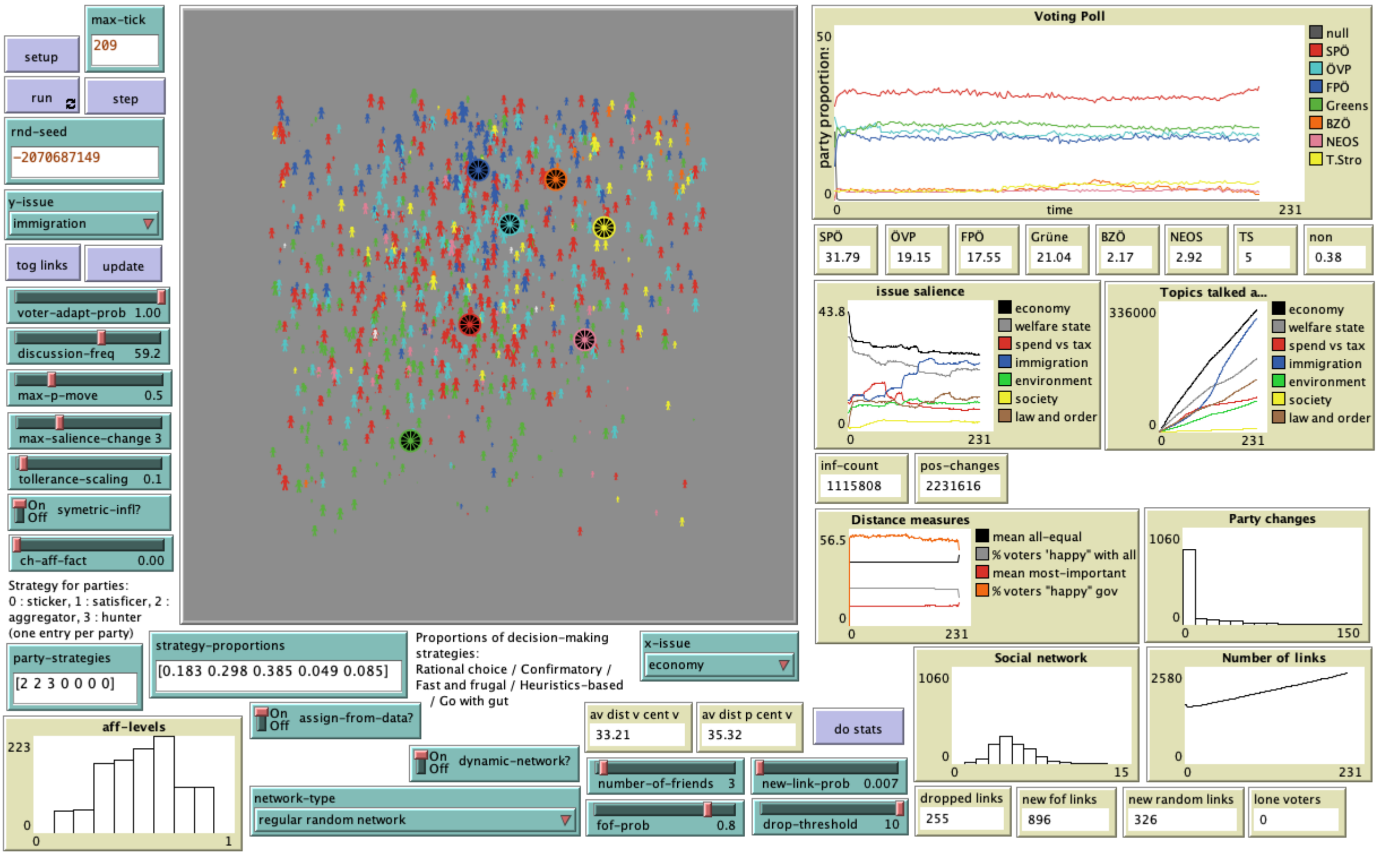

Figure 10. Only interaction with those with similar attitudes

The fact that agents only 'talk' to those who are very similar to themselves means there is little overall convergence of attitudes, even when voter agent saliences and topics 'talked' about do.

The second way, is to make it so that those that are more involved with politics generally are more certain of their opinions and hence adapt these less than those less interested using the *ch-aff-fact* parameter (in a manner similar to the Bounded Confidence Model of Deffuant, 2002).

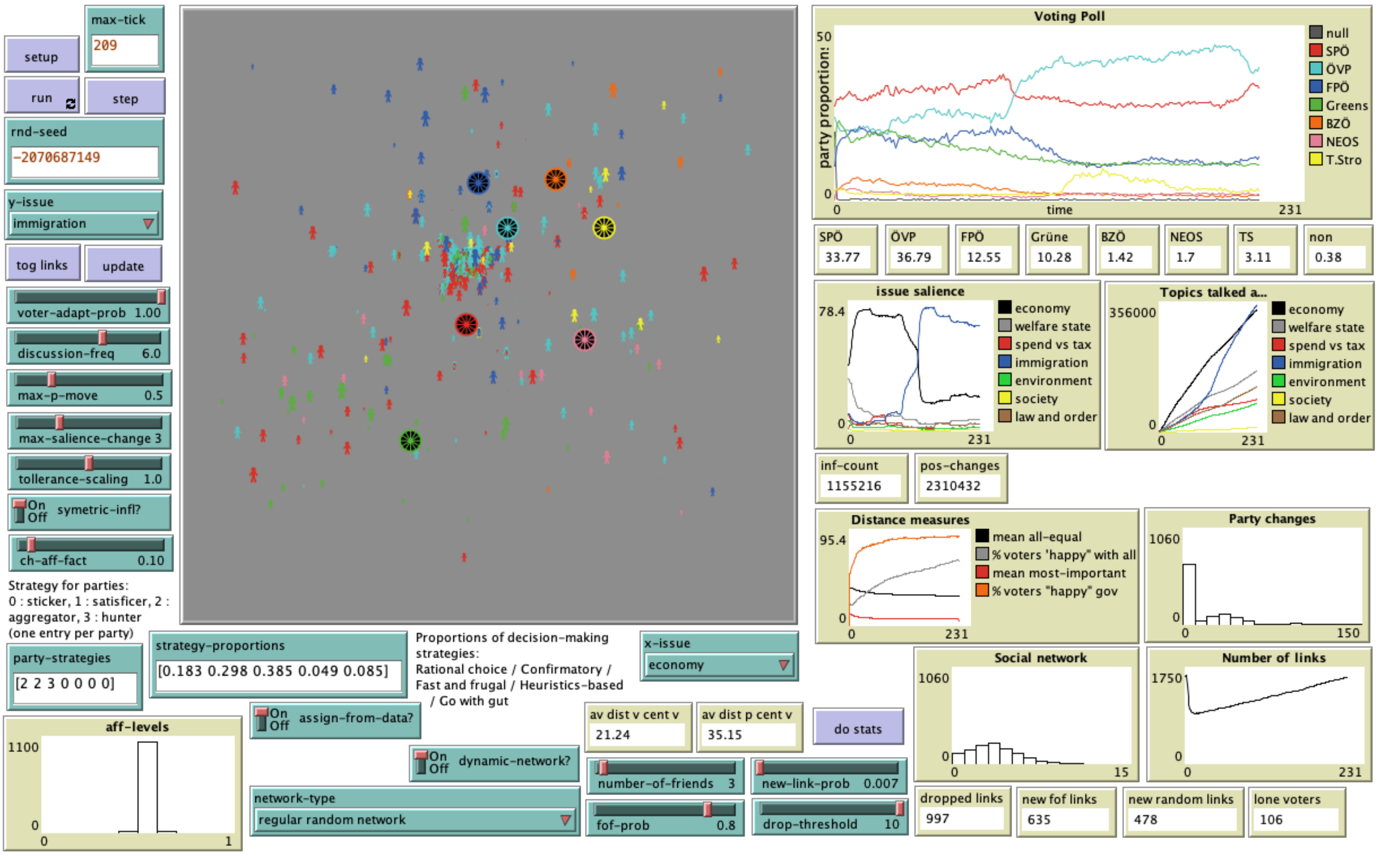

Figure 11. Adding a bias so those more involved in politics change their mind less

The result is an increased convergence of voter agent attitudes, but supporters of the FPO are less influential here, due to the relative lack of interest in politics by their supporters. This may be unrealistic since a general lack of interest in politics may not be the same as a lack of interest in the particular, nativist, politics of the FPO. However, adding the above bias in openness to influence (away from those more involved in politics) has the effect of increasing the concentration of SPO and OVP supporters in the centre ground and the dominance of those two parties in the polls. Thus even though there is a general change in global salience to the issue of immigration it does not benefit the FPO party agent.

Sub-section conclusion. The conclusion of this subsection is that the attitude influence, saliency influence and broader interaction between agents (not just between those with *very* similar opinions) makes a significant difference – we do not get examples like the reference case with 'plausible' polling dynamics. Note that all of the above are where social influence about attitudes and voter agent saliences co-evolve with that of the network dynamics. In the next section we look at their impact without network dynamics and find they have a considerably reduced impact.

Varying the Network Dynamics

In this model, it is not only the attitudes and saliences of voters that adapt but also the social network itself. There are two components to this. (a) With a small probability, (given by *new-link-prob*) new random links can be formed either with a 'friend-of-a-friend' – another agent linked to those you already link with – or with a truly random other. The proportion of 'friend-of-a-friend' new links is set by the parameter: *fof-prob*. (b) the agents count how many 'disagreements' there have been with those they link with (interactions rejected due to attitudes being too far apart), so that when this number goes above a given threshold (set by the parameter: *drop-threshold*), that link is killed off (so no more interaction occurs between the previously connected agents). When active, these two mechanisms means that the social network itself adapts so that more similar agents tend to be connected.

Thus the first experiment is simply to turn off all network change (both network change processes). This is the next case illustrated below. Note that the social influence processes still operate, but over the same social network that it is initialised with at the start.

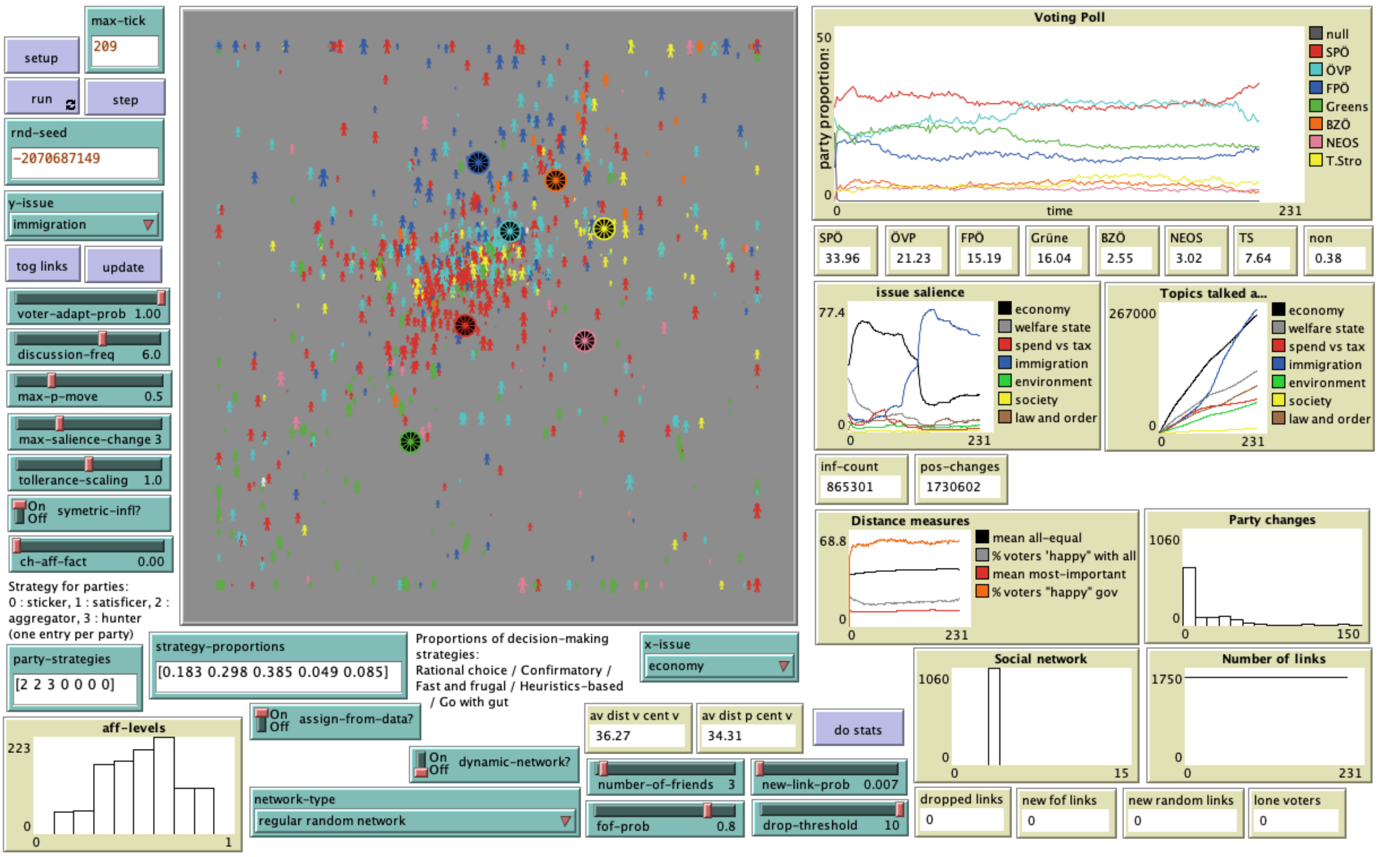

Figure 12. The case of no network change (a fixed social network)

The effect of this is to greatly lessen the convergence of voter agents, indeed some become more extreme in their attitude (the voter agents at the edge of the world view). This produces some polling dynamics but does not allow the FPO to break through and capitalise on the later global salience of the immigration issue. Unlike in the reference case, more of its first-wave of supporters remain at the extreme of thee immigration issue, but this means that the FPO party agent does not shift towards the centre ground to gain votes so it does not then appeal to those towards the centre.

The next experiment is to turn off only new link creation, setting *new-link-prob* to zero. This case is illustrated next.

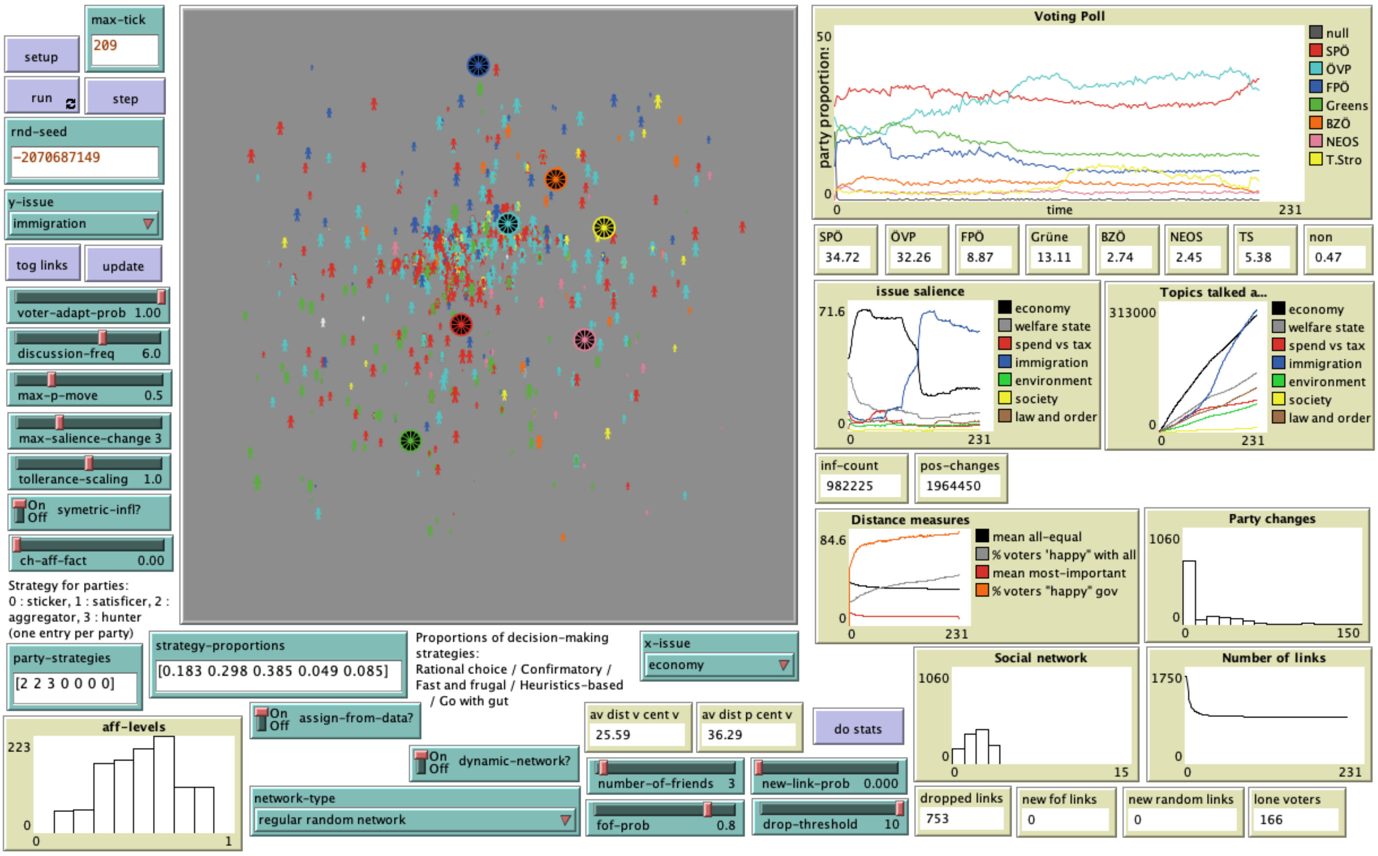

Figure 13. No link creation

Dropping some links but not creating new ones means that the social network becomes somewhat less connected over time, and this seems to prevent the creation of extreme voter agents, but also does not allow FPO supporters much purchase upon the majority of other voter agents.

The next comparison is with link dropping stopped (this is caused by setting drop-threshold to zero). Thus those that are linked at the start remain so for the duration of the simulation (although interaction over any particular link may decrease due to new links being created).

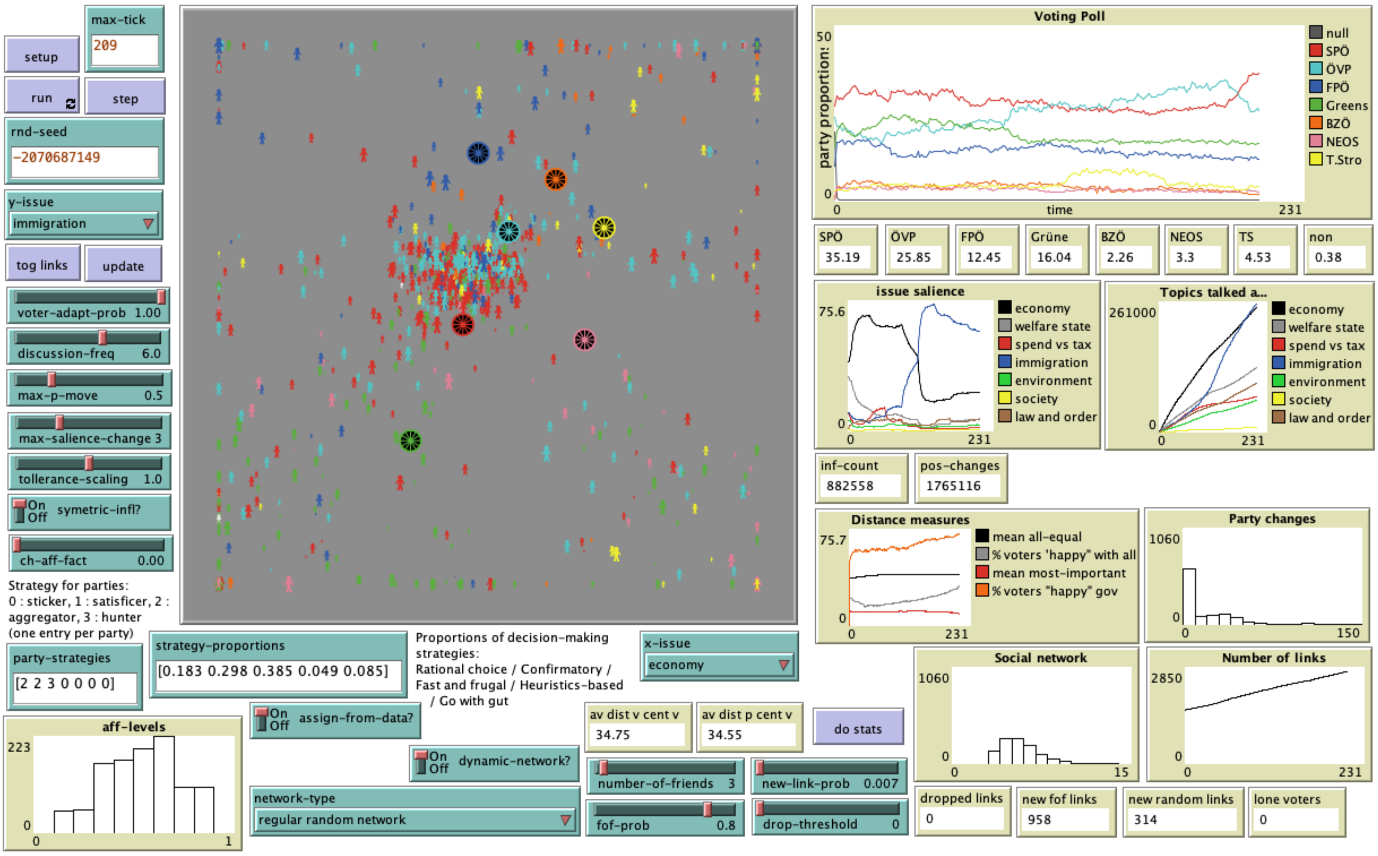

Figure 14. No link dropping

This looks very similar to the case of no network change at all (above), with some concentration of voters centrally (maybe a bit more than the no change case) but also some spread to the extremes.

Sub-section conclusion. The results of this subsection show that many aspects of a *changing* social network matter to the results (though less the rate of link dropping). There is very little in the political science literature that looks at how political discussion networks change over time as a result of interactions, but these simulations show that it can matter to the outcomes. This mirrors the results from a previous abstract toy model, that sometimes one needs both opinion change and network change to get certain results (Edmonds 2020). In this model, the network dynamics seem to be important in letting voters influence each other about a new concern (in this case immigration) whilst also allowing a convergence of attitudes. It is this combination that, in this model, allows a party like the FPO to suddenly gain votes by moderating its policies. The role and importance of the network dynamics are more systematically compared in the next section.

Varying Party Adaptation

The next aspect to look at is the adaption on behalf of the parties to voters. In this model four of the parties have fixed policies (Greens, BZO, NEOS and T.Stro), the other three of them adapt to the landscape of voter agents. Two of these use the 'aggregator' strategy (SPO and OVP) – that is, the party moves towards the average position of its current supporters, but only on the dimensions that are important to it. The FPO agent uses the 'hunter' strategy, a greedy hill-climbing method – that is, it keeps moving in same direction if it gained vote share with last move, otherwise they turn around. In other words this is a highly dynamic and opportunist strategy.

The next illustrates the case when dynamic party adaption is turned off (*max-p-move* set to zero).

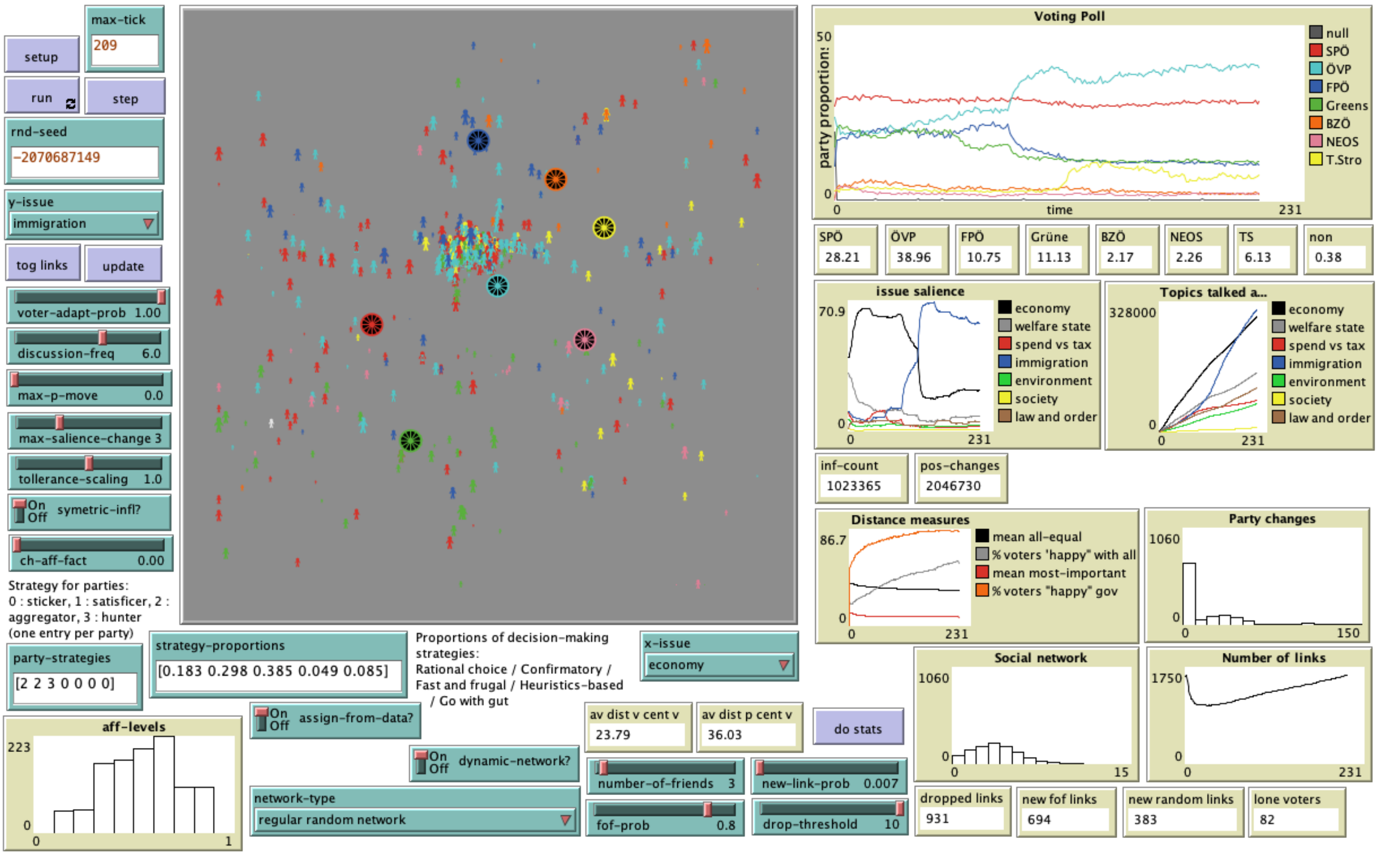

Figure 15. No adaption of policies by parties

Here, although the voters have influenced each other as in the reference case, resulting in a convergence of attitudes and voter agent salience, the three mobile party agents have not moved to take advantage of this, letting the OVP win, due to its initial policy position.

If we allow party adaption from these three parties but assign the FPO an aggregator strategy (like the SPO and OVP), we get the following.

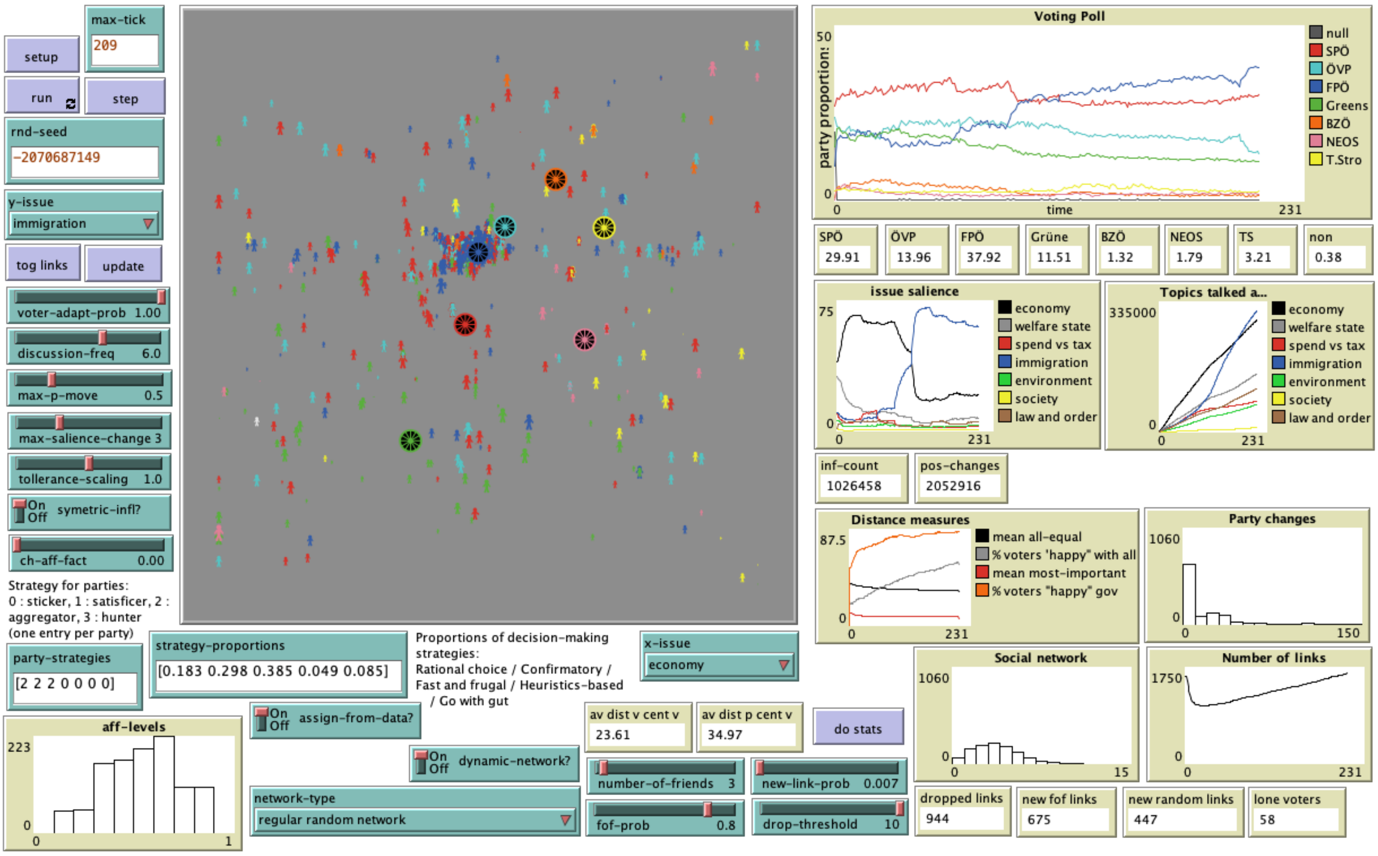

Figure 16. FPO using same strategy as SPO andd OVP – the *'aggregator'* strategy

This gives the FPO agent a solid advantage – both due to its position on immigration as this become increasingly salient, but also in its ability to locate itself at the centre of its developing supporter base. In this case (unlike the reference case) there is no sharp flipping of support between OVP and FPO and the FPO is acting in a non-populist manner by seeking the centre ground of support more cautiously like other parties (rather than opportunistically shifting).

Sub-section conclusion. The conclusion of this subsection is that how parties adapt to the voter landscape is important to the results making a significant difference.

Varying the Initial Network Characteristics

In the last subsection we look at the impact of different political discussion networks, as initialised in the model at the start. Exploring these possibilities is important as the overall structure of such networks in reality is largely unknown (due to the sheer difficulty of getting data about this).

The first contrast is to considerably increase the number of initial links voter agents have. This makes for a very densely connected network (in contrast to some evidence which suggests it is relatively sparse, e.g. Huckfeldt et al 1995). This is done by changing number-of-friends to 10. This is illustrated next.

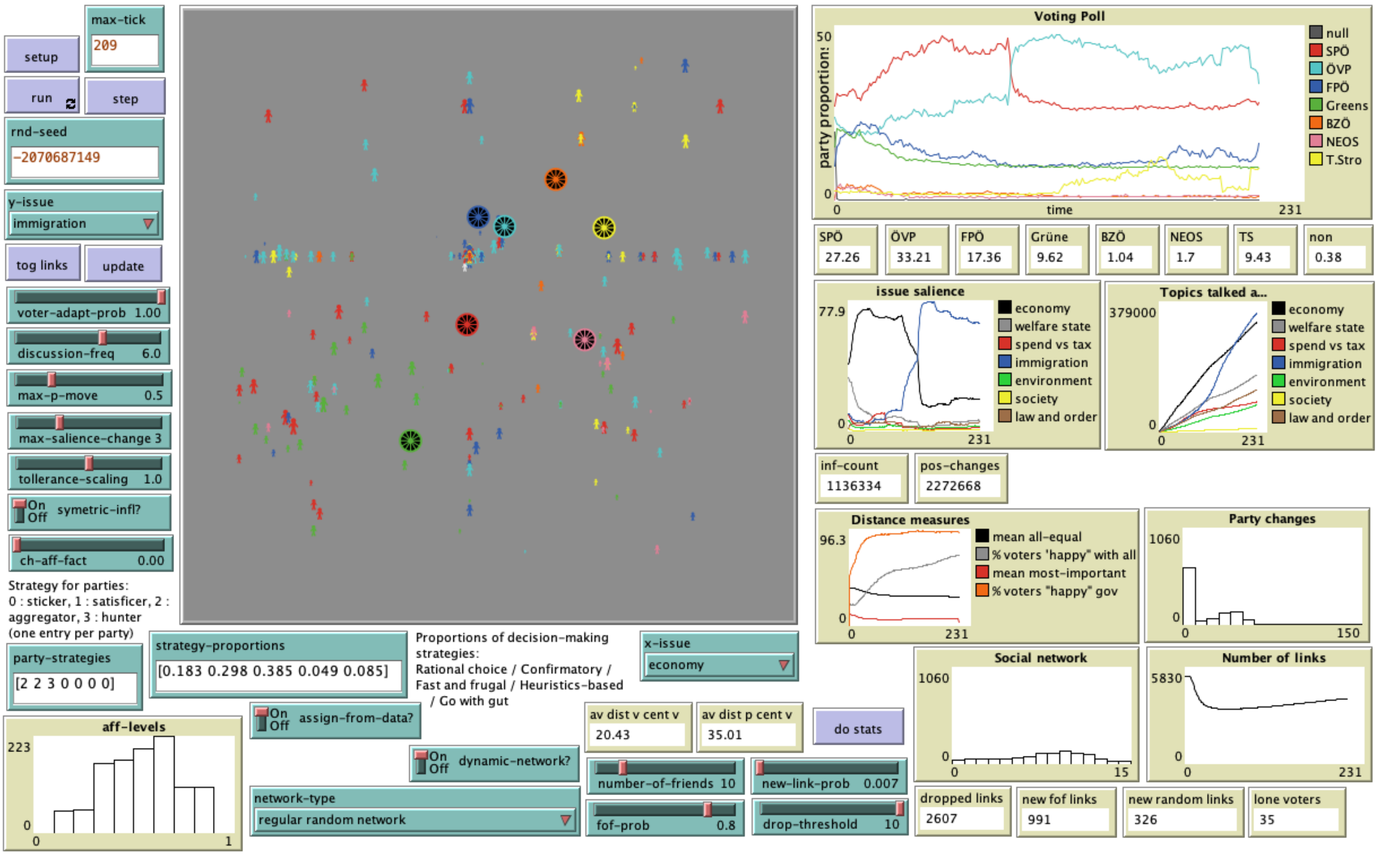

Figure 17. The case of a densely connected network

This shows a vastly increased amount of voter agent convergence, with the result that there is competition between only SPO and OVP party agents, with the latter capturing more votes as immigration becomes salient among voter agents.

The reference case was a random network, which meant that it was well connected to all other nodes (albeit through a jump or two) – in other words, it has a uniform rate of linking and a low network 'diameter'. This is in contrast to a network made through a preferential attachment process, which produces a few nodes with lots of connections, some with more than average and many with far fewer. This is the kind of structure found in some known influence structures (such as links in web pages). This results in a low-diameter but irregular network. This is illustrated in the next case.

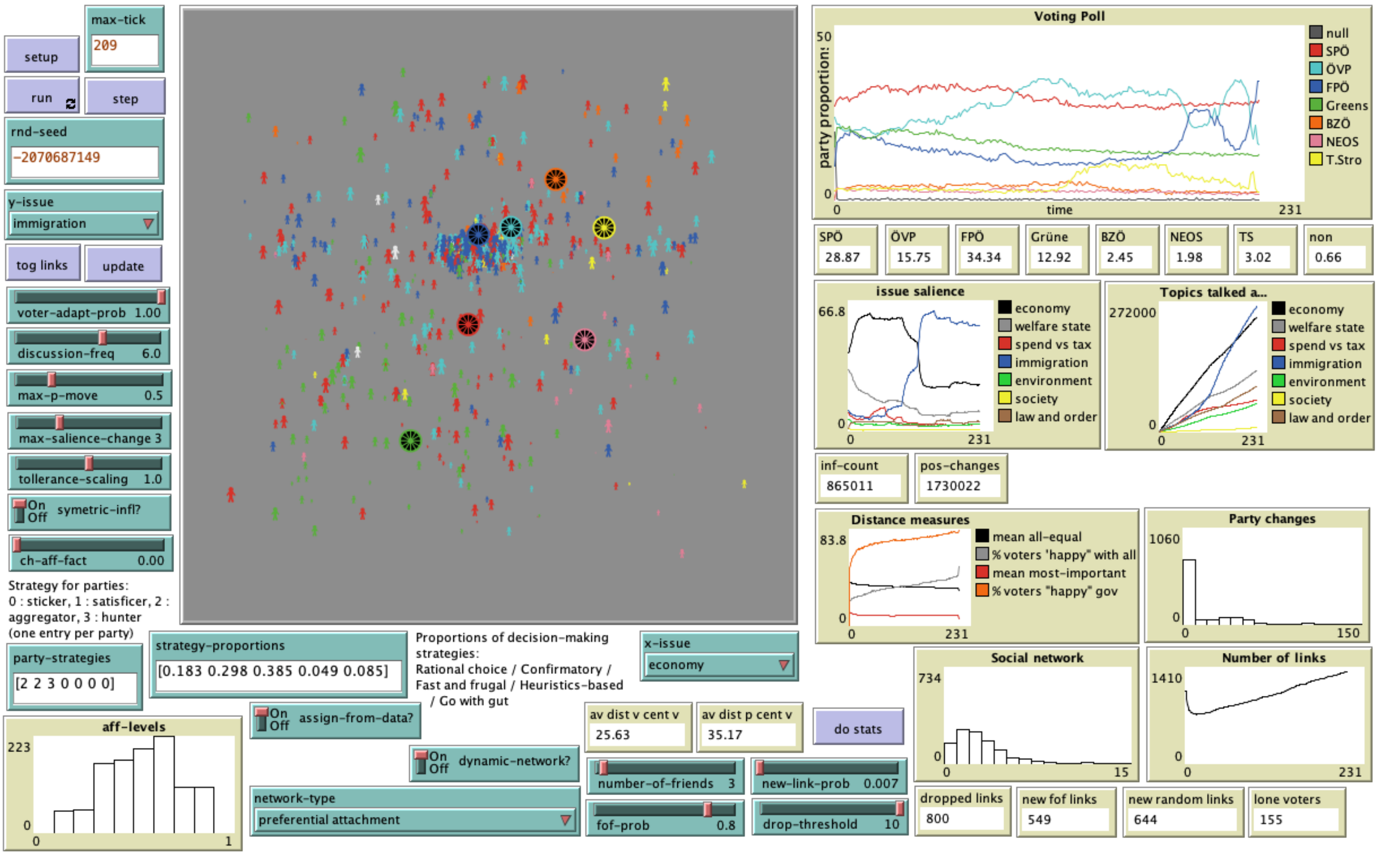

Figure 18. A preferential attachment network

This results in outcomes that are qualitatively similar to the reference model. The initial skewed distribution of links is evened out by the network dynamics over time.

Another network variation is that of a homophily-based network. This is where agents with similar attitudes tend to be connected at the start. This results in a regular but high-diameter network. This is illustrated in the next run.

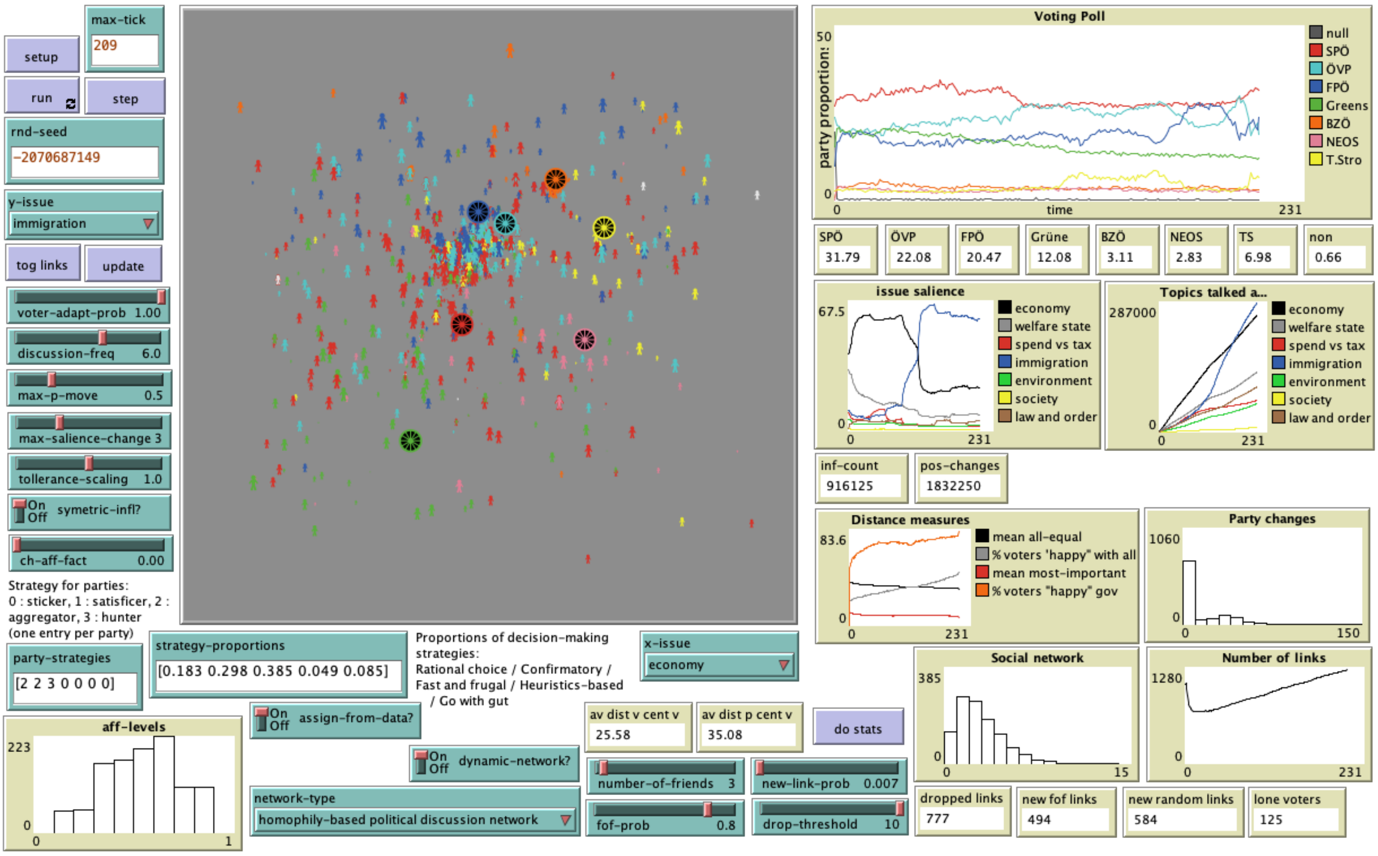

Figure 19. A homophily-based network

Again, this does not make a significant qualitative difference, with somewhat less of a convergence of voter agents, due to influence being less able to spread through the whole network, which allows the SPO to capture slightly more votes.

Sub-section conclusion. The conclusion of this sub-section is that whilst the overall connectivity matters to the outcomes, the other initial network characteristics do not matter critically to the results (in the same way as the variations in the other sub-sections). The diameter of the network seems to have had a slightly greater effect than the regularity of the network due to this being an aspect that is not 'flattened out' by the subsequent network dynamics. However, apart from saying that the initial network does not seem to impact the outcomes so much, it does not suggest any hypotheses other than the importance of network density.

General Sensitivity Analysis, with respect to: *what might increase the representativeness of democratic outcomes?*

Here we look at the sensitivity of results more systematically. The key output measure of interest we use is the proportion of voters within 10% of the centroid of policies in a notional elected government – that is: (a) see which parties get what numbers of votes according to current state of simulation, (b) pick parties starting from those with most votes until they represent a majority of votes (c) work out the average policy position of these parties, weighted by their number of votes (d) see what proportion of voters have attitudes, in the dimensions that are most salient to them, close to this. Although very imperfect (parties do not form coalitions regardless of their policies and the policy of the government is not the centroid of those of the coalition partners). However, this is acceptably close for our purposes.

Each point in the following line graphs (each value for each kind of network) is an average over 20 independent simulation runs.

The Non-Dynamic Network Case

We start the sensitivity analysis with the non-dynamic case of the simulation, so that the impact of various parameters without this level of complication. In each case we get each point on each line (i.e. for each network type) using 10 independent simulation runs to get an average value – typically 440 runs each graph.

Varying the rate of the rate with which voters adapt on interaction (*voter-adapt-prob*) is shown in the below.

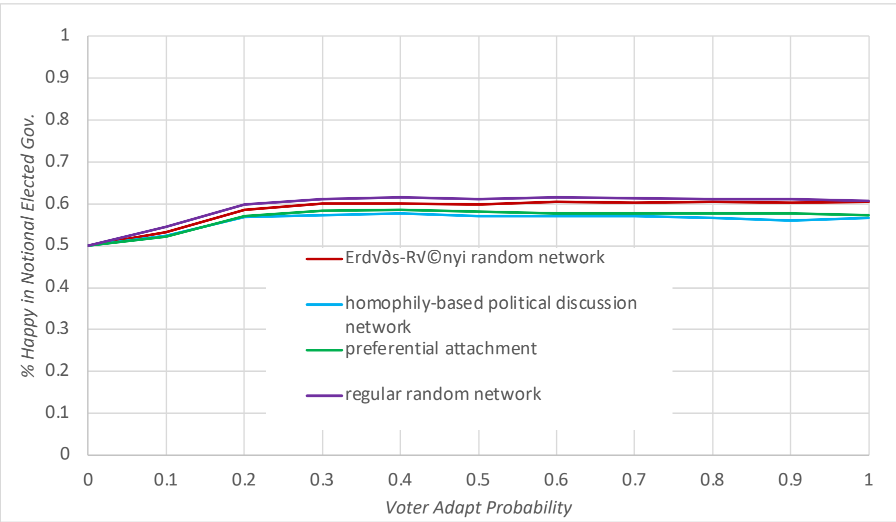

Figure 20. Comparing rates of voter adaption in the non-dynamic network case

Only small values of voter adapt probability (< 0.2) seem to matter here. Probably above that, voters are able to adapt enough to converge.

The two networks for which the *number-of-friends* parameter is varied is shown next.

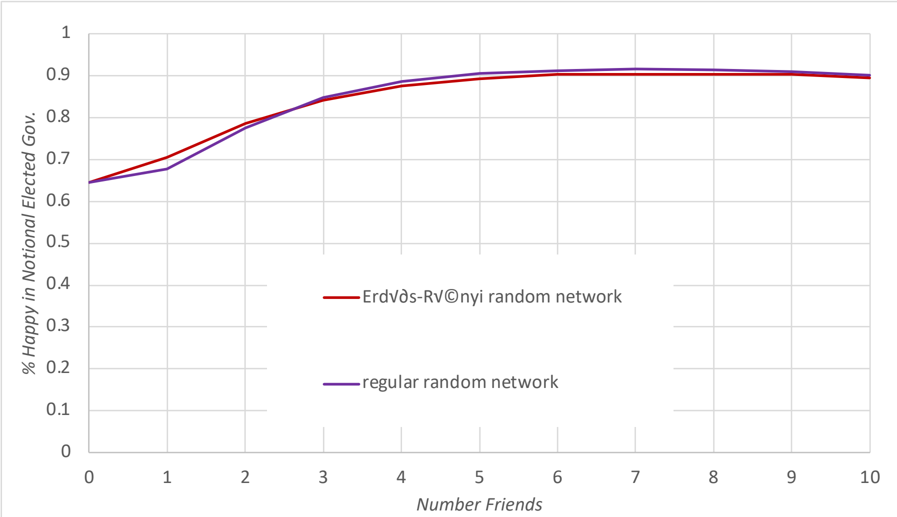

Figure 21. Comparing the number-of-friends in the non-dynamic network case

This seems to make a difference up to a value of 5. Above 5 the network is sufficiently dense that there is no significant diameter left and all agents can effect each other.

The max-salience-change parameter controls how much the voter agent salience of issues can adapt in agents on interaction with another (strictly, the maximum it can adapt).

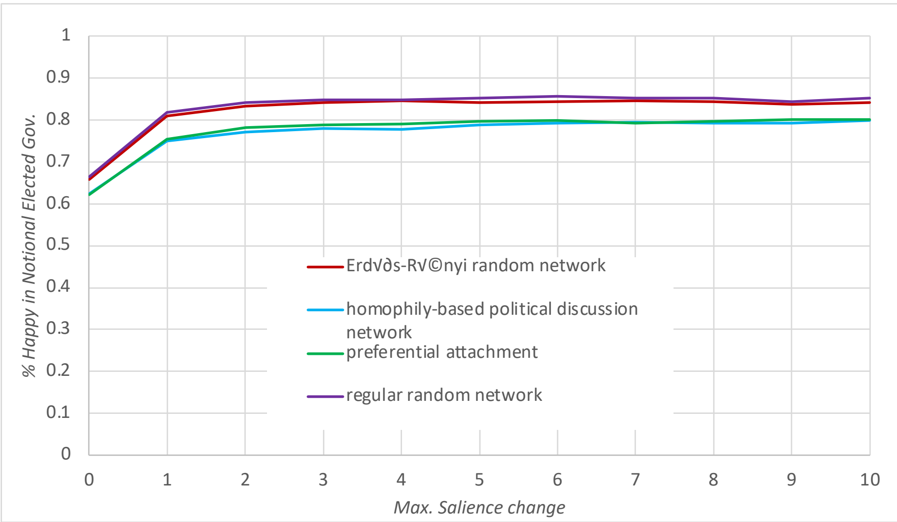

Figure 22. Comparing maximum rate of salience change in the non-dynamic network case

Again, some and no global salience change makes some difference but not at levels above this.

Next the amount that party agents can adapt (for ‘aggregator’ or ‘hunter’ strategies).

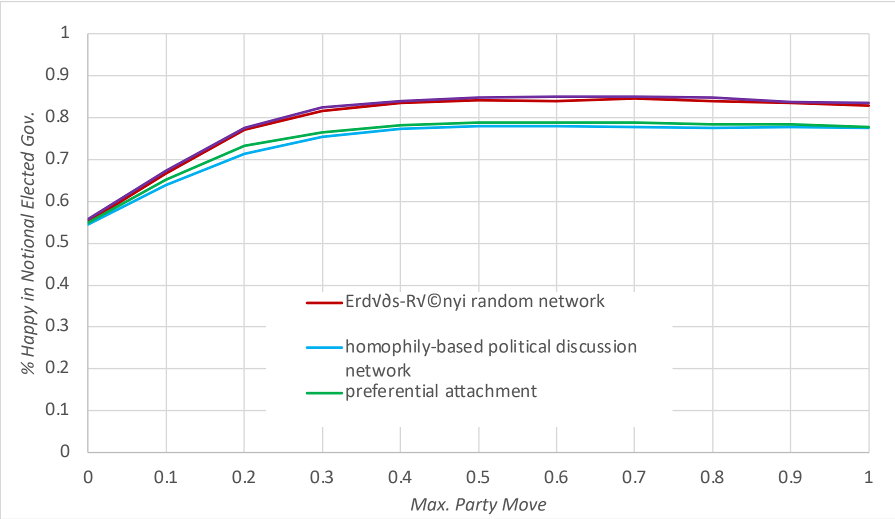

Figure 23. Comparing the rate of party adaption in the non-dynamic network case

Here values up to 0.4 make a difference, but not above that.

The value of the drop-threshold makes almost no difference in any case, so we will not show this.

Sub-section conclusion. As expected, the impact of any of these is not great in all cases where the network is not dynamic. There is also not much detectable difference between kinds initial network, however the two random networks tend to be very similar and are slightly higher than the homophily-based and preferential attachment versions (which are similar in this dimension to each other).

Dynamic/Non-Dynamic Network Comparisons

To see the impact of various non-network settings in dynamic and non-dynamic network cases we varied the following parameters with 10 independent runs for each set (192,000 simulation runs in total). Since the initial network is largely unknown, this is shown for each of these.

- *discussion-freq*: {1, 2}
- *max-p-move*: {0.5, 1}
- *voter-adapt-prob*: {0.5, 1}
- *max-salience-change*: {1.5, 3}
- *dynamic-network?*: {true, false}
- *network-type*: {"homophily-based political discussion network", "regular random network", "preferential attachment", "Erdös-Rényi random network"}
- *number-of-friends*: {1, 3, 5}

In each such diagram the error bars show one standard deviation either way.

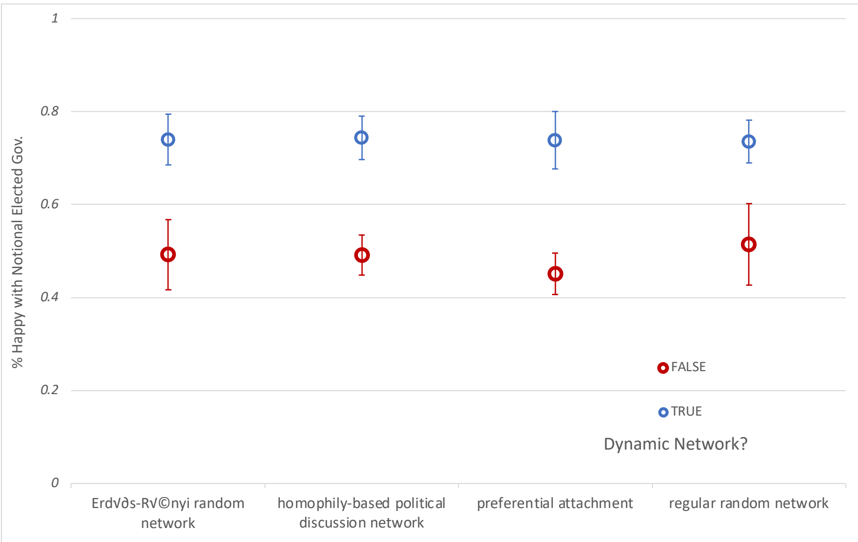

Figure 24. Overall contrast of dynamic vs. non-dynamic networks for each of four initial network configurations

The significance of the dynamism of the network is very clear[1]. For all the different kinds of initial interaction network it clearly makes a difference whether the network is dynamic or not,

The sub-case where it made the most difference was with the following settings:

- max-salience-change: 3
- voter-adapt-prob: 1
- max-p-move: 1
- discussion-freq: 2

[1] Note that standard approaches to statistical significance are not suitable for analysing simulation output. It is easy to get any level of statistical significance one wants, simply by running the simulation enough times. Indeed, it would be surprising if there was *any* change in settings in a simulation that would *not* show a statistically significant difference in the outcomes if run enough times. Statistical significance is used to rule out the case that apparent connections are due to noise rather than any identifiable mechanism – here we *know* there are such mechanisms since we programmed these into the simulation. The question here is to see which settings and processes have a notable effect on macro outcomes.

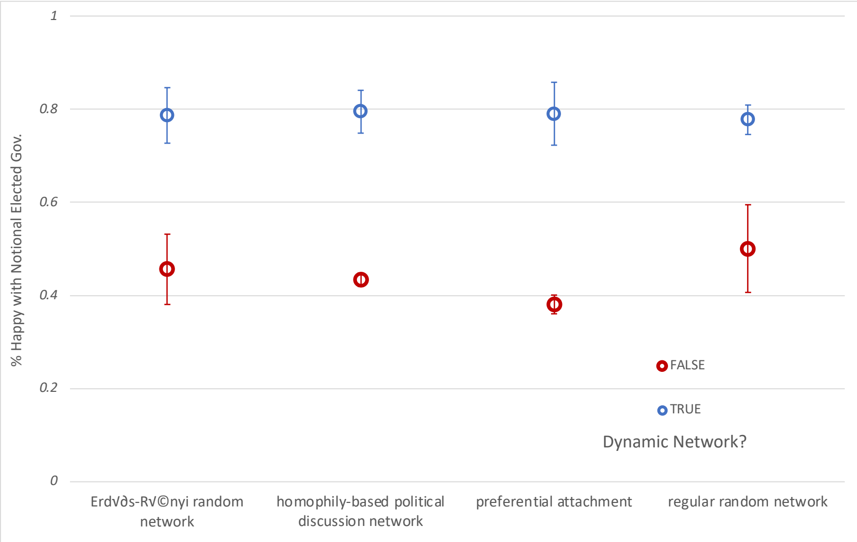

Figure 25. The sub-case where the dynamism of the network made the greatest difference

The sub-case where it made the least difference was as follows:

- max-salience-change: 1.5
- voter-adapt-prob: 0.5
- max-p-move: 0.5
- discussion-freq: 1

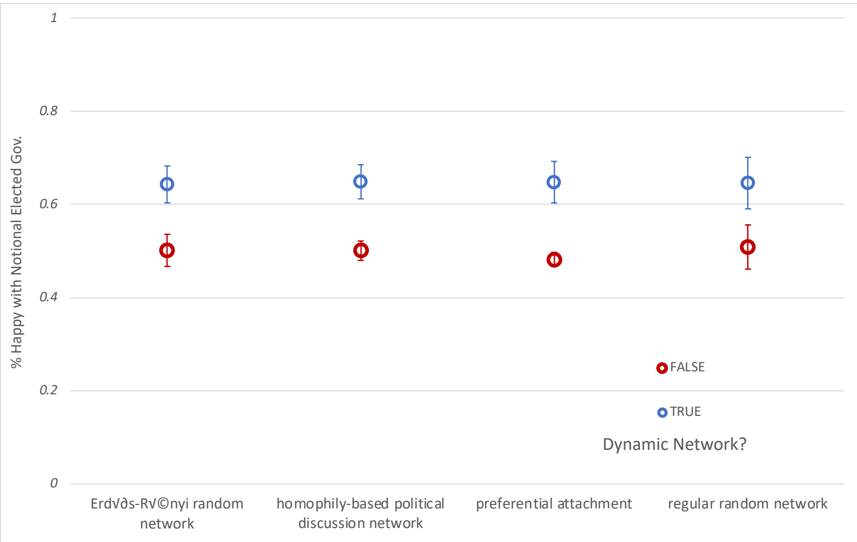

Figure 26. The sub-case where the dynamism of the network made the least difference

Lower values of *max-salience-change*, *voter-adapt-prob*, *max-p-move* and *discussion-freq* result in less difference between dynamic and non-dynamic networks, but this is still a clearly identifiable difference.

Sub-section conclusion. The conclusions from this sub-section is that the dynamism of the interaction network matters a lot in this model – we get very different results if this aspect is not present and not anything like the reference case (see the first section for more vivid illustrations of this).

Only Dynamic Networks

Here we assume a dynamic network but look at the impact of some of the parameters, namely the following (10 independent runs for each setting combination, making 64,800 simulation runs in total). Again, for each other aspect we do the comparison for each network type separately.

- network-type: {"homophily-based political discussion network" "regular random network" "preferential attachment" "Erdös-Rényi random network"}
- number-of-friends: {1, 3, 5}
- new-link-prob: {0.001, 0.003, 0.007, 0.01}
- drop-threshold: {1, 3, 10}
- tollerance-scaling: {0.5, 1, 1.5}
- ch-aff-fact: {0, 0.1, 0.2}

Firstly the *number-of-friends* value used in initial network construction (note this does not affect the homophily-based or preferential attachment networks as this parameter is not used for them as coded in this version of the model).

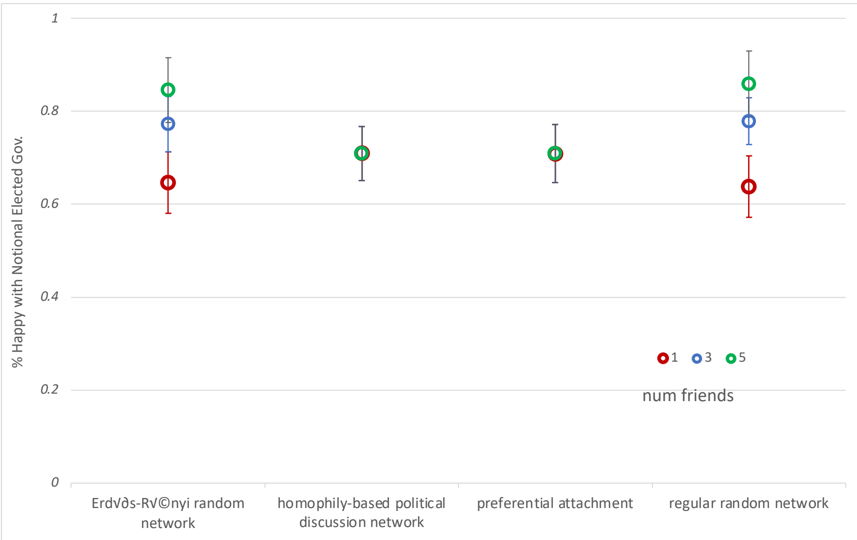

Figure 27. Comparing effect of number of initial links for different kinds of network

Here we see that a greater number of friends (the more densely connected the network), causes a significant increase in the % of 'happy' voter agents (in the networks where this parameter makes a difference).

Next the impact of the rate of new link creation (*new-link-prob*).

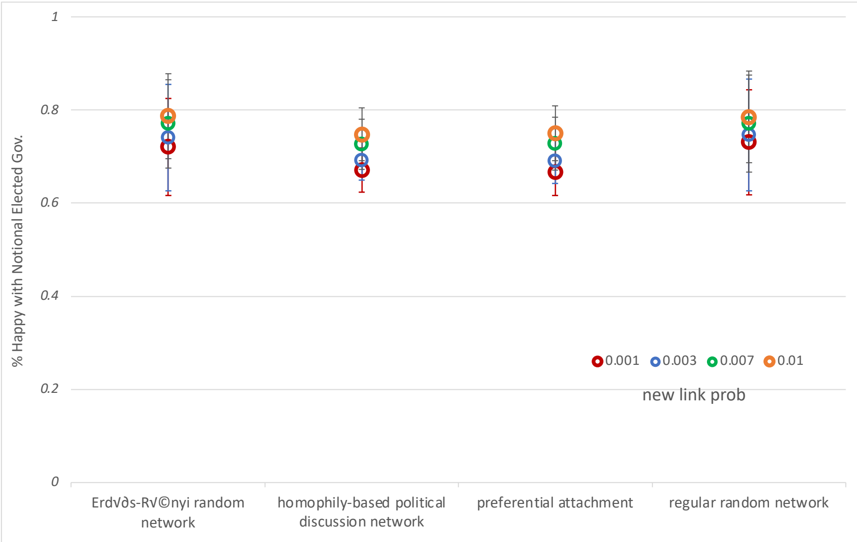

Figure 28. Comparing the effect of the rate of new link creation

Again, an increased rate of link creation results in a denser network, hence more voter convergence and thus it is more likely that a notional government would be closer to voter attributes, however the effect is not large (compared to the variation within each such set).

Next the threshold for dropping links.

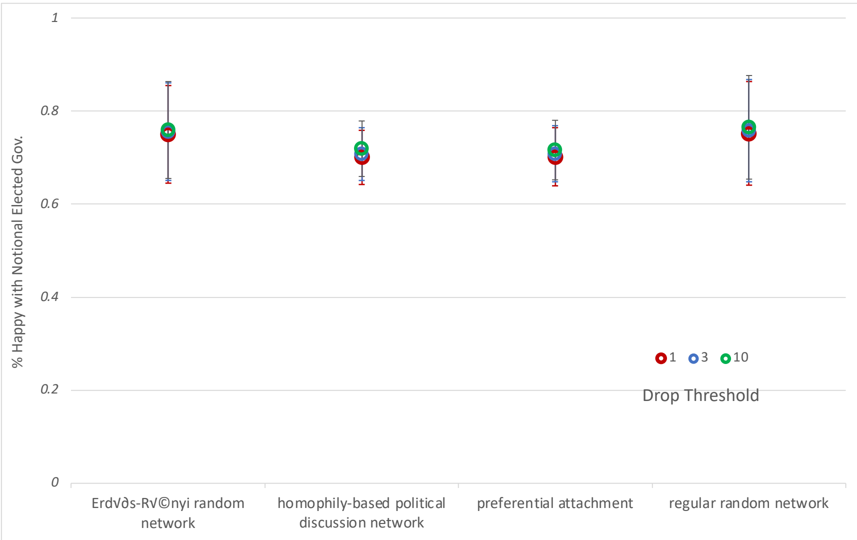

Figure 29. Comparing impact of threshold for dropping links

This seems to make very little difference to the results (at least using this target measure).

The tolerance scaling parameter (called *tollerance-scaling* in the model due to a typo) changes the scale over which attitudes of potentially interacting voter agents are compared to see if they are close enough to influence each other. Thus a small value of *tollerance-scaling* means that only very close agents will interact, a large value means that they interact even if distant.

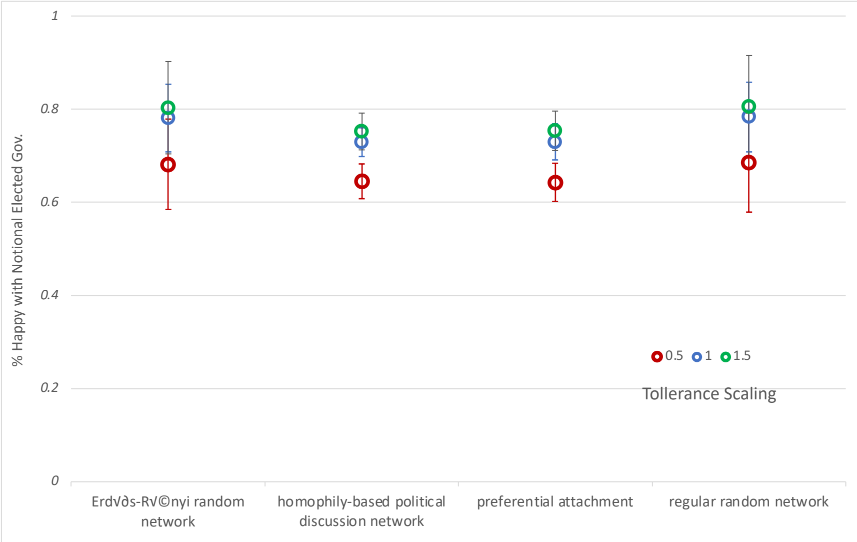

Figure 30. Comparing the effect of tolerance scaling

This seems to make a difference for small values of this parameter, but above a value of 1 most pairs of agents seem to interact so bigger values ceases to have an effect.

The *ch-aff* parameter controls the bias that affects the extent to which agents with higher level of political involvement change their voter agent salience less on interaction with another agent.

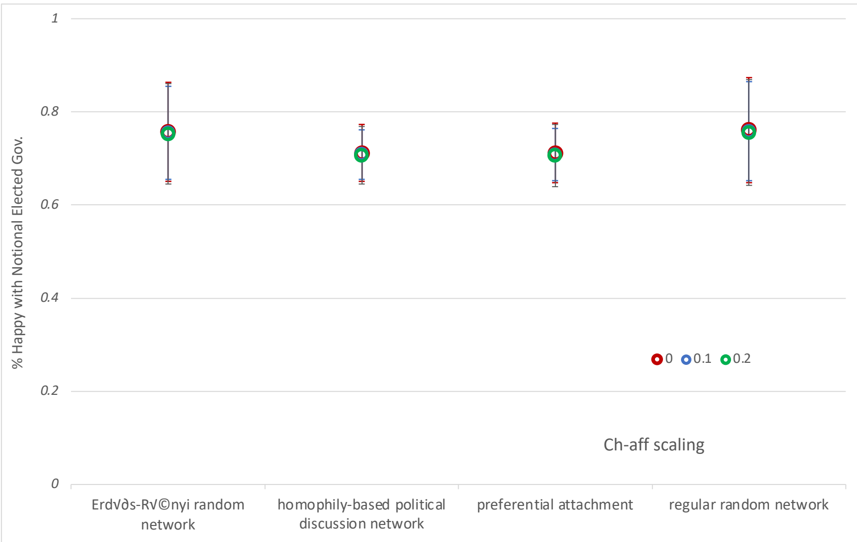

Figure 31. Comparing the impact of the *ch-aff* bias

The effect of this is not distinguishable.

The explorations in this sub-section show that (by this target measure) the parameters that effect either the density of the network or the amount of interaction make a difference, with higher densities (up to a point) and higher amounts of interaction increasing the proportion of voters happy with a notional elected government.

Conclusions

Let us look at each of our original research questions in turn.

How important are the social influence processes in producing plausible outcomes?

The social influence processes that cause attitude and salience change in voters is essential to the outcomes. In conjunction with the network dynamics, the number of initial friends (this affects network density), tolerance scaling (so agents interact with those of more different attitudes) and link formation (again affecting developing network density) make a noticeable difference. However, they only make a significant difference in combination with network dynamics.

How important are the network dynamics in producing plausible outcomes?

The network dynamics are essential for producing results like that of the reference case in this model. The opinion dynamics and network dynamics co-evolve and reinforce each other. The dynamism of the network makes most difference with more discussion between agents, more adaptivity by voters in terms of attitudes and salience change, and a greater adaptivity from those parties who change policies in response to the voter attitude landscape.

How important are party dynamic behaviour in producing plausible outcomes?

The co-evolution of party agent policies and voters is another key element of this simulation. This is necessary, in this model, in order to get results similar to that of the reference case. Furthermore, if the FPO agent uses the same strategy as the SPO or OVP we do not get similar outcomes. The opportunism of the FPO agent is necessary in contrast to the more cautious adaptive strategies of the SPO or OVP agents.

How sensitive is the model to its various settings and parameters?

The simulation model is moderately sensitive to a number of its parameters, but these are not arbitrary parameters and all are, at least in principle, empirically derivable. The level of sensitivity reflects our understanding of the situation represented. None of the continuous parameters requires very particular values for the simulation to ‘work’. Notably the simulation is not so sensitive to particular kind of initial interaction network, as long as it is sufficiently connected and network dynamics (particularly new ‘friend-of-a-friend’ link creation) occur. This is fortuitous as there is very little data upon what this network looks like.

What impact do the settings have on the convergence of voter attitudes and how 'effective' is a resulting (notional) government in reflecting these attitudes?

In the model, the convergence of voters on issues, but particularly a new issue (e.g. immigration), helps parties adapt so as to better represent the issues the voters care most about in a subsequent notional government. Adaptivity in terms of all of: voter attitudes and saliences, the interaction network and party policies all play a part in this, and thus enable the system to better represent voters. In the case of the adaptivity of the network this can vary between an average of 80% 'happy' (the proportion of voters whose attitudes, that are important to them, are close to those of a notional elected government) down to below 50%. In this model, party agents like the FPO can play a role in this adaptation, but lose out as a party compared to their electoral success if they adopted more cautious strategies that are more similar to the SPO and OVP.

Are there any insights from this analysis that relate to populism?

Clearly, this simulation does not cover all the processes and structures that are necessary to understand populism. However, it does touch on *some* important aspects.

In the model, the (final) success of the FPO agent was down to a number of factors:

1. The quickness with which a newly important issue (in this case immigration) can gain *widespread* voter salience within a population as a result of a combination of social interaction and network change.
2. The ability of a population of voters to converge in terms of attitudes on a newly important issue, shifting the centre of general opinion and providing a newly positioned reservoir of voters that populists might appeal to.
3. The speed with which a party agent (such as the FPO) can respond to such a new attitude landscape so as to pick up a surprising number of votes (sometimes just before an election).

It is the co-evolution of (1) and (2) above that can result in a very rapid change in the electorate. In this model this ability is inherent but can be mobilised under the right circumstances. It can be hard for more traditional political parties to adapt quickly to such new developments due to them being more constrained by liberal norms and due to the fact that their policy is formed as a result of a complex internal negotiation. Many populist and nativist parties are dominated by a single leader and are less constrained in these ways.

Suggested Areas for Further Empirical Research

Whilst these explorations do not prove anything directly about observed political phenomena, they are suggestive of hypotheses that might merit further empirical investigation. These include:

- How does social influence impact upon voters' attitudes but also the importance they attribute to them?
- Under what circumstances are voters open to social influence from those they interact with?
- How does the social interaction network co-adapt with those of voters' attitudes?
- When do different kinds of parties adapt their policies and by how much?

Acknowledgements

The agent-based simulation was developed by Ruth Meyer. Data and expertise used to inform the model was provided by Marco Fölsch, Martin Dolezal & Reinhard Heinisch. Bruce Edmonds did the analysis reported in this document using the model. The research in this report was part of the Populism and Civic Engagement project (PaCE), that received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 822337. For more about PaCE and access to the other deliverables mentioned in this report see the project website at http://popandce.eu

References

Deffuant, G., Amblard, F., Weisbuch G. and Faure, T. (2002), How can extremism prevail? A study based on the relative agreement interaction model, *Journal of Artificial Societies and Social Simulation*, 5(4):1. https://www.jasss.org/5/4/1.html

Edmonds, B. (2020) Co-developing beliefs and social influence networks – towards understanding socio-cognitive processes like Brexit. *Quality & Quantity*, 54(2):491-515. DOI:10.1007/s11135-019-00891-9

Edmonds, B., le Page, C., Bithell, M., Chattoe-Brown, E., Grimm, V., Meyer, R., Montañola-Sales, C., Ormerod, P., Root H. & Squazzoni. F. (2019) Different Modelling Purposes. *Journal of Artificial Societies and Social Simulation*, 22(3):6. http://jasss.soc.surrey.ac.uk/22/3/6.html.

Huckfeldt, Robert, Paul Allen Beck, Russell J. Dalton and Jeffrey Levine (1995) Political Environments, Cohesive Social Groups, and the Communication of Public Opinion. *American Journal of Political Science*, 39(4):1025-1054.

Lau, Richard R, Mona S Kleinberg, and Tessa M Ditonto. 2018. "Measuring voter decision strategies in political behavior and public opinion research." *Public Opinion Quarterly* 82 (S1): 911–936.

Laver, Michael, and Sergenti, Ernest. 2012. *Party Competition: An Agent-Based Model*. Princeton and Oxford: Princeton University Press.

Meyer, R., Fölsch, M., Dolezal, M. & Heinisch, R. (2021) An Evidence-Driven Model of Voting and Party Competition. In *Advances in Social Simulation*, Springer, pp. 261–273. https://link.springer.com/chapter/10.1007/978-3-030-92843-8_20

Muis, Jasper, and Michel Scholte. 2013. "How to find the 'winning formula'? Conducting simulation experiments to grasp the tactical moves and fortunes of populist radical right parties." Acta Politica 48 (1): 22–46.