The Duration of Multilateral Negotiations in the Council of the European Union

WORK IN PROGRESS

Frank M. Häge[1]
University of Limerick

Abstract

Most existing formal theories of group choice focus on predicting the outcome of decision-making. However, the duration of decision-making is also of substantial interest, as it indicates the ability of a collective actor to respond to changes in the social, political, and economic environment in an efficient and timely manner. I use an agent-based model of coalition building to generate hypotheses about the duration of multilateral international negotiations. The model produces predictions about the effects of changes in the number of negotiators, the initial preference distribution, and the general level of impatience of negotiators. According to the model, increases in the number of actors and decreases in impatience lead to increases in the duration of decision-making. In these instances, the model provides a clear mechanism for generally accepted cause-effect relationships. However, it also provides somewhat counter-intuitive predictions: preference heterogeneity affects decision-making duration only marginally and preference polarization not at all. In this respect, the model provides novel alternative hypothesis that can be tested against existing theoretical accounts.

[1] Email: frank.haege@ul.ie; web: www.frankhaege.eu.

The duration of EU legislative decision-making

The duration of legislative decision-making plays an important role for the functioning and legitimacy of the European Union (EU). Like any other political system, the EU needs to be able to process its legislative workload efficiently and to provide policy responses to economic, social, and political problems in a timely manner. Efficient decision-making averts legislative paralysis and improves the public's perception of the EU as a valuable form of governance. Corresponding to its practical relevance, the duration of decision-making has long been of interest to scholars of EU politics (Golub 1999, 2007; Golub and Steunenberg 2007; Hertz and Leuffen 2011; König 2007; Schulz and König 2000). Existing studies pay particular attention to the effect of changes in preference configurations and institutional rules, and despite major methodological disagreements (Golub 2008; König 2008) have produced a number of consistent empirical findings. In particular, all study results indicate that the stronger powers of the European Parliament under the cooperation and codecision procedure prolong the decision-making process, and that the possibility of adopting Council decisions by a qualified majority of member states' votes shortens it. Notwithstanding fundamental differences in measurement, the findings are also in agreement that diverging member state preferences increase the time it takes to adopt European laws.

This paper does not challenge the validity of these empirical findings but provides an alternative theoretical underpinning. Existing research develops hypotheses mainly from the spatial theory of voting (for the most explicit discussions, see Golub 2007; Schulz and König 2000), which makes predictions about the conditions under which policy change will or will not occur. The usefulness of this theoretical framework for the study of decision-making duration is however questionable, as essentially static spatial models do not provide predictions about the time required to adopt a decision.[2] In the spatial theory of voting, fully informed decision-makers compare a proposal to the status quo policy and either vote in favour or against it. Collective decisions are reached instantaneously. Spatial theory predicts a change in policy if the status quo point lies outside the core and no change if it lies inside the core. The core represents the set of policies that cannot be defeated by another proposal in a pair-wise voting contest. The size of the core depends on the preferences of actors whose agreement is required to pass legislation. Increased preference divergence, more inclusive

[2] In addition to the spatial theory of voting, Golub (2007) also refers to coalition theory to justify his hypotheses about the effect of enlargement and Council voting rules. Coalition theory makes predictions about the proportion of possible winning coalitions, not about the duration of the decision-making process. In this respect, coalition theory is just as inadequate as an explanatory framework for decision-making speed as the spatial theory of voting.

voting rules, enlargement, and the empowerment of the EP are all associated with a potential increase in the size of the core and therefore with an increase in policy stability. Existing studies of decision-making speed suggest that the size of the core is not only related to policy stability, but to the length of the decision-making process as well (Golub 2007, 157; Schulz and König 2000, 656). However, the spatial theory of voting provides no justification for such an expectation. Indeed, such a connection can at best be justified by invoking auxiliary assumptions. According to Schulz and König (2000, 656), the policy gridlock predicted by the simple spatial model can "in reality" be overcome by side payments or issue-linkages.[3] Similarly, Golub (2007: 158) states that agreements are delayed until a "complex package deal involving the introduction of a new dimension" is crafted.[4]

Recourse to these additional assumptions acknowledges deficiencies of the simple spatial model as an explanation for decision-making speed but does not provide a solution. If it is correct that the gridlock predicted by the model is in practice regularly overcome through various compromise mechanisms, then the model misses these essential features of decision-making. If the compromise mechanisms do the real theoretical work, they should be explicitly modelled. As they stand, the informal auxiliary arguments indicate that those mechanisms affect policy stability but they still do not allow the derivation of expectations about decision-making speed. Why should the adoption of a decision through side-payments, package deals, or issue-linkages require more time than the adoption of a proposal involving fewer issue dimensions and less vote trading? If preferences are common knowledge, decisions could be made just as instantly as in the simple spatial model; if they are not, decision-making in the simple spatial model could take just as much time as in a model allowing for higher-dimensional exchanges. Thus, while many of the empirical findings of the extant literature are very persuasive, the theoretical rationales and causal mechanisms underlying those correlations require further theorizing.

In this paper, I contribute to the theoretical literature on the duration of decision-making in the EU and international organisations more generally by presenting a dynamic agent-based model of coalition-building in multilateral negotiations. The agent-based modelling framework is especially suitable for studying multilateral negotiations, a highly

[3] See also König (2008, 154).

[4] Golub (2007) also mentions the possibility that a shift in an actor's position, for example through the arrival of a new government, overcomes the predicted gridlock. An exogenous change in actors' preferences seems to be the only mechanism for overcoming gridlock consistent with the basic assumptions of the spatial theory of voting. Given that preference change must affect not just any member state but one of the pivotal actors for this mechanism to work, and given that even governments of different political persuasion often represent the same 'national' interest in the Council, the practical relevance of this mechanism seems marginal.

complex phenomenon that involves a large number of heterogeneous actors and their path-dependent interactions over time. The basic idea underlying the model is that negotiators' coalition-building behaviour is driven by their desire to form a blocking coalition. Only being part of a blocking coalition ensures that the negotiator's views are taken into account when the final compromise outcome is negotiated. Although negotiators have a clear incentive to join other negotiators to form blocking coalitions, they would also like the blocking coalition to represent a position as close as possible to their own most preferred policy. Thus, negotiators incur a policy cost if they change position. In addition, negotiators might be more or less patient for various reasons. The more patient negotiators will be less likely to make policy concessions than the less patient ones. Finally, the shadow of the vote increases during the process of negotiations and might induce more flexibility in the positions of previously more reluctant negotiators. Thus, in the model, a negotiator's probability to change her position is determined by a simultaneous evaluation of policy costs, impatience, and the fear of being outvoted.

In the next section, I describe the implementation of the agent-based model in more detail.[5] Special attention is given to the way agents determine the probability of joining the policy position of other negotiators. The general dynamics of the model and individual negotiators coalition-building decisions are then illustrated through a description of a simulation run. Following the illustrative example, I describe the setup and results of computational experiments to derive empirically testable hypotheses. The experiments assess the effects of different preference constellations, general levels of impatience, and changes in the number of member states. I conclude the article by discussing the main theoretical results of the analysis as well as limitations and scope conditions of the theory.

Modelling the duration of multilateral negotiations

Multilateral negotiations usually take place in a high-dimensional issue space and involve a large number of actors with heterogeneous attributes. Amongst other things, these individual differences include varying policy preferences, power resources, and impatience levels. In addition, coalition-building is a path-dependent process in which current decisions are crucially affected by earlier choices as well as initial conditions at the start of the negotiation process. As negotiations unfold over time, system properties first change as a result of the interactions of individuals; but subsequently, these changed system properties in turn affect the future attributes and behaviour of individuals, generating continuous causal feedback

[5] The model is implemented in Netlogo 4.1.3 (Wilensky 1999).

loops between the macro- and micro-level of the negotiation system. Capturing the essential features of such complex systems in mathematical models that are analytically tractable is increasingly difficult and computational models offer clear advantages in this respect.

The computational model employed here was originally developed to explain the apparently consensual nature of decision-making in the Council under qualified majority voting (Häge 2010). The basic idea underlying the model is that negotiators' behaviour can largely be explained by their desire to form a blocking minority. Negotiators are policy-seekers, which means that they are generally reluctant to make policy concessions. At the same time, negotiators are aware that they might be outvoted if their counterparts are able to strike a compromise that has the support of a winning coalition. In order to avoid isolation, they successively join the positions of other negotiators with similar policy views until their coalition is large enough to block any decision. In more technical terms, negotiators in the model employ the following simple coalition-building rule: If the closest coalition in terms of policy similarity is equal to or larger in size than the current coalition, then support the position of the other coalition, else stick to the current position. A collective decision is adopted by vote if one coalition reaches winning majority size before any other coalition has reached blocking minority size and by consensus if all negotiators are part of blocking coalitions. Based on these assumptions, the original model successfully reproduced the high consensus rate of decisions in the Council, including its insensitivity to changes in the number of member states over time.

For the current purposes, the model is extended to generate hypotheses about the duration of Council decision-making. While the original model assumed that negotiators joined larger coalitions with similar policy views deterministically in each negotiation round whenever the conditions outlined above were met, the revised version models this move stochastically. Two individual-level factors and one system-level factor are assumed to affect a negotiator's probability to join another coalition. In line with existing theorising on decision-making and bargaining, policy similarity and patience are individual-level factors that should decrease the probability of a negotiator changing positions. At the same time, a growing risk of being outvoted should increase the probability that a negotiator changes position. Amongst these three factors, only impatience is determined exogenously through a random draw from the beta distribution with shape parameters α and β. The resulting variable takes values between 0 and 1.

$$Impatience \sim Beta(\alpha, \beta)$$

Impatience is a convenient summary variable for all individual-level factors that affect the propensity of a negotiator to make concessions, are given at the outset of negotiations, and stay relatively stable over time. Amongst other characteristics, such factors may include the salience attached to a proposal by government, the efficiency of national coordination processes, or the involvement of domestic veto players. At the aggregate level, the shape and especially the central tendency of the distribution of impatience levels may reflect deadline effects, levels of politicization, or a collective 'sense of urgency'. The two parameters of the beta distribution allow for a wide variety of shapes to be investigated in the computational experiments.

Policy similarity is operationalized as the distance between the negotiator's ideal point and the position of the closest alternative coalition in the current point in time. Which coalition is currently the closest alternative coalition depends on the negotiators previous coalition choices and the resulting path through the policy space. To avoid arbitrary effects of the overall size of the policy space, the distance is normalised by dividing it by the maximally possible distance. The maximally possible distance is the diagonal distance between two corner points of the policy space. Finally, the resulting fraction is subtracted from 1 to transform the variable to indicate similarity rather than dissimilarity values. More formally, policy similarity is determined through the following equation, where *x* and *y* refer to policy space coordinate values of focal negotiator *i* and alternative coalition member *j*. The numerator represents the Euclidean distance between the negotiator's ideal point and the position of the alternative coalition;[6] and the denominator represents the scaling factor, that is the Euclidean distance between the lower left and upper right corner points of the policy space.

$$Policy\ similarity = 1 - \frac{\sqrt{\left(x_j - x_i\right)^2 + \left(y_j - y_i\right)^2}}{\sqrt{(x_{max} - x_{min})^2 + (y_{max} - y_{min})^2}}$$

Finally, the prospect of being outvoted or the 'shadow of the vote' is operationalized by first calculating the difference between the number of negotiators in blocking coalitions plus the largest non-blocking coalition and the number of negotiators required to adopt a decision by vote. This term is then normalized by the difference between the minimum number in the largest non-blocking coalition (i.e. 1) and the number of negotiators required to adopt a

[6] Note that apart from early negotiation rounds, the ideal point of a negotiator is usually different from her current position.

decision by vote. The normalized term provides a measure of distance to the majority threshold, which decreases with more blocking coalitions and larger coalition sizes. To produce a measure of closeness to the majority threshold that increases with coalition sizes, the normalized distance measure needs to be subtracted from 1. The idea underlying this operationalization is that a negotiator's pressure to form a blocking coalition increases the closer other negotiators are to form a winning majority. As the size of coalitions depends on negotiators' coalition-building behaviour, the value of the shadow of the vote variable is endogenous to the negotiation process. However, in contrast to the also endogenous policy similarity variable, which takes different values for different negotiators, the shadow of the vote is the same for all negotiators not in a blocking minority coalition. To illustrate the calculation of variable values, consider the situation of 27 member states and a qualified-majority voting threshold of 72 per cent, where the number of negotiators required to adopt a decision by vote is 20.[7] As long as negotiators have not formed a blocking coalition, the shadow of the vote is simply measured as 1 minus the normalized difference between 20 and the currently largest coalition. As soon as some member states have formed a blocking coalition, which in this case consists of 8 member states, the shadow of the vote is measured as 1 minus the normalized difference in the number of negotiators in blocking coalitions plus the number of negotiators in the largest non-blocking coalition. The equation for the calculation of the shadow of the vote variable is given below. In this equation, *n* refers to the number of negotiators in blocking coalition k and non-blocking coalition *l*, respectively. The number of negotiators required to form a winning coalition is represented by $n*$.

$$Shadow\ of\ vote = 1 - \frac{n^* - (\sum n_k + \max(n_l))}{n^* - 1}$$

All three variables take values between 0 and 1, but apart from those boundary points, it is not obvious how their values should map onto the probability to change policy position. Given that impatience scores are already defined as probability scores, they can be treated without any further transformation as values of negotiators' innate propensity to change position. How the values of the other two variables translate into probability scores is less clear-cut. Treating probability scores as linear functions of those variable values seems implausible. In the case of policy similarity, a linear mapping would mean that negotiators who have to traverse 20 per cent of the policy space to join a larger coalition still had an 80 per cent probability to do so. In general, policy distances between negotiators of more than 30

[7] For simplicity, the current version of the model does not allow for variation in voting weights.

per cent of the policy space are quite rare. Thus, a linear mapping would result in probability scores that are almost always higher than 70 per cent. A non-linear relationship, with probability scores decreasing more strongly than policy similarity values, is more realistic. The power transformation $y = x^k$, with $k > 1$ satisfies this requirement while keeping the end points of the variable's scale unchanged.[8] The higher the exponent k, the stronger the decrease in the probability scores resulting from deviations from perfect policy similarity. The model's default mapping with a relatively high exponent of $k = 7$ is illustrated in the left panel of Figure 1.[9] According to this equation, a policy similarity score of 0.9, which corresponds to a distance of 10 per cent of the policy space, reduces the probability score to just under 0.5. A further reduction to a policy similarity score of 0.8, which corresponds to a distance of 20 per cent of the policy space, reduces the probability score to about 0.2.

A linear mapping onto probability scores is also implausible for shadow of the vote values. Such a relationship would imply that the addition of one negotiator to a coalition has the same effect on the probability score at the start of negotiations as when coalitions are close to forming a winning majority. However, as long as the number of negotiators with or approaching blocking minority status is nowhere near winning majority size, the shadow of the vote should have minimal effect. At the same time, when a winning majority is close to being formed, the shadow of the vote should be overwhelming. The S-shaped function $y = x^k/(x^k + (1 - x)^k)$, with $k > 1$ ensures that shadow of the vote values on the lower half of the scale are translated into disproportionately low probability scores and values on the higher half are translated into disproportionately high probability scores. When $k = 1$, the equation reduces to a linear relationship, and the larger k gets, the more disproportionate the mapping of shadow of the vote values into probability scores becomes. The right panel of Figure 1 illustrates the relationship with a relatively modest exponent $k = 2$.

---

[8] For various functional forms relating variables with range 0 to 1, see Taagepera (2008).

[9] As further described below, the model makes rather unintuitive predictions about a non-effect on decision-making duration of differences in preference constellations. A large exponent ensures that even small differences in policy result in large changes in the probability to change position. Thus, such a calibration maximizes the potential for policy similarity to have an effect on duration and the chosen exponent value is quite conservative in this respect.

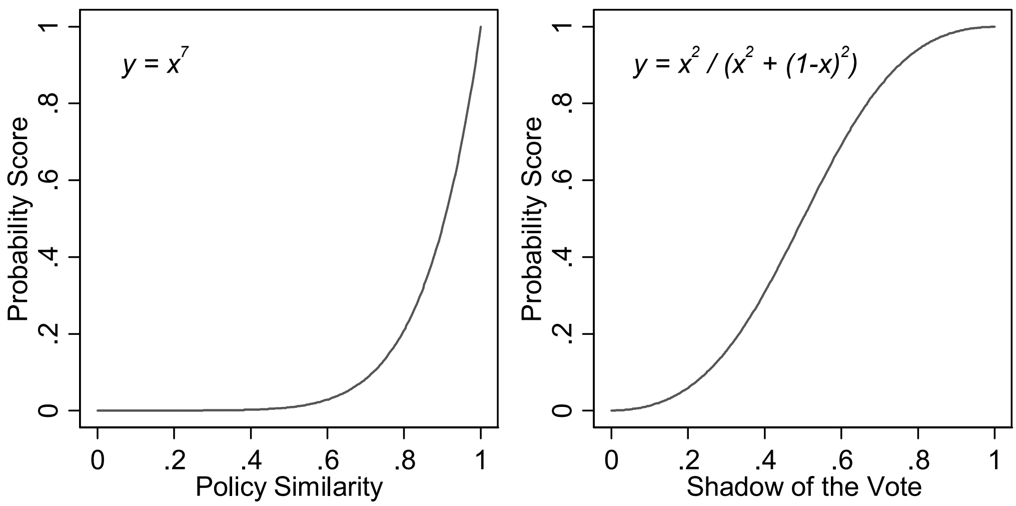

Figure 1 Mapping of variable values onto probability scores

Left panel: Small decreases in policy similarity values result in relatively large decreases in probability scores. Right panel: Changes in the shadow of the vote values have disproportionately small effects on probability scores at the lower end of the shadow of the vote scale and disproportionately large effects at the higher end of the scale.

Given that the probability scores for all three variables range between 0 and 1, the overall probability of a negotiator to change position can be determined through a fuzzy set aggregation rule (Goertz 2006; Ragin 2000):

$$Probability = \max\big(shadow\ of\ the\ vote, \min(impatience, policy\ similarity)\big)$$

To determine a negotiator's position change probability, we first determine the minimum of her individual-level variable values. In fuzzy set logic, the minimum operation aggregates necessary conditions. According to this part of the aggregation rule, both impatience and policy similarity are considered necessary for a high position change probability. Conversely, either low impatience or policy dissimilarity is sufficient to prevent negotiators from changing position. The next step of the aggregation rule consists of finding the maximum of the shadow of the vote value and of the minimum value of the individual-level variables. In fuzzy set logic, the maximum operation aggregates sufficient condition. Thus, a strong shadow of the vote is sufficient for a high position change probability and can override patience and concerns about policy dissimilarity. At the same time, impatience and policy similarity are jointly sufficient as well. If both impatience and policy similarity are high, the position change probability will be high as well.

Illustration of model dynamics

Figure 2 illustrates the general dynamics of the model. For ease of exposition, the figure presents a somewhat unusually short simulation run that lasted only five negotiation rounds. In the model, coalition-building takes place in a two-dimensional policy space, which is represented as a 100 x 100 square lattice. Each panel presents a snapshot of negotiators' positions on that square lattice. The upper left panel plots the initial distribution of negotiators' positions, which are assumed to correspond to their ideal points. At the start of the simulation, the number of negotiators is specified and their ideal points are randomly distributed over the policy space according to a beta distribution. For the current example, the number of negotiators was set to 27 and the shape parameters of the beta distribution were set to ensure that their ideal points had a uniform probability of being located anywhere in the policy space, which results in a rather heterogeneous initial dispersal of ideal points. The voting threshold for all simulations reported in this paper was set to 72 per cent, which corresponds closely to the real-world voting threshold in the Council.

In each round of the model, negotiators have the opportunity to join the position of an alternative coalition. Negotiator will move to an alternative coalition with a probability equal to their position change probability if the alternative coalition has at least as many members as the current one. Negotiators act sequentially according to a random schedule in each round to avoid artificial order effects. The numbered arrows in the panels indicate the destination and sequence of moves of negotiators that led to the configuration of positions depicted in the next panel. Individual- and system-level variables are continuously updated, which means that negotiators do not base their decision to move to another coalition on the state of the world at the beginning of the negotiation round, but on the state of the world at the point in time during the negotiations round when it is their turn to decide about changing positions. This means the panels provide a depiction of the relevant decision-making situation only until the first negotiator has moved as part of the next negotiation round.

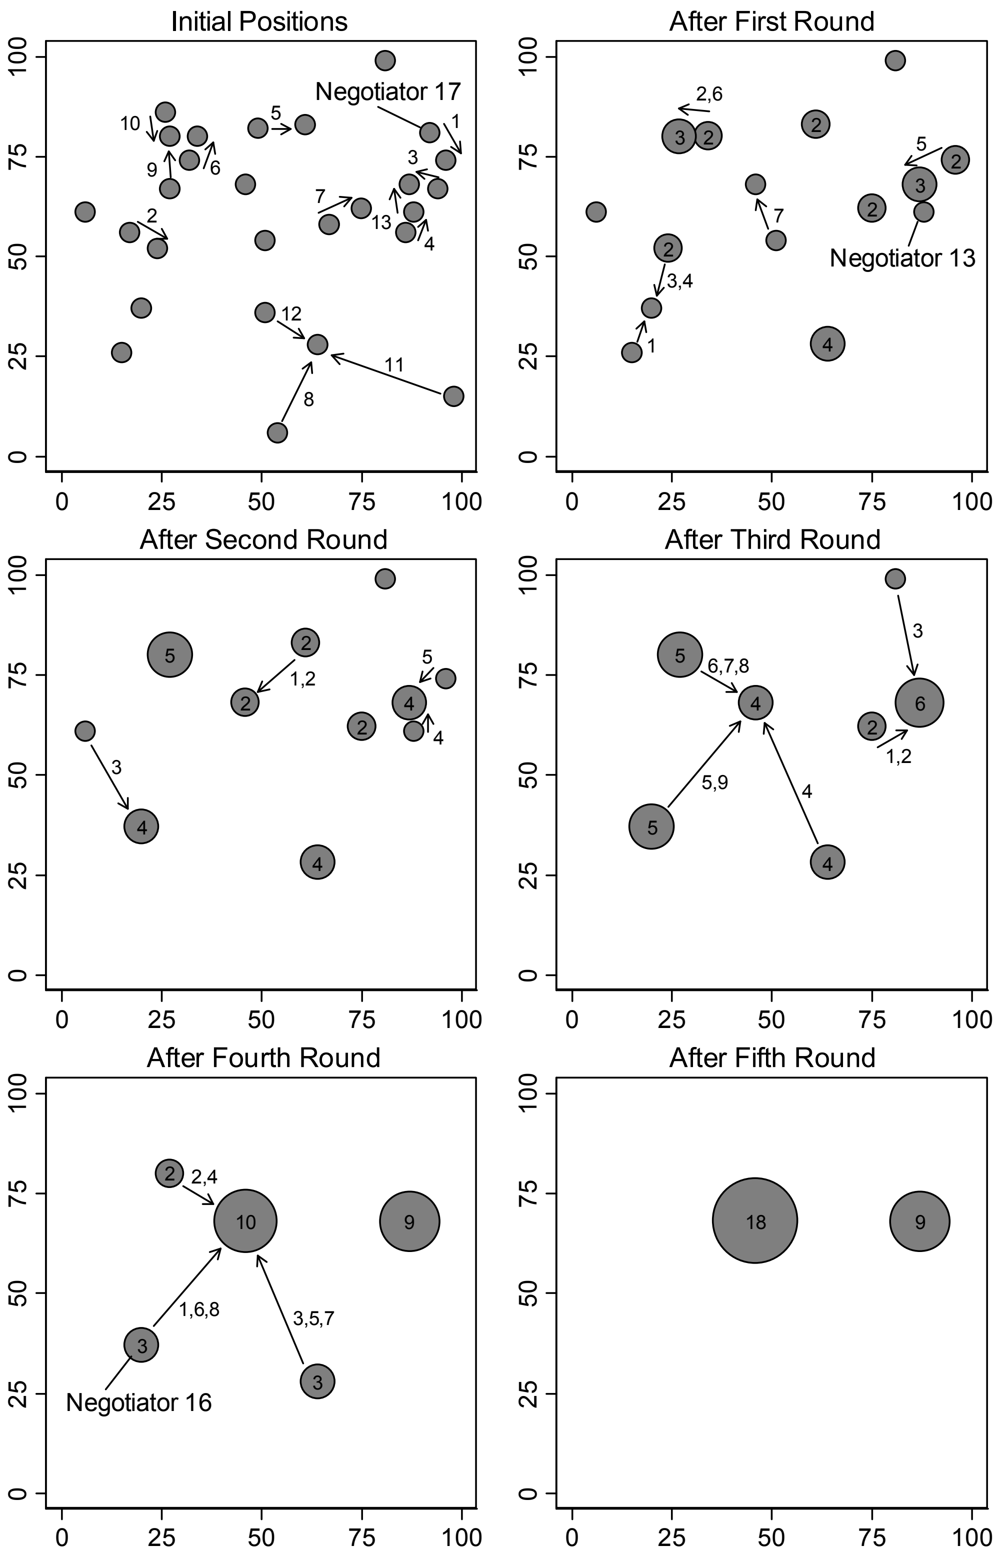

Figure 2 Illustration of model dynamics

The panels provide snapshots of the positions of negotiators in the policy space after initialization of the model and subsequent negotiation rounds. The arrows indicate changes in positions negotiators make in the next negotiation round and the accompanying numbers the sequence in which those moves occur. The size of the circles increases with the number of negotiators occupying that position. The number in the circle provides their precise number.

Although this means we cannot follow a specific individual and trace all its decisions through the entire negotiation process, we can illustrate the decision-making calculus of negotiators by considering the decision-making of the negotiator who moves first in each negotiation round. This negotiator bases its decision on the position configuration as depicted in the respective panel. For example, after the initialization of the model, negotiator 17 is randomly chosen to act first. Negotiator 17 is not in a blocking coalition yet, so it surveys the position of other negotiators and identifies the position of negotiator 8 as the closest one. Furthermore, this 'coalition' is at least as large as its own, which means it is a viable alternative. Whether or not negotiator 17 joins negotiator 8 depends on the former's position change probability. When the model was initialized, negotiator 17's impatience value was randomly drawn from a beta distribution with a somewhat left-skewed shape, producing generally rather high impatience values. Negotiator 17's impatience value was randomly set to about 0.62. The policy distance to negotiator 8 is quite modest, resulting in an even higher policy similarity probability score of 0.66. Finally, after surveying the positions of other negotiators, negotiator 17 has also recognized that no negotiators have formed coalitions yet, resulting in a probability score for the shadow of the vote of zero.

Either patience or a lack of policy similarity is sufficient to make a negotiator less likely to move to another position, therefore the minimum of the impatience and the policy similarity scores are first determined. In this case, the impatience score is smaller than the policy similarity score, resulting in a minimum value of 0.62. In the next step, the maximum of this minimum value and the shadow of the vote value is identified. High shadow of the vote scores increase the position change probability, but lower ones do not reduce it. As the shadow of the vote is zero, the maximum value and therefore the position change probability is 0.62. Technically, a random draw from a uniform distribution with support 0 to 1 now determines whether negotiator 17 moves to the alternative coalition or stays at its current position. In this simulation run, the random draw resulted in a number smaller than or equal to 0.62, so the negotiator moved. Taking into account the new state of the world after negotiator 17 joined the position of negotiator 8, the next negotiator performs the same calculations. Once each negotiator had its turn, the current negotiation round ends.

The top right panel provides the configuration of policy positions after the first round. Randomly determined, negotiator 13 is now the first actor to decide about joining another coalition. The closest coalition consists of three negotiators, which is clearly larger than negotiator 13's one-member 'coalition'. Negotiator 13 has an exogenously determined impatience score of 0.60. A policy similarity score of 0.54 indicates that the alternative coalition is quite far away from negotiator 13's original ideal position. At this stage, the largest coalition consists of four negotiators, generating a negligible shadow of the vote of less than 0.04. Given our fuzzy set decision-making rule, the concession probability is thus max(0.04, min(0.60, 0.54)) = 0.54. The shadow of the vote has no effect yet, but the relatively low policy similarity value leads to a lower concession probability despite the somewhat higher impatience score. In this case, the random draw results in a value larger than 0.54 and negotiator 13 remains at its current position. For the remainder of this and subsequent negotiation rounds, all negotiators not in a blocking coalition apply this decision-making rule.

In general, the shadow of the vote is small in early negotiation rounds but increases constantly over time. In the example presented here, the shadow of the vote increases from zero after initialization to 0.04 after the first round, 0.07 after the second round, 0.11 after the third round, and 1.00 after the fourth round. At the end of round three, the largest coalition consists of only six negotiators, still smaller than the required blocking coalition threshold of eight and far off the majority coalition threshold of 20. However, during round four, two blocking coalitions of size nine and ten form rapidly. The size of the largest non-blocking coalition is three. If the members of this coalition join one of the blocking coalitions, they could adopt a compromise agreement. Thus, the shadow of the vote reaches its maximum at this point in time and overrides patience and policy cost considerations. For example, the first negotiator to consider moving to another coalition at the beginning of the fifth round is negotiator 16. Negotiator 16 has an innate impatience score of 0.73. However, the alternative coalition with 10 current members is rather far away, resulting in a low policy similarity score of 0.09. Because of the strong policy differences, the position change probability would be extremely low in the absence of the shadow of the vote. But given the imminent possibility of being outvoted, the shadow of the vote score is 1.00, resulting in a position change probability of 1.00, making sure that negotiator 13 moves to the alternative coalition as soon as possible. The same is true for the other negotiators in round five that are not in blocking coalitions yet.

If the coalition-building behaviour results in a majority size coalition or a compromise among a sufficiently large number of negotiators in blocking coalitions after a negotiation round, the simulation stops. The lower left panel depicting positions after the fifth negotiation round illustrates the latter situation. Here, both coalitions are of blocking minority size, but neither is large enough to unilaterally adopt a decision by vote. The only way a decision can be adopted is by reaching a compromise agreement between the two coalitions. How this compromise agreement is reached is not modelled explicitly. The underlying idea is that compromise building in the endgame of negotiations plays a relatively negligible role for explaining decision-making duration compared to the coalition-building process that precedes it.

Computational experiments

In contrast to classical mathematical models that use comparative statics analysis to derive hypotheses analytically, agent-based and other computational models rely on numerical simulation. Such computational experiments identify the effects of variables by observing simulation outcomes under varying treatment conditions. In contrast to real experiments, all other things can be kept constant in computer simulations to unambiguously identify the 'causal' role that a particular variable plays in the model. The theoretical literature on bargaining and collective decision-making often stresses the role of policy conflict and impatience for the duration of negotiations. Thus, the following experiments generate hypotheses about the effects of different preference configurations and impatience distributions on decision-making duration. The effect of group size on decision-making duration is an especially interesting and contested topic in the EU context, where the number of member states has increased continuously over time through various enlargement rounds and is likely to continue to do so at least in the near future. Therefore, group size is the third factor investigated here.

With respect to the effect of preference configurations, two important views can be distinguished. As discussed earlier, spatial theories of voting lead to the expectation that preference heterogeneity inhibits policy change (Tsebelis 2002). In contrast, bargaining theories argue that the polarization of preferences leads to more political conflict (Esteban and Schneider 2008). However, none of these bodies of theories makes direct predictions about decision-making duration. In the following, I examine whether the agent-based model also generates these hypotheses. When the model is initialized, the locations of negotiators' ideal points are randomly generated from a beta distribution with shape parameters α and β. The beta distribution has a symmetric shape when $\alpha = \beta$, is right-skewed when $\alpha > \beta$ and left-skewed when $\alpha > \beta$. As a special case, the beta distribution reduces to the uniform distribution when $\alpha = \beta = 1$. Different preference configurations can thus be generated by setting the shape parameters of the beta distribution to different values.

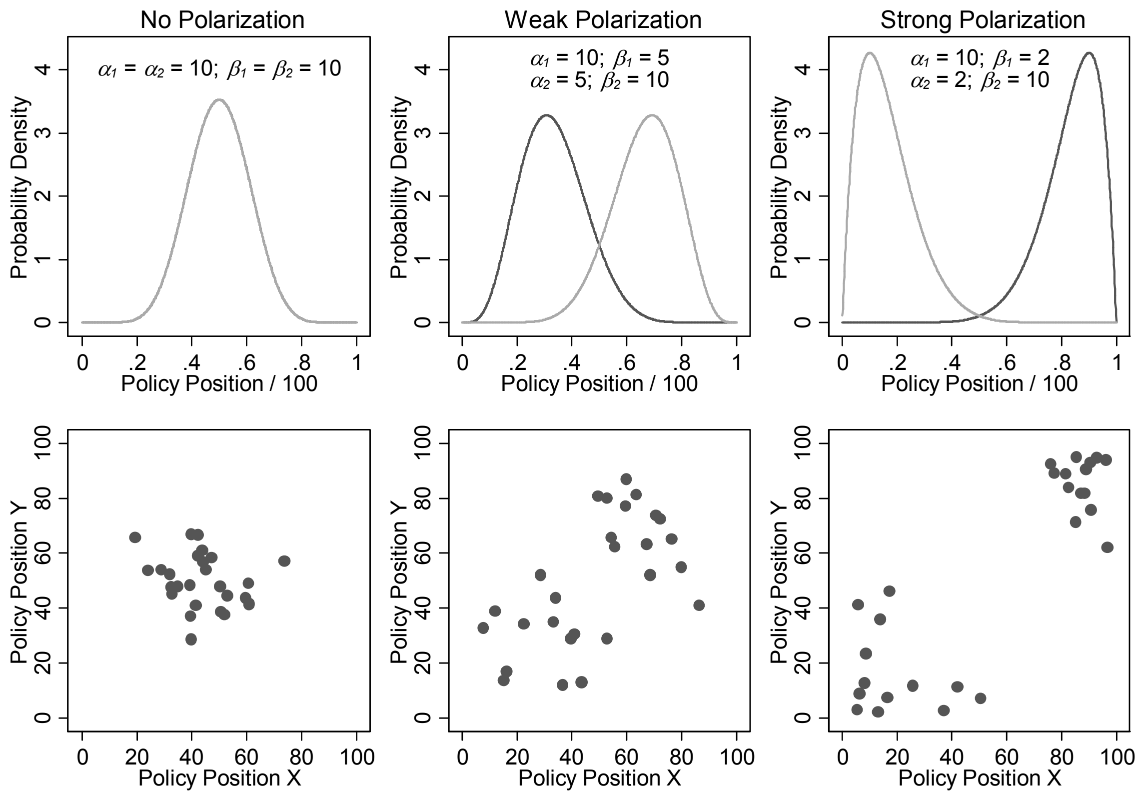

Figure 3 Different degrees of preference polarization

This figure illustrates the different preference polarization conditions examined in the computational experiment. The first row of panels shows the shape of the probability density distributions from which ideal points are drawn at the start of the simulation. The distributions are beta distributions, whose shape is determined by two parameters α and β. The parameters for each of the two groups' distributions are given in the respective panel. The second row of panel provides illustrative examples of the resulting preference configurations.

Polarization is the degree to which negotiators form groupings with relatively distinct policy positions. Polarization is larger the more inter-group differences in policy positions outweigh intra-group differences. For simplicity, I only consider the case where negotiators form two groupings, implying that preferences on the two dimensions are correlated. If polarization affects decision-making speed, this should become most obvious in the most extreme case of bipolar group differences. At the start of the simulation, negotiators are randomly divided into two groups. For each group, ideal points are drawn from a separate beta distribution. I distinguish three scenarios as shown in Figure 3. The first row of panels in the figure plots the shapes of the probability density distributions and the second row provides example preference configurations produced by those distributions. The 'no polarization' scenario sets both beta distribution parameters to a value of 10, yielding a centrist symmetric distribution that is the same for the two groups. This scenario is identical to assuming that no group differences exist. The ‘weak polarization’ scenario reduces the second parameter of group 1 and the first parameter of group 2 to a value of 5, resulting in slightly right- and left-skewed preference distributions, respectively. Finally, the ‘strong polarization’ scenario reduces these parameters to a value of 2, further increasing the skew of the preference distributions and the overall polarization.

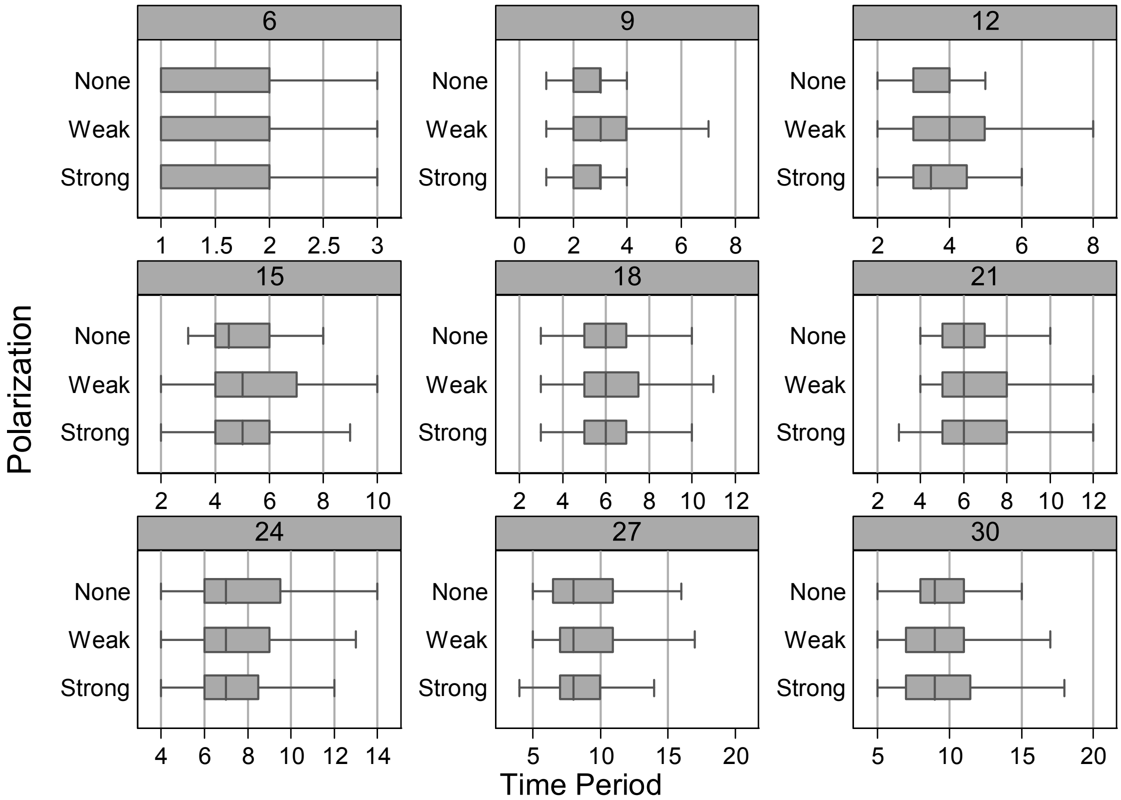

Figure 4 Effect of preference polarization on duration by group size

The figure compares decision-making duration across preference polarization conditions for different number of member states. For the reported simulations, impatience values were randomly drawn from a uniform distribution. Each box plot is based on 100 simulation runs with the same specification of conditions but varying random seed. Outlying observations were omitted from the figure to increase readability.

Figure 4 reports the effect of polarization on decision-making duration for different group sizes. The figure is based on model runs with impatience values drawn from a uniform distribution, but results with different distributional shapes yield similar results. Each box plot is based on 100 simulation runs with the same specification of conditions but different random seed. The main message from this figure is that preference polarization does not have any effect on decision-making duration in the model. At first sight, this result seems counterintuitive. However, if negotiators are divided into two roughly equally large groupings, the policy concessions they have to make to form blocking coalitions are actually relatively small; and once all negotiators have formed blocking coalitions, they strike a compromise without delay, regardless of the distance of inter-group differences. Thus, the first empirically testable hypothesis of the model can be stated as follows:

H1: Changes in preference polarization do not lead to changes in decision-making duration.

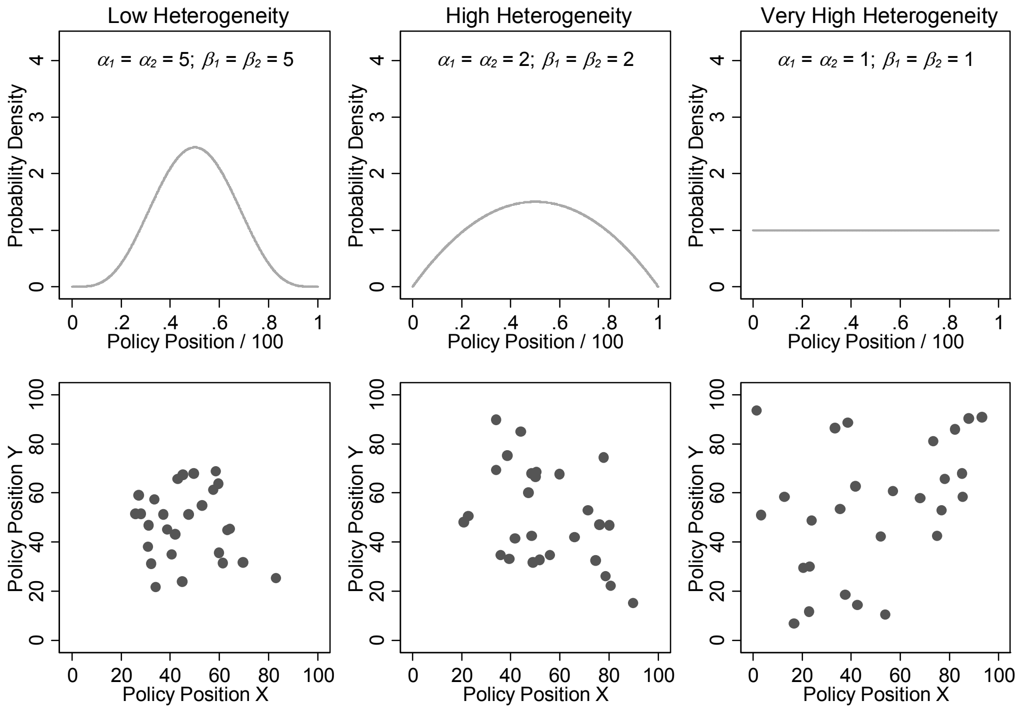

Figure 5 Different degrees of preference heterogeneity

This figure illustrates the different preference heterogeneity conditions examined in the computational experiment. The first row of panels shows the shape of the probability density distributions from which ideal points are drawn at the start of the simulation. The distributions are beta distributions, whose shape is determined by two parameters α and β. The parameters for each of the two groups' distributions are given in the respective panel. The second row of panel provides illustrative examples of the resulting preference configurations.

Preference heterogeneity is the extent to which negotiators' ideal points are dispersed across the entire policy space. To generate varying degrees of preference heterogeneity, the two shape parameters of the beta distributions are set equal to each other but varied in size. Drawing ideal points of both groups from distributions with the same shape parameters is equivalent to drawing all ideal points from the same distribution. As mentioned above, setting both distribution parameters to a value of 1 results in a uniform distribution, with ideal points having a constant probability of being located anywhere in the policy space. Preference constellations drawn from the uniform distribution constitute a situation of high preference heterogeneity. Increasing the values of the shape parameters decreases preference heterogeneity. Besides a value of 1, values of 2, 5, and 10 are considered in the experiment.

The shapes of the probability density distributions and examples of the resulting preference configurations are provided in Figure 5.[10]

Figure 6 presents the results of the experiment. The plots show the effect of preference heterogeneity on decision-making duration for different group sizes. The presented results are again based on impatience values drawn from a uniform distribution, but results with other impatience distributions yield the same conclusions. In contrast to preference polarization, the heterogeneity of preferences increases the duration of decision-making somewhat, although the effect size is very modest. This finding leads to the second hypothesis:

H2: Increases in preference heterogeneity lead to only minor increases in decision-making duration.

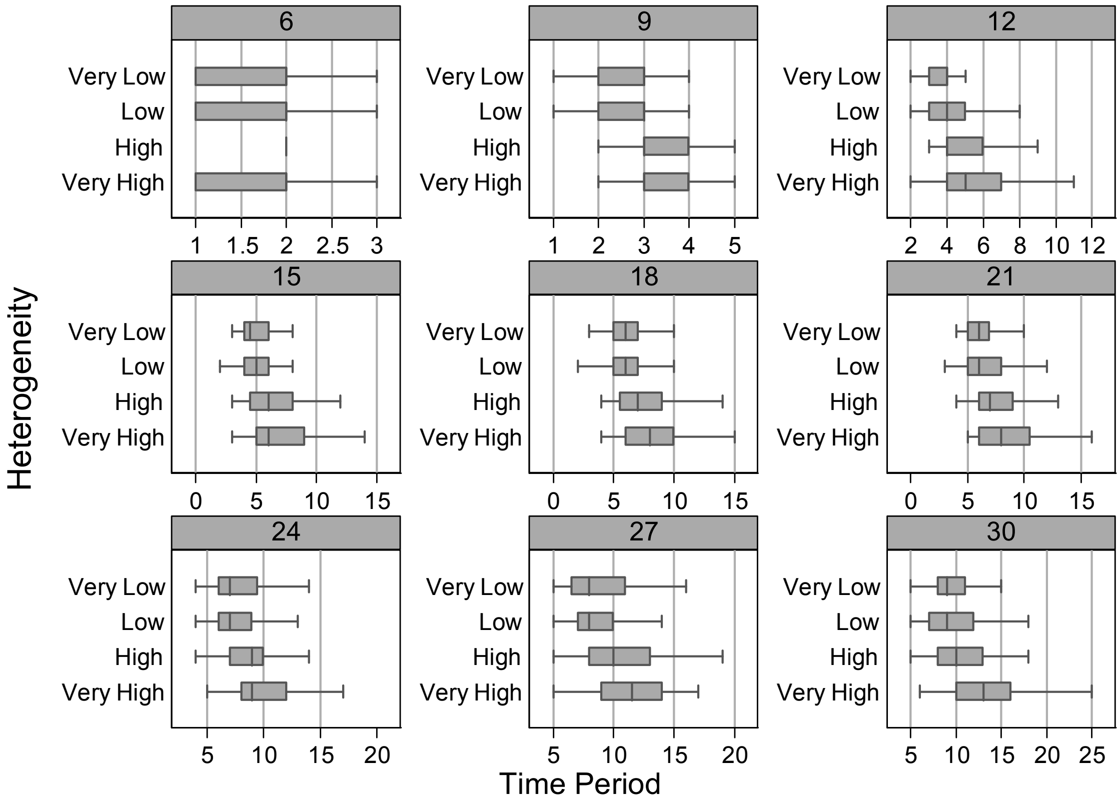

Figure 6 Effect of preference heterogeneity on duration by group size

The figure compares decision-making duration across preference heterogeneity conditions for different number of member states. For the reported simulations, impatience values were randomly drawn from a uniform distribution. Each box plot is based on 100 simulation runs with the same specification of conditions but varying random seed. Outlying observations were omitted from the figure to increase readability.

[10] For reasons of space, the 'Very Low Heterogeneity' condition is not shown in Figure 5. This condition is the same as the 'No Polarization' condition illustrated in the left panel of Figure 3.

The third factor investigated is the general level of impatience of negotiators. Rather than investigating whether individual differences in impatience lead to different success rates for individual negotiators, the experiment investigates whether differences in the central tendency of the impatience distribution affects the duration of decision-making. At the start of each model run, individual impatience values are randomly drawn from a beta distribution. To see whether it matters for decision-making duration if negotiators are generally more or less impatient, the shape parameters of the beta distribution can be varied. Figure 7 presents the different impatience distributions used in the experiment. The seven conditions range from 'Extremely Low' ($\alpha = 1;\ \beta = 10$) over 'Medium' ($\alpha = 10;\ \beta = 10$) to 'Extremely High' ($\alpha = 10;\ \beta = 1$).

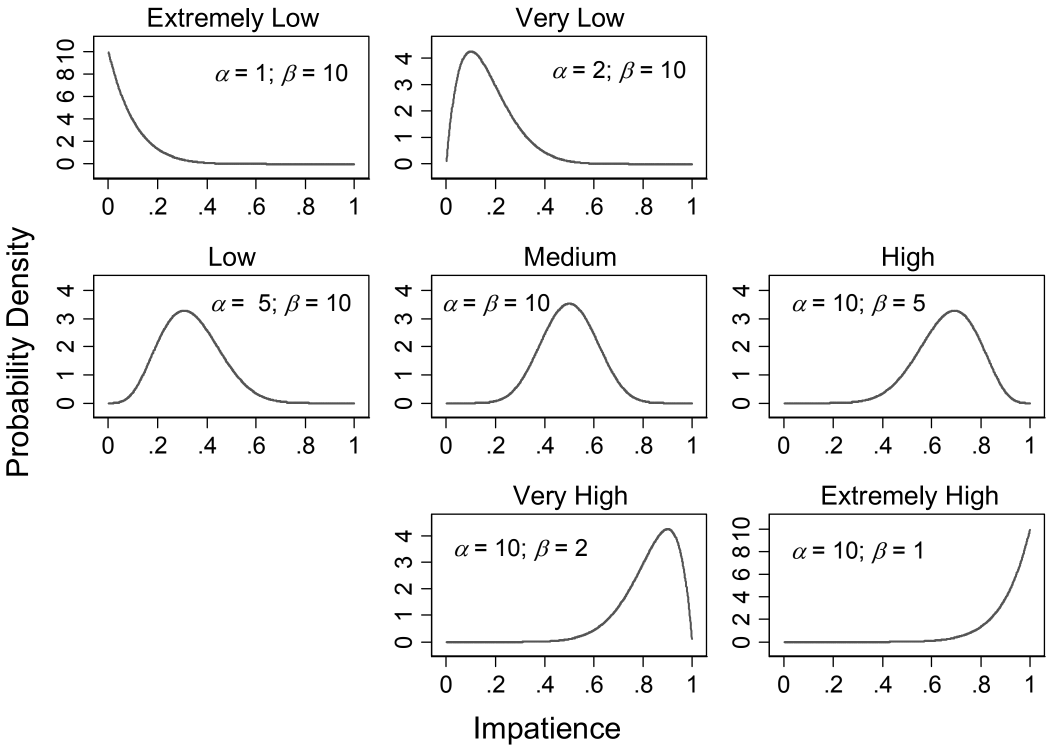

Figure 7 Different degrees of impatience

This figure illustrates the different impatience level conditions examined in the computational experiment. The panels show the shape of the probability density distributions from which impatience values are drawn at the start of the simulation. The distributions are beta distributions, whose shape is determined by two parameters α and β. The precise parameters for each distribution are given in the respective panel. Parameter values were chosen to reflect a wide range of 'typical' impatience values, from extremely low (upper left panel) to extremely high (lower right panel).

Figure 8 presents the experimental results for the effect of impatience on decision-making duration for different group sizes. In all reported scenarios, negotiators' ideal points were drawn from a uniform distribution, but the results are robust to alternative distributional forms. The figure clearly demonstrates that for any number of member states the general level of impatience has a major effect on decision-making speed. This result has a quite intuitive explanation. In the model, the duration of decision-making depends crucially on how quickly negotiators form blocking coalitions. Thus, changes in duration in response to changes in the typical impatience of negotiators are a straightforward consequence of all negotiators being prepared to change positions more readily. This finding yields the third hypothesis:

H3: Increases in the general level of negotiators' impatience lead to decreases in decision-making duration

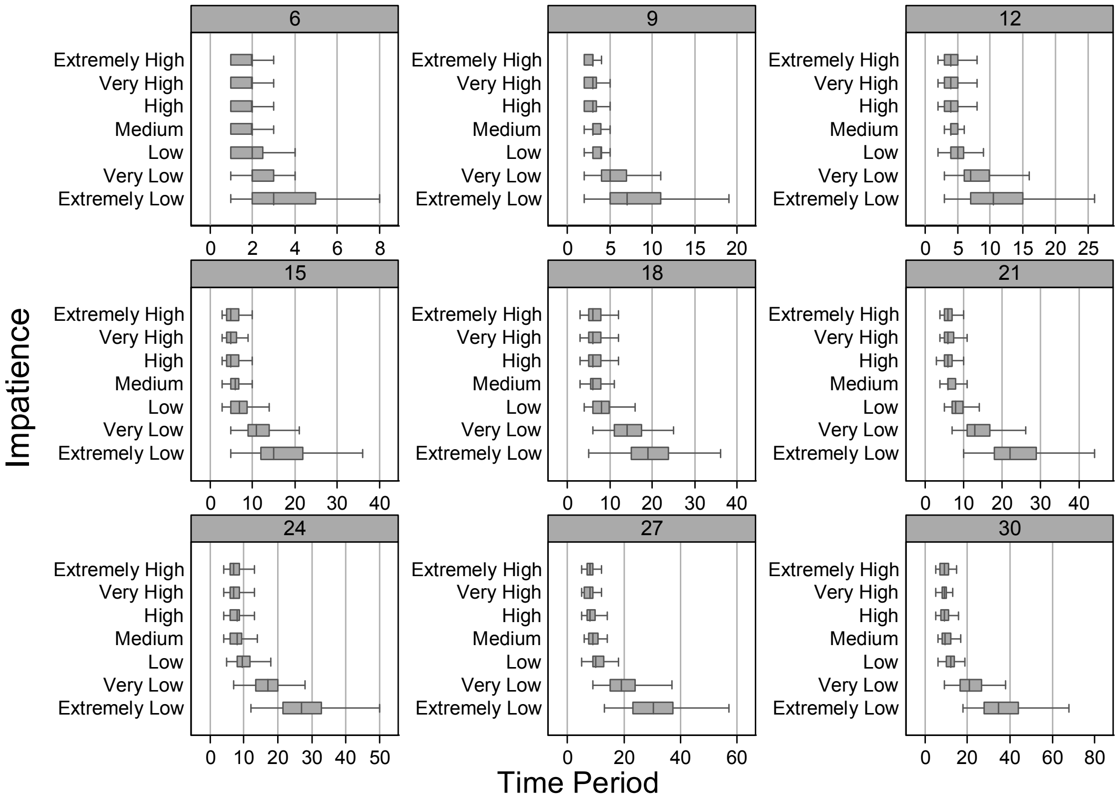

Figure 8 Effect of impatience on duration by group size

The figure compares decision-making duration across impatience conditions for different number of member states. For the reported simulations, the location values of negotiators' ideal points were randomly drawn from a uniform distribution. Each box plot is based on 100 simulation runs with the same specification of conditions but varying random seed. Outlying observations are omitted from the figure to increase readability.

Although the negative effect of impatience on duration is discernible regardless of the number of member states, comparison between panels in Figure 8 also indicate differences in duration between different numbers of member states. Figure 9 presents the effect of group size on decision-making duration more clearly. The figure shows a strong, almost linear positive relationship that holds regardless of negotiators general level of impatience. Again, in the simulations underlying this figure, negotiators' ideal points were drawn from a uniform distribution, but the general pattern holds for other distributional forms as well. Thus, the final hypothesis reads as follows:

H4: Increases in the number of member states lead to increases in decision-making duration.

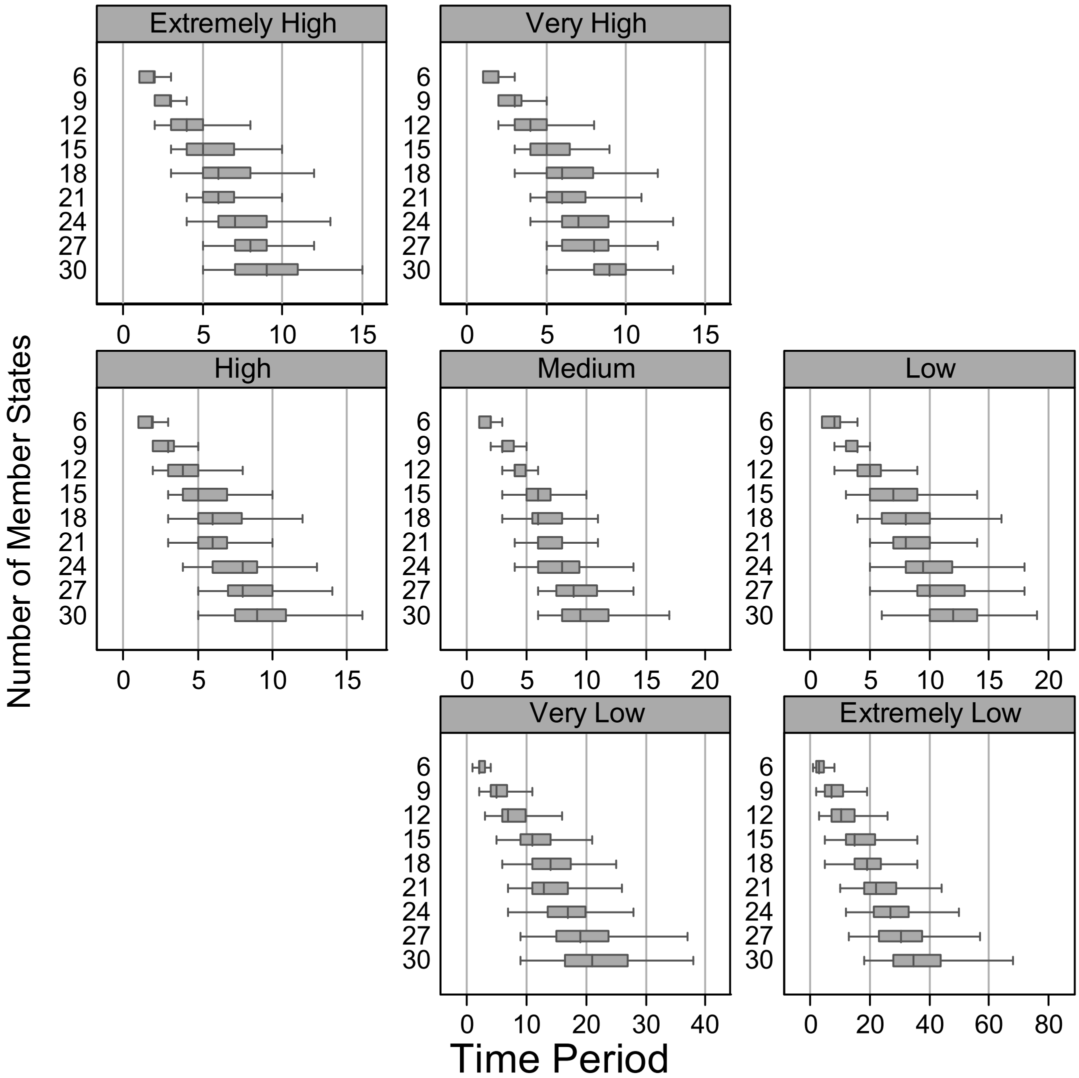

Figure 9 Effect of group size on duration by impatience

The figure compares decision-making duration across group size conditions for different impatience values. For the reported simulations, the location values of negotiators' ideal points were randomly drawn from a uniform distribution. Each box plot is based on 100 simulation runs with the same specification of conditions but varying random seed. Outlying observations are omitted from the figure to increase readability.

The extended duration of decision-making when the group size increases is an almost mechanical outcome of a larger coalition size required to block a decision. In the model, coalition-building is a sequential process in which negotiators form larger and larger coalitions until they reach blocking minority size. This process will be completed faster when the coalition size required to block a decision is smaller and take longer if the required coalition size is larger.

Conclusion

The ability to respond to economic and social problems in a timely fashion is an important feature of a political system. Delayed policy responses can create considerable costs for affected constituents. Furthermore, citizens have a justified expectation that law-makers adopt policies efficiently, using the least amount of resources required. In this respect, legislative paralysis has a negative effect on citizens' perceptions of the legitimacy of political decision-making. In line with the importance of this factor for evaluations of the functioning of the EU political system, decision-making duration has received considerable attention in the existing literature. Past studies have produced a number of robust empirical findings using increasingly sophisticated statistical methods and datasets with ever expanding coverage. Unfortunately, the theoretical underpinnings of these empirical findings are somewhat less clear-cut. Existing research has relied on the spatial theory of voting and coalition theory to generate hypothesis about the length of decision-making. However, these formal theories make predictions about policy stability, not the duration of decision-making.

As a step towards improving our understanding of the causal mechanisms underlying the duration of decision-making in the Council, this paper presents an agent-based model of the coalition-building behaviour of negotiators. The model is founded on the idea that negotiators strive to form blocking minority coalitions to ensure that their views are taken into account when a compromise agreement is negotiated in the Council. Negotiators join the positions of other negotiators with similar views until their group is large enough to block a qualified majority vote. Computational experiments based on this model provide a number of predictions consistent with existing empirical findings. First, decision-making duration increases strongly with decreases in the general level of impatience of negotiators. Politicization is likely to result in negotiators being less willing to make policy concessions quickly. Empirical studies have shown that decision-making under codecision results in both more politicized Council negotiations (Häge 2011b) and longer decision-making processes (Häge 2011a).[11] The agent-based model provides a theoretical rationale for this finding. The

[11] Note that this finding is different from the results of other existing studies in that the dependent variable is the duration until the first Council decision rather than the duration until the conclusion of the entire legislative process.

involvement of the EP increases the politicization of policy issues, and the stronger public scrutiny makes negotiators in turn more reluctant to make concessions.

Second, decision-making duration increases with the number of member states. This result is consistent with the latest findings about the effect of enlargement by Hertz and Leuffen (2011), which corroborates most earlier empirical work. According to the agent-based model, the increase in the duration of decision-making brought about by a larger number of member states has a simple, almost mechanical explanation. A larger number of member states means that a larger number of states need to come together in order to form a blocking minority. The step-wise process of forming larger and larger groups to constitute a blocking minority takes longer if the required number of member states is larger.

While these two hypotheses are in line with the theoretical expectations and results of most existing research, some more counterintuitive findings of the computational experiments indicate that preference configurations play at best a minor role for explaining the length of decision-making. Different degrees of preference polarization have no effect on duration at all, and different degrees of preference heterogeneity have a very modest one compared to levels of impatience and group size. This result is somewhat remarkable, as the probability to change position and join another negotiator in the model depends strongly on the distance of that negotiator's position from the ideal point of the focal individual. The findings indicate that, even in cases where the average policy distance between negotiators is large, negotiators are often able to form blocking minorities with close-by negotiators without the need to traverse large areas of the policy space. The non-effect hypothesis poses a clear alternative to existing theoretical arguments, which expect preference heterogeneity to be a major determinant of decision-making duration. Although past research has found empirical support for the effect of preference heterogeneity, these studies have operationalized the concept either very indirectly or at a higher level of aggregation, making their findings somewhat suspect. A more appropriate test to compare the predictive power of alternative theoretical models would involve detailed case-specific preference data at the proposal or even issue level. Such competitive model testing is a promising avenue for future research.

While the model makes a theoretical contribution by explicitly formulating the causal mechanisms through which decision-making is faster or slower, we should also be clear about its limitations. In its current form, the model is mainly applicable to win-win situations, in which 'gains from trades' can be made in negotiations. Such a situation underlies many instances of multilateral international negotiations; otherwise states would not engage in negotiations in the first place. However, in everyday EU legislative decision-making, member states do not have this opt-out opportunity. Sometimes, a group of member states sufficiently large to form a blocking minority might find themselves in a situation in which they prefer the existing status quo over any possible negotiation outcome. The model clearly does not capture such situations. At the same time, the low failure rate of proposals of about 10 per cent (Häge 2011a) indicates that such situations are the exception rather than the rule. The goal of the model is not to capture the characteristics of all possible types of negotiation situations, but the essential features of a typical one.

Also, the model is confined to decision-making under some form of majority rule. Coalition building is of secondary importance when decisions need to be adopted by unanimous consent. A major advantage of coalition building is that it reduces the complexity of the negotiation situation considerably. Under unanimity, where each member state forms its own blocking minority, uncertainties about individual negotiators' preferences and reversion points multiply, making it hard to find a generally acceptable compromise agreement. These uncertainties are likely to play a much larger role for explaining decision-making duration under unanimity than majority vote. Ideally, a general theory of the duration of multilateral negotiations should be applicable to both unanimity and majority voting situations, but such generality will most likely come at the cost of a far more complex model. Arguably, if the causal mechanisms operating in different domains differ strongly, more can be learned from distinct, simple domain-specific models than from a more general but highly complex model that needs to specify every contingency.

References

Esteban, J., and G. Schneider. 2008. Polarization and Conflict: Theoretical and Empirical Issues. *Journal of Peace Research* 45 (2): 131-141.

Goertz, Gary. 2006. *Social Science Concepts: A User's Guide* Princeton: Princeton University Press.

Golub, Jonathan. 1999. In the Shadow of the Vote? Decision Making in the European Community. *International Organization* 53 (4): 733-764.

Golub, Jonathan. 2007. Survival Analysis and European Union Decision-Making. *European Union Politics* 8 (2): 155-179.

Golub, Jonathan. 2008. The Study of Decision-Making Speed in the European Union. *European Union Politics*: 167-179.

Golub, Jonathan, and Bernard Steunenberg. 2007. How Time Affects EU Decision-Making *European Union Politics* 8 (4): 555-566.

Häge, Frank M. 2010. "Coalition Building and Consensus in International Negotiations." Paper presented at the Pan-European Conference on EU Politics, 23-26 June, Porto.

Häge, Frank M. 2011a. The European Union Policy-Making Dataset. *European Union Politics*: forthcoming.

Häge, Frank M. 2011b. Politicising Council Decision-Making: The Effect of European Parliament Empowerment. *West European Politics* 34 (1): 18 - 47.

Hertz, Robin, and Dirk Leuffen. 2011. Too Big to Run? Analysing the Impact of Enlargement on the Speed of EU Decision-Making. *European Union Politics* 12 (2): 193-215.

König, Thomas. 2007. Divergence or Convergence? From Ever-Growing to Ever-Slowing European Legislative Decision Making. *European Journal of Political Research* 46 (3): 417-444.

König, Thomas. 2008. Analysing the Process of EU Legislative Decision-Making: To Make a Long Story Short... *European Union Politics* 9 (1): 145-165.

Ragin, Charles C. 2000. *Fuzzy Set Social Science*. Chicago: University of Chicago Press.

Schulz, Heiner, and Thomas König. 2000. Institutional Reform and Decision-Making Efficiency in the European Union. *American Journal of Political Science* 44 (4): 653-666.

Taagepera, Rein. 2008. *Making Social Sciences More Scientific: The Need for Predictive Models*. Oxford: Oxford University Press.

Tsebelis, George. 2002. *Veto Players: How Political Institutions Work*. Princeton: Princeton University Press.

Wilensky, Uri. 1999. Netlogo 4.1. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston.