

Diplomatic Relations in a Virtual World

Aaron Bramson^{1,2}, Kevin Hoefman^{2,3}, Koen Schoors² and Jan Ryckebusch⁴

¹ *Laboratory for Symbolic Cognitive Development, RIKEN Center for Biosystems Dynamics Research, 6-7-3 Minatojima-Minamimachi, Chuo-ku, Kobe 650-0047, Japan*

² *Department of General Economics, Ghent University, Tweekerkenstraat 2, 9000 Ghent, Belgium*

³ *Howest University College West Flanders, Marksesteenweg 58, 8500 Kortrijk, Belgium*

⁴ *Department of Physics and Astronomy, Ghent University, Proeftuinstraat 86, 9000 Ghent, Belgium*

Abstract

We apply variations and extensions of structural balance theory to analyze the dynamics of geopolitical relations using data from the virtual world *Eve Online*. The highly detailed data enable us to study the interplay of alliance size, power, and geographic proximity on the prevalence and conditional behavior of triads built from empirical political alliances. Through our analysis, we reveal the degree to which the behaviors of players conform to the predictions of structural balance theory and whether our augmentations of the theory improve these predictions. In addition to studying the time series of the proportions of triad types, we investigate the conditional changes in triad types and the formation of polarized political coalitions. We find that player behavior largely conforms to the predictions of a multipolar version of structural balance theory that separates strong and weak configurations of balanced and frustrated triads. The high degree of explanatory power of structural balance theory in this context provides strong support for both the theory and the use of virtual worlds in social science research.

Keywords: diplomatic relations, balance theory, coalition formation, virtual worlds

1 Introduction

In order to better understand the dynamics of real-world political interactions, we analyze the characteristics and dynamics that naturally emerge within a virtual world political environment. We first describe an expanded version of structural balance theory that we use as a guiding principle governing the dynamics of political relations. We then evaluate the degree to which the virtual world dynamics cohere to the predictions of this theory in multiple ways. We further describe how differences between the virtual and real world affect the translation of social theories between virtual and real-world environments.

Data regarding the stability and dynamics of real-world political relations are scarce, because the number of real-world countries is small and because political relations tend to be relatively stable over the span of history for which reliable data are available. This paucity of data on the dynamics of real-world political relations makes it nigh impossible to test general theoretical claims about the drivers of change in international relations. Datasets such as the Correlates of War (CorrelatesOfWar 2019) and the conflict in the middle east (Economist 2015), for example, are too sparse to provide strong support for general sociopolitical theories. Furthermore, changes in external factors, such as technology and economic activity, alter the context in which political changes occur, thus further complicating or even invalidating the analysis of real-world political relations across time and space. We address these problems by considering an alternative body of empirical data on political relations that does not suffer from these limitations. In particular, we employ data on political relations from a virtual world called *Eve Online* (EVE) to study the dynamics of political relations (CCP 2019).

The core of our analytical approach derives from Structural Balance Theory (Heider 1946; Cartwright and Harary 1956; Harary 1959). The principle of balance theory is that interpersonal

relations have valences—some people you like and some you dislike—and we can capture those as positive and negative links in a social network. Structural balance theory (sometimes “social balance theory” or simply “balance theory”) provides a characterization of when such a signed social network is balanced—and if not, then how frustrated it is. The implication is that frustrated relations are more likely to change than balanced ones, and so the theory implicitly also provides a theory about the dynamics of signed social networks. There are already many variations of this methodology, and we further extend it to explore patterns of frustration dynamics in political networks as they evolve.

The community structure of the *EVE* alliance standings network reveals multiple mutually competing coalitions. By treating this virtual world as a multipolar system engaged in structural balance dynamics, we can explain many patterns observed in the data. Deviations from the predictions of structural balance theory (like persistent strong frustration) reveal the limitations of applying the theory to a conflict-driven game context (Belaza *et al.* 2017; Belaza *et al.* 2019). Still, the high degree of explanatory power of structural balance theory in this context provides strong support for both the theory and the use of virtual worlds in social science research.

2 Data: *EVE* Virtual World

Appreciating the limited data available on real-world political and/or trade networks and large-scale conflicts, we explore the appropriateness of data collected from a massively multiplayer online game called *EVE*. The advantages of using data from an immersive, sandbox computer game (Squire 2008) are: (1) there are more data, (2) the data are reliable, (3) events in a computer game occur more frequently than in the real world, and (4) most of the activity is already quantified in a consistent way. One concern about data from a computer game is that it may lack sophistication, because the game world is necessarily a simplified version of the real world. Along these lines, one may worry that the behaviors of players reflect constraints and interventions from the developers, rather than natural diplomatic actions. We desire data that are simpler than the real-world data in some respects, yet still complicated in the right ways to be useful.

Few games provide the sociopolitical structure, variety of activities, and realistic motivations necessary to be candidates for serving this role. *EVE* does seem to provide all the right ingredients to generate realistic (and yet constrained) dynamics in large hierarchical political structures. Players of *EVE* engage in activities such as mining, harvesting, research, industry, trading, couriering, protection, piracy, and politics in an open-ended sandbox environment. Players are free to join corporations, which can form into alliances, which are in turn part of implicit coalitions in a multitiered hierarchy of political arrangements. Alliance membership determines friends and foes, which regulates conflicts on a scale ranging from single-player economic sabotage to prolonged territorial wars involving thousands of players.

Because we wish to use these data to understand why and how certain political, economic, and social activities occur, it is important to ensure that the underlying forces driving them—and not just aggregate patterns—match up with corresponding real-world motivations. The political actions we examine are generated by human players reacting to each other, and to an environment with fixed and known rules. The game offers players a wide range of economic activities (Mildenberger 2013). Because these activities are connected in a complex interplay of dependencies, players inside the game are motivated in similar ways as those real-world actors whose behavior we wish to study. Players invest time and real-world money to acquire virtual resources in the game world. These resources are then used to enhance the players’ capabilities, influence, and/or status. When players lose these resources in conflict or through mistake, the time, effort, and money they used to acquire them are permanently lost. This feature creates genuine scarcity and risk aversion that fuels realistic economic behaviors (BBCNews 2014; Goh 2018; Hoefman *et al.* 2019).

The constraints imposed by a virtual world are also a blessing. Although the scale of the game is large, with hundreds of thousands of players, it is still manageable. For example, alliances set sales taxes for their stations and income taxes for their members, but there is just a single rate for each, rather than an encyclopedic tax code. Political standings between pairs of alliances are captured by a single value between -10 and +10, rather than a complicated mixture of treaties, international laws, and historical conventions. All the necessary ingredients for complex sociopoliticoeconomic dynamics exist, but they are all simplified. A virtual world is like a computer simulation, except with humans instead of algorithms as the driving force. What we get is a “frictionless pulley” type of environment in which to test social and economic theories without the noise, confounding factors and size limitations that plague real-world datasets. Below we highlight some relevant aspects of EVE, with additional details in the Supplementary Material.

2.1 Players

There are between 400,000 and 500,000 accounts created for playing *EVE*. On average, 33,000 accounts are signed on simultaneously during our analysis period ([Eve-Offline 2018](#)).

2.2 Corporations

All characters belong to a corporation, some of which are default nonplayer character corporations run by the game itself, and others which are owned and run by players. Players can create and close corporations, so the number of corporations and their memberships change over time; however, there are consistently around 380,000 distinct corporations. Many of the successful, long-lived corporations have memberships in thousands ([EveWho 2017](#)).

2.3 Alliances

Corporations can unite to form alliances. The largest alliance has more than 28,000 player members combined from more than 500 corporations. Most successful and long-lived alliances have thousands of player members and dozens of corporations ([EveWho 2017](#)). The player who is elected to run the alliance has the authority to set the political standings toward other alliances, corporations, and/or individual players.

2.4 Geography

The game universe consists of 7,930 solar systems, located in three-dimensional space and connected by a network of (mostly) short-range transportation channels called “gates.” Although some expensive ships are capable of ignoring the connections network by directly jumping to systems within a certain range, most travel is done by traveling to neighboring systems through the gate network.

2.5 Security Levels

Of the 7,930 solar systems, 5,201 are normally navigable. These are further divided into three categories by their security level: protected high security systems (1,090), low security systems (817), and unregulated null security systems (3,294).

2.6 Sovereignty

Of the 3,294 null security systems, 2,712 are conquerable by players, meaning alliances hold sovereignty (ownership) over a system and can control access to its stations and resources. Many alliances do not hold sovereignty over any systems, for example, because they operate solely within high-security space or in abnormal wormhole space. We perform separate analyses for (1) all alliances with more than 200 members and (2) alliances holding sovereignty over at least one system (all of which have more than 200 members).

2.7 Coalitions

Alliances sometimes further coalesce into emergent social superstructures called “coalitions” that cooperate for purposes of mutual protection and coordinated attack. Their existence was neither planned, nor foreseen by the game designers, and there is no in-game support for these structures nor any inclusion in the game mechanics. The existence of these social superstructures can only be seen via player chat logs and online forums; however, in Section 5, we show that the structure of the coalitions mirrors the polarized groups predicted by structural balance theory.

3 Methods: Structural Balance Theory

As stated earlier, the core of our analytical approach is based on structural balance theory (Heider 1946; Cartwright and Harary 1956; Harary 1959). The technical aspects of balance theory have evolved over time, and we further develop the method in this paper. Starting with Abell (1968), it has become standard practice to examine only triads (aka triangles, connected triplets, and 3-cycles) in the network (although also see Facchetti, Iacono, and Altafini (2011)). Typically, the triads are considered balanced (or stable) when there are zero or two negative edges, and unbalanced (or unstable) if there are either one or three negative edges. However, it can be more insightful to look at all four types of triads separately (see Figure 1). Rather than classifying a network as balanced or unbalanced, we are interested in the proportions of each type of triad.

From a static perspective, we can determine whether the proportion of frustrated triads is smaller or larger than expected compared to random signed networks. However, since balance theory is really about network dynamics, it has become common to measure aggregate frustration over time. This approach has been applied to both simulations (Hummon and Doreian 2003; Antal, Krapivsky, and Redner 2006) and empirical networks (Leskovec, Huttenlocher, and Kleinberg 2010; Szell, Lambiotte, and Thurner 2010; DuBois, Golbeck, and Srinivasan 2011). Balance theory implies that not only should the frustrated triads constitute a small proportion of the total network, but also their prevalence in the system should exhibit a decreasing trend (unless some event injects new information into the system). Although the theory operates at the level of individual triads, because triads share edges, the change of a single edge can set in motion a series of subsequent edge changes, as the frustration cascades throughout the network (Bramson, Hoefman *et al.* 2017).

EVE offers a great deal of anecdotal evidence suggesting that players face situations and choices matching the premises of social balance. When an alliance sets a positive or negative standing to another alliance, this has a direct effect on the behaviors of everybody in both alliances (often many thousands of players), as well as indirect effects on other alliances. One such example is described in an anecdote from the history of the BRUCE alliance:

When BRUCE first arrived in Y4Y7, a constellation in southern Syndicate at the invitation of COE, they were forced to also become friends with Anarchy Empire – a rather unpleasant group of pirates that COE had befriended. BRUCE held a closer philosophy to that of the OSS at the time and became blueboxed [i.e. had positive standing] to both OSS and the COE/AE pair. This was a strange step as the OSS faction and COE/AE faction were hostile to each other, and led almost immediately to problems in the area.

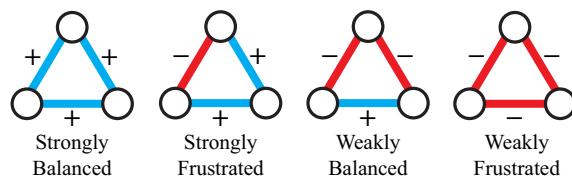


Figure 1. The four distinct triad types of balance theory and their looser categorization into weaker and stronger levels of frustration.

This very early diplomatic challenge for BRUCE was, in their own words, handled poorly. BRUCE ended up dumping the OSS bluebox less than a week after gaining it because they, understandably, did not want to fire on COE whom invited them to Syndicate. Despite BRUCE's loathing of Anarchy Empire's philosophies they valued their word to COE more. [Soon after] ... war erupted between BRUCE/COE and the OSS (EVE-history.net 2011).

Such a story strongly indicates that theories for the dynamics of international relations in general, and structural balance theory specifically, have gainful application to alliance relations in EVE.

3.1 Making the Edges Symmetric

Although early versions of balance theory were based on directed networks, and Hart (1974) proposed a measure of frustration as the proportion of balanced 2- and 3-semicycles, nearly all work on structural balance theory uses undirected networks. In EVE, alliance leaders set alliance standings toward other alliances, so standings are directed edges. Although not completely symmetric, they are nearly so with an average matched sign reciprocity of 84% (see the Supplementary Material for additional details of directed standings and a plot of reciprocity over time).

Near symmetry has been observed in other signed networks as well; for example, Facchetti *et al.* (2011) report that the directed edges in Epinions, Slashdot, and WikiElections signed networks are nearly symmetric. The reason for symmetry is clear in the game's context: if X is an ally of Y , while Y is an enemy of X , then players in Y may get penalized if they return fire when players in X attack them. Such reasoning/motivation would be present in any aggression-modulating political standings. Therefore, in the current analysis, use symmetrified standing weights in the counts and types of triads. To determine the valence for our undirected network, we set the links to be negative if either direction is negative/neutral, and positive if both directions are positive, or one is positive and the other is unset.

3.2 Strong and Weak Frustration

As mentioned above, we examine the proportions of all four types of symmetric triads, rather than classifying them as balanced or unbalanced. We consider the triple negative triads as being weakly frustrated (essentially neither frustrated nor balanced) compared to strongly frustrated single negative triads (Davis 1967; Szell *et al.* 2010). We can also consider triple positive triads as being more strongly balanced than the weaker "mutual enemy" single positive triads. The reasoning here is that if a strongly balanced triad changes via a single edge flip, then it would necessarily turn into a strongly frustrated one. However, a weakly balanced triad may become strongly or weakly frustrated, the latter of which exerts less pressure on the stability of the system of standings. Figure 1 shows a breakdown of all the four types of triads according to this looser version of balance theory.

3.3 Multipolar Political System

Compared to the tenets of classic balance theory, one specific deviation we expect to see is that triple negative triads will be over-represented. Or, more precisely, that the original structural balance theory overstates the amount of frustration, because triple negative triads do not actually represent an unstable situation. Essentially, classic balance theory implicitly assumes that there are only two real teams, and the bulk of early work focused on this "bipolar world" (Waltz 1964, 1979). In this bipolar case, one pair of nodes in the triple negative triad would find it preferable to team up against a mutual enemy, thus converting into a weakly balanced triad, the so-called common enemy effect. We see temporary cooperative behavior like this in the World Wars (Antal *et al.* 2006), and this dynamic certainly does occur in EVE. But EVE also exhibits long-standing, three-way mutual antagonism among independent political entities (Davis 1967; Axelrod and

Bennett 1993). Separating strongly and weakly frustrated triads allows one to capture the structure and dynamics of multipolar political systems in a more refined way.

3.4 Triadic Network Transformation

Our second analytical method aims to better understand triad *dynamics* and contingent behaviors. Although solutions for studying dyadic network dynamics already exist (e.g., Snijders 2017), we needed to develop a way to capture the system as a *network of triads* changing through time. To build a triadic network, we generate a *triadic node* for each triplet of alliances in the network. These triadic nodes store the relevant properties for the state of the triad that we are interested in (e.g., how many positive and negative edges or whether it is frustrated), allowing us to track triad changes over time, as well as apply weights to the triads based on their properties. Additional details of the triadic network can be found in the Supplementary Material and in Bramson, Hoefman *et al.* (2017).

4 Results: Structural Balance in EVE

Figure 2 shows the stacked proportions of triad types for sovereign alliances in EVE from February 4, 2015 to April 17, 2016. The proportions for large alliances (those with more than 200 members) are very similar (see the Supplementary Material). Although the proportions clearly fluctuate over time, they are remarkably stable in light of the volatile game environment (Belaza *et al.* 2017; Belaza *et al.* 2019). Considering the mean over time, the strongly balanced triads make up 20.26% of the total, while the strongly frustrated triads make up only 5.69%. According to the bipolar interpretation of classic balance theory, that first number is too small and the second too large. We also see a persistent level of 27.76% weakly frustrated and 46.29% weakly balanced triads. Nearly two-thirds of the triads are balanced, but that also means one-third of the triads are consistently unbalanced.

Although coalitions are not officially recognized within the game, the unofficial data (Chuggi and Sky 2017) report that strongly balanced triads are mostly (but not exclusively) internal links within coalitions—mutually friendly relations are what makes it a coalition. Coalition members nearly all share the same enemies as well as friends, so every alliance in coalition A is aggressive to both alliances X and Y in coalition B, thus creating a large number of weakly balanced triads. In addition, nearly every alliance in coalition A is an enemy with all alliances in both coalitions B and C, and this generates the large number of triple negative (weakly frustrated) triads. The existence of

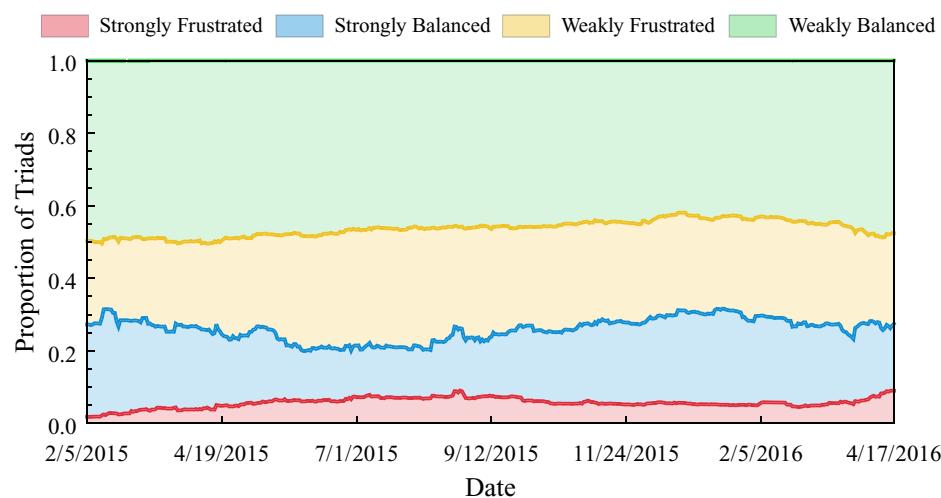


Figure 2. The stacked daily unweighted proportions of each triad type among alliances that hold sovereignty during our time frame.

coalitions thus helps make sense of why these two kinds of weak triads typically make up around three-quarters of the triads in the system.

The distributions of triad types are similar for large and sovereign alliances, indicating that the same general forces are operating in both contexts. Furthermore, the results in both cases support a modified structural balance analysis: grouping the weakly and strongly frustrated triads together drastically overestimates the number of triad changes we should expect to see. That is, strongly frustrated triads always make up less than 10% of the total, thus indicating that these arrangements are indeed avoided. The large percent of weakly frustrated (triple negative) triads support our conjecture that they do not actually inject frustration into a multipolar political system like this one.

In order to establish a baseline for comparison, we randomly assign the valence properties of the network edges while maintaining the same network structure and numbers of positive and negative links. This is repeated 100 times for each day's network, and we capture the time series of proportions of each triad type. Plots of stacked mean proportions, and a line plot showing the mean and three standard deviations, are shown in the Supplementary Material. Due to the large numbers of negative links in the system, randomly assigning their location dramatically increases the number of strongly frustrated triads (from 5.69% (sov) to 26.92%) while dramatically decreases the number of strongly balanced triads (from 20.26% (sov) to 5.6%). Although less dramatic, there is a substantial increase in the triple negative weakly frustrated triads as well (27.76% (sov) to 43.66%). Clearly, the political dynamics of EVE add meaningful structure to the distribution of edge valences in a way that reflects the social tensions of strong frustration and presence of friendly coalitions.

The numbers of both large and sovereign alliances grew dramatically during our analysis period, and the specific alliances making up each set experienced significant churn (see the Supplementary Material for details). Despite these changes in alliance numbers and participants, the overall fluctuation in proportions is small. Collectively, these results support our hypothesis that the alliance politics in EVE are consistent with multipolar balance theory. On the other hand, we expected to find (1) even less strong frustration, (2) more dynamics in frustration over time rather than a steady level, and (3) a larger difference between sovereign and large alliances' strong frustration (which are roughly equal here). In order to explain these discrepancies, we extend the analysis in two directions. First, in Section 4.1, we start from the intuition that not all alliances matter equally for systemic frustration and weigh the triads based on a variety of importance indicators. Then, in Section 4.2, we dive into the contingent behavior of the alliance triads to evaluate whether resolving frustration is actually driving edge changes in the triads.

4.1 Extension 1: Weighted Structural Balance

In the previous analysis, we discovered more strong frustration than expected, and nearly similar levels for large and sovereign alliances. Here, we consider that not all alliances are equally important in assessing systemic frustration levels. We use properties of the components of each triad to weight that triad's contribution to the counts of each type of triad. For example, alliances that are far away from each other may form a strongly frustrated triad without it actually interfering with their activities, and so it should be weighted less in computing total system frustration.

Consider a triad made up of three alliances A , B , and C . First, we standardize each property for each node or edge i to be between 0 and 1 using the minimum and maximum values over the entire time series. (Additional measurement details can be found in the Supplementary Material.)

1. $M = \sqrt[3]{M_A \cdot M_B \cdot M_C}$ for the number of member players (M_i),
2. $S = \sqrt[3]{S_A \cdot S_B \cdot S_C}$ for the number of systems over which an alliance has sovereignty (S_i),
3. $D = 1 - \sqrt[3]{D_{AB} \cdot D_{BC} \cdot D_{CA}}$ for Euclidean distances between alliances' centroids (D_{ij}),

4. $\mathcal{W} = \frac{1}{60} \sum_{i=1, j=1, i \neq j}^{3,3} |W_{ij}|$ for directed edge standing weights (W_{ij}), and
5. C is the arithmetic mean of the four above measures.

Our reasoning for (1) and (2) is that the larger the alliance, the greater the impact its involvement in a frustrated situation will have, and hence the more pressure it will experience to alleviate the frustration (McDonald and Rosecrance 1985). Distance clearly plays a role in the importance of political ties (Neumayer and Plümper 2010; Sommerer and Tallberg 2019), and we propose that being far apart spatially, (3) implies little *de facto* inconvenience, and hence less pressure to alleviate the frustration. If the roughly 6% of triads that are strongly frustrated involve small alliances, or alliances that are far away from each other, then the actual frustration in the system would be lower than the unweighted analysis above suggests. Although our construction of the triads uses symmetrified edges to determine the triad type, we use the sum of the absolute value of the weights of the six included directed edges to determine the standing weight (4). The directed edges have weights between -10 and 10 , so this weight amplifies triads among close friends (coalition members), sworn enemies, and mixtures of the two. We focus on the geometric mean, because it assigns a small weight to a triad if any of the three values is small, making it a natural match for this application. The larger alliances A and B are, the less likely they are to care about their relations with small alliance C .

- 4.1.1 *Weight-Adjusted Frustration.* For each of the five weightings above, we recalculate the number of each of the four types of triads by summing up the weights of the triads on each day. Although the sums of the weights are always smaller than the sums of the counts, by looking at the proportions of each type of triad, we can directly assess their relative incidence in a parsimonious way. Figures 3–6 show the triad proportions using each weighting scheme (solid line) compared to the unweighted sovereign proportions (dotted line) for sovereign alliances (plots for large alliances appear in the Supplementary Material). Tables 1 and 2 summarize the differences in average proportions for both large and sovereign alliances for comparison along with a statistical measure of the significance of that difference.

For sovereign alliance triads, weighting by membership (Figure 3) yields 0.95% fewer strongly frustrated and 2.18% fewer strongly balanced triads. This indicates slightly less participation in coalitions, and slightly stronger pressure to alleviate situations that would be awkward for its members. The total effect of membership weighting is more than twice the size for sovereign

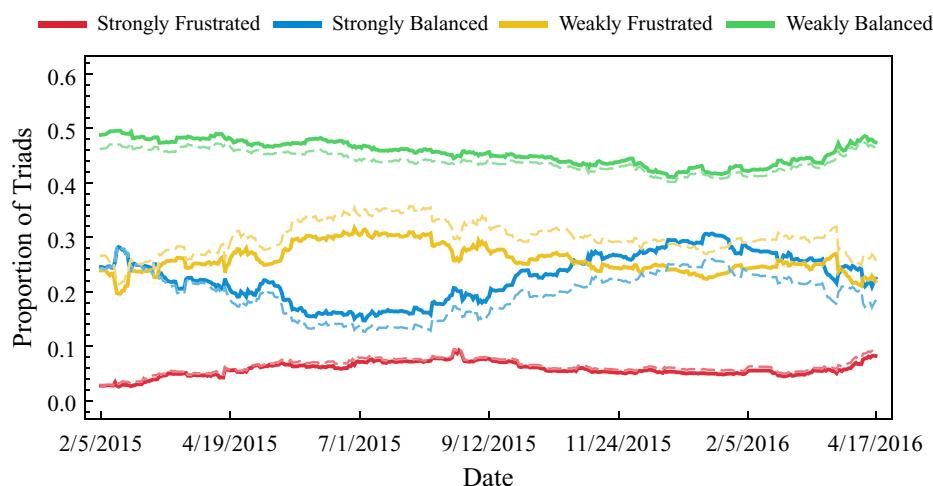


Figure 3. The daily membership-weighted proportions of each triad type among sovereign alliances (bold line) compared to the unweighted proportions (dashed line).

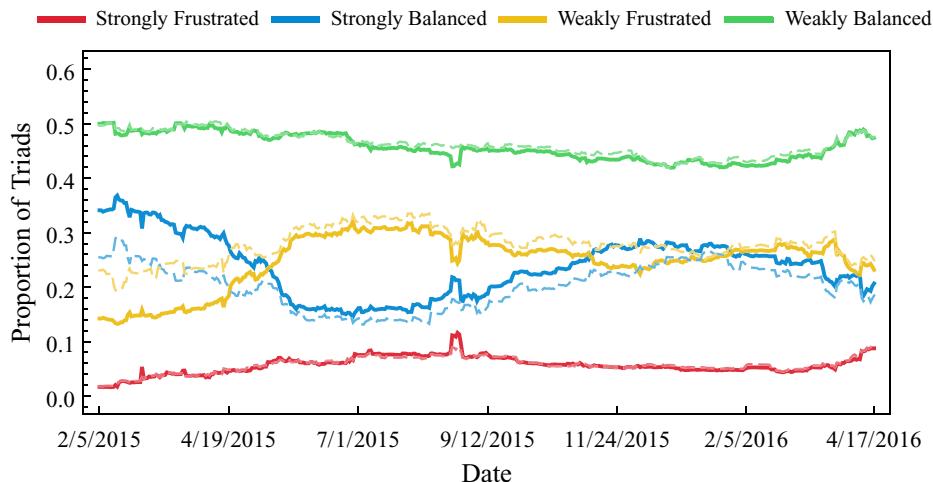


Figure 4. The daily *sovereignty-weighted* proportions of each triad type among sovereign alliances (bold line) compared to the unweighted proportions (dashed line).

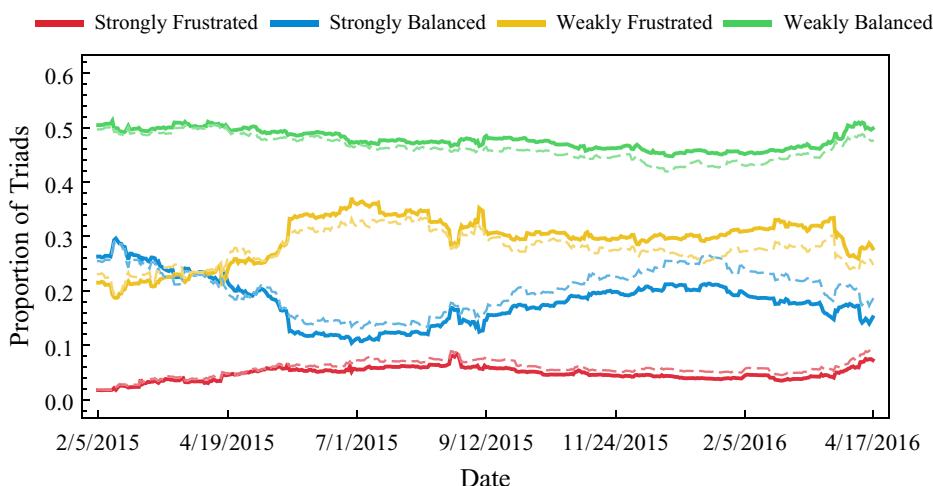


Figure 5. The daily *distance-weighted* proportions of each triad type among sovereign alliances (bold line) compared to the unweighted proportions (dashed line).

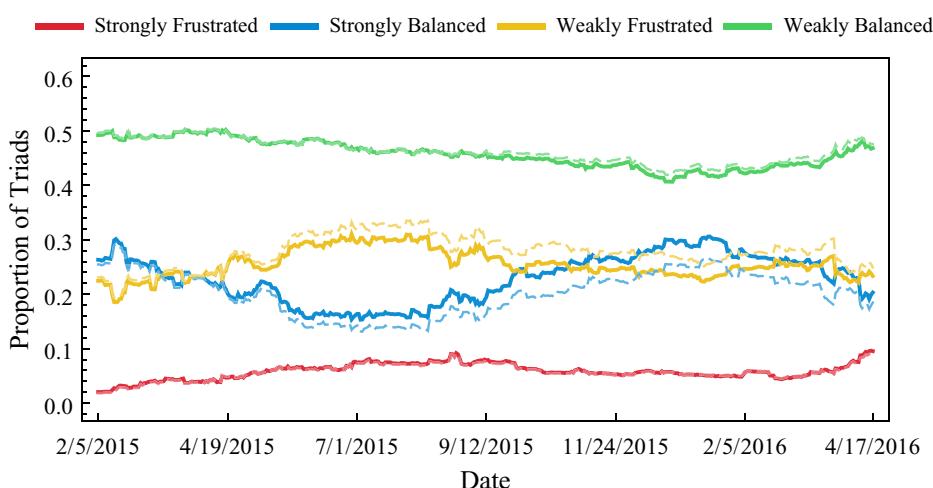


Figure 6. The daily *standing-weighted* proportions (bold line) of each triad type among sovereign alliances compared to the unweighted proportions (dashed line).

Table 1. The mean percent change (proportions $\times 100$) between the unweighted triads and each type of weighted triad for the *large alliances*.

Weighting	Strongly balanced	Strongly frustrated	Weakly balanced	Weakly frustrated	Sum of abs differences
Membership	+ 1.06	- 0.67	- 0.86	+ 0.47	+ 3.06
Sovereignty	+ 9.78	+ 0.4	- 4.49	- 5.69	+ 20.37
Distance	+ 0.18	- 0.05	+ 0.29	- 0.42	+ 0.95
Standing	+ 2.19	- 0.63	+ 1.86	- 3.43	+ 8.11
Combined	+ 1.01	- 0.27	+ 0.77	- 1.5	+ 3.55

Table 2. The mean percent change (proportions $\times 100$) between the unweighted triads and each type of weighted triad for the *sovereign alliances*.

Weighting	Strongly balanced	Strongly frustrated	Weakly balanced	Weakly frustrated	Sum of abs differences
Membership	- 2.18	- 0.95	+ 1.46	+ 1.67	+ 6.26
Sovereignty	+ 3.59	- 0.06	- 0.57	- 2.96	+ 7.19
Distance	+ 2.66	+ 0.1	- 0.5	- 2.26	+ 5.52
Standing	+ 3.28	- 0.72	+ 1.94	- 4.5	+ 10.44
Combined	+ 2.72	- 0.26	+ 0.49	- 2.95	+ 6.42

alliances than large alliances; however, a stricter adherence to balance theory is not clearly the best explanation.

Weighting by sovereignty (Figure 4) has nearly zero (-0.06%) effect on strong frustration for the sovereign alliances, and the effect on large alliances is only marginally positive ($+0.4\%$). Weighting the sovereign alliance triads by the number of systems they hold sovereignty over has an effect similar in total size to membership weighting, but in the reverse direction. For example, instead of a -2.18% change for strongly balanced triads, with sovereignty weighting, we see a $+3.59\%$ increase. This means that alliances with large populations are less likely, but alliances with large territories are more likely, to be in the same coalition.

The total effects of sovereignty for large alliances ($+20.37\%$) are nearly three times the effect on the sovereign alliances ($+7.19\%$). So, among large alliances (and recall that the sovereign alliances are also large), holding sovereignty has the largest impact on triad types (especially strongly balanced at $+9.78\%$), although conditional on being sovereign the effects are smaller ($+3.59\%$). The effect of sovereignty weighting on the strongly frustrated triads is too small to tout it as a validation of balance theory; however, the sensitivity of large alliances to sovereignty lends strong support to the idea that impacting alliance members with alliance-induced frustration does affect political decision-making within EVE.

For the distance-weighted triads (Figure 5), the main effect is more strongly balanced ($+2.66\%$) and fewer weakly frustrated (-2.26%) triads. Although the total effect of distance is small for sovereign alliances ($+5.52\%$), it is essentially negligible for large alliances ($+0.93\%$). We initially expected distance to have a larger impact, but two factors appear to mitigate the effect. First, the territory of some alliances is very spread out, and as a result, our use of the mean location places that alliance in a location that is unrepresentative for the actual situation, sometimes even placing them in a spot where they do not own any territory. Second, although there do exist clusters of nearby systems owned by a coalition of friendly alliances, when such a cluster is under a focused attack, the attacking alliance(s) often set up a base adjacent to the enemy's territory, or even

relocate there altogether. Combined, we can see that the mean location of an alliance's systems may be a poor indicator of who its neighbors are, and alliances often have large numbers of friendly *and* enemy neighbors.

As an alternative to the centroid-based distance measure, we also investigate a distance weight using the distance between the two closest systems owned by each alliance. There is evidence that closer physical relationships are positively correlated with increased violence (Maoz *et al.* 2007; Gibler and Braithwaite 2013) due to territorial disputes or long-standing cultural differences. What we find is that the closest-system distance generates triad proportions very similar to centroid-based distance but slightly closer on average to the unweighted triads. The combined weight incorporating closest-system distance also reveals a slightly smaller change from unweighted triads (details in the Supplementary Material). For either version, distance has the weakest overall effect. The high mobility of alliance locations in EVE implies that, unlike geographically bounded states, close proximity (by centroid or by nearest system) does not generate substantially increased conflict or partnership.

The standing weights (Figure 6) provide the strongest support for balance theory in two ways: (1) the total effect is the largest for sovereign alliances and second largest for large alliances, and (2) the directions of the effects all correspond to those implied by the theory. That is, it is only for the standing weightings (and the combined weighting in virtue of them) that we see more of both types of balanced triads and less of both types of frustrated triads. The largest effects occur on the triple positive and triple negative triads, which lends support to the existence of a multipolar political structure among coalitions.

The plots of the time series reveal that the effects of weights depend on dynamic features of the world. However, to ease further comparison, Table 1 shows the mean differences in proportions of all triad types between the unweighted and each weighted analyses for the large alliances, while Table 2 does so for the sovereign alliances.

We do not show the time series for the combined weighting here (see the Supplementary Material), but the results in the tables are enough to understand the most interesting points. The changes for the combined weights are *smaller* than most of the single weightings. For the large alliances, this is particularly striking: despite the sovereign weighting having a 20% impact, the combined weighting (of which sovereignty makes up 1/4) only has a 3.55% impact. What this means is that, in those triads for which sovereignty had a big impact, the other weights were negligible/countervailing, and similarly for other weightings. Rather than finding that closely located, large alliances with many members and/or territories are strongly bonded (or some other combination effect), we find for both large and sovereign alliances that the four weightings are strongly complimentary. That said, the inclusion of the standing weights does pull the combined weights in the right direction to conform to the predictions of balance theory.

The largest surprise for the weighted triads proportion analysis is the lack of a clear effect on the strongly frustrated triads. Not only are the changes among the smallest, only the strongly frustrated triads reveal insignificant changes from weighting (for distance). If any of the weightings had brought strong frustration to near zero, this would have offered decisive support for our modified structural balance in this context. Although finding persistent frustration across all the weightings is consistent with empirical data on political relations (McDonald and Rosecrance 1985), to more deeply explore the source of this structural frustration, we move on to the next extension.

4.2 Extension 2: Conditional Behavior of Alliances

The previous analysis revealed a persistent level of strong frustration in the network, belying the tenet of balance theory that frustrated triads indicate relations that need to be straightened out, and that we should expect to see decreasing levels over time. One explanation in our domain is

that there exists a constant stream of exogenous events that push triads into frustrated political relations. That is, the level of frustration may be rather consistently around 6%, but the particular alliances involved in frustrated triads may be changing due to exogenous influences. This is especially plausible considering the churn we saw in the set of alliances holding sovereignty at a given time. Here, we test the ideas that (1) alliances in frustrated triads tend to resolve those triads into balanced configurations, and (2) that the persistence of strong frustration occurs primarily through an influx of new relations.

As the anecdote about the BRUCE alliance in Section 3 illustrated, frustrated triads can lead to problems in collective action. Many factors typically go into the decision-making process of which allies to keep and which to drop, but structural balance theory implies that the change which minimizes local frustration is the most likely, even if this is not explicitly part of the decision process. Keep in mind that alliances can only change their own standings to other alliances, a change which may or may not be (but usually is) reciprocated. The triad network is extremely dense, and a change in one edge to balance a frustrated triad can (and usually does) propagate across the triad network, inducing more triads to become frustrated (Bramson, Hoefman *et al.* 2017).

In the previous two analyses, we looked at changes in the proportions of triad types. Here, we look at individual triads, and the propensity of each kind of triad to change to each other kind. Balance theory implies that balanced triads will tend to stay balanced, and strongly frustrated triads will be either avoided in the first place, or be resolved. Because the proportion of strongly frustrated triads remains fairly steady throughout the analysis period, the hypothesis we test here is that the existing, frustrated triads are indeed balancing out, but there are other, exogenous sources of frustration that are being injected into the system, causing the whole to maintain a systemic, aggregate level of frustration.

4.2.1 Triad Dynamics. Here, we focus on sovereign alliances, to determine whether the behavioral predictions of balance theory materialize at the microlevel, and to what degree the pressure to resolve frustration is affected by alliance sizes and distances. Our triadic network framework allows us to examine specific changes in the state of each triad in the system. Unlike previous analyses that examine changes in the system through dyadic changes in the network, using our triadic network construction, we directly examine the triadic dynamics of the system. First, we examine differences in the creation, persistence, and dissolution rates for each type of triad in Table 3. Dissolution of a triad means that at least one of the three symmetrified edges is removed (requiring both directed edges no longer exist), while triad creation implies a link formation that suffices to create a new triad.

The strongly frustrated triad creation rate (8.61%) is slightly higher than its persistence rate (5.88%), and slightly lower than its dissolution rate (9.17%). This confirms that there is a proportional influx of frustration when links/triads are created, and a proportional outflux of strong frustration through removed links/triads. The proportions of both types of balanced triads are highly similar across entrance, persistence, and removal, but the weakly frustrated triads have a 28% persistence level, despite a 24% creation rate. This implies that triads are becoming weakly frustrated from the three other types. Next, we examine conditional changes in triad types, that is, what kind of triad turns into what other kind of triad.

A table of the results is available in the Supplementary Material, but the most salient feature is that more than 98% of the triads stay in the same state from day to day. The stability in the proportions of triad types we saw in the time series above can be attributed to the proportions of triads when created combined with the high persistence of triad types. That said, strongly frustrated triads are the least persistent and most likely to become nonexistent. Here, we again see

Table 3. Summary results showing the percentages of unweighted sovereign alliance triad types when they are created, that persist in the system over time, and when they are removed. This only includes adding and removing triads through link creation and destruction, that is, excluding nodes entering/leaving the system.

	Strongly balanced	Strongly frustrated	Weakly balanced	Weakly frustrated
Triad creation	21.96	8.61	45.53	23.91
Triad persistence	20.24	5.88	45.93	27.95
Triad dissolution	20.43	9.17	44.23	26.17

Table 4. Percentage summary of unweighted sovereign alliance triad type changes, including changes through the deletion of edges. Because we use daily data, it is possible for triads to change multiple edge valences in one iteration.

From	To strongly balanced	To strongly frustrated	To weakly balanced	To weakly frustrated	To nonexistent
	balanced	frustrated	balanced	frustrated	nonexistent
Strongly balanced	–	39.78	8.85	0.04	51.33
Strongly frustrated	29.13	–	35.44	0.74	34.69
Weakly balanced	2.17	11.98	–	36.32	49.53
Weakly frustrated	0.03	0.35	56.62	–	43.00

the relatively high dissolution of strongly balanced triads (more below) and the greatest stability occurring for weakly frustrated triads.

We now examine changes in triad types conditional on there being a change. We show these conditional triad change proportions, including triads that dissolve, in Table 4. (The analogous tables for large alliances and weighted triads can be found in the Supplementary Material.)

The first observation is the size of the proportions of triads that are dissolved through link removal (right-most column of Table 4). It is worth noting that strongly balanced triads are more likely to be dissolved altogether, rather than to change type. The second observation is that when a triad does change its state, it is proportionally much more likely to change by a single valence flip. Because we use daily data, more than one edge change can occur in a single time step. And although this does happen, we find that 94.5% of changes (excluding deletion) are a single edge flip away.

This second observation is especially interesting in the context of the first one. Because a single edge change will always bring a strongly balanced triad to the strongly frustrated state, the fact that strongly balanced triads are disproportionately dissolved rather than changed can be explained by the avoidance of the frustration it would otherwise create. Comparing this dynamic to the weakly balanced triads further supports this conclusion: weakly balanced triads are more than three times more likely to transition to a weakly frustrated state (36.32%) than to a strongly frustrated state (11.98%). Dissolving the triad is more than four times more likely (49.53%) than changing to a strongly frustrated state. Although the persistence of strongly frustrated triads seems to conflict with the predictions of balance theory, we find that when there are changes, strongly frustrated triads tend to be preferentially avoided. This result provides reasonable support for balance theory as a partial explanation of triad dynamics.

- 4.2.2 *Weighted Triad Dynamics.* Here, we augment the unweighted analysis with a brief treatment of the effects of weights on triad changes. The tables for triad change rates for all weightings, for both large and sovereign alliances, appear in the Supplementary Material. Here, we present tables

Table 5. Summary results showing percentage changes resulting from applying the combined weighting to sovereign alliance creation, persistence, and removal rates (i.e., combined weighting rates minus unweighted rates). This only includes adding and removing triads through link creation and destruction, that is, excluding nodes entering/leaving.

	Strongly balanced	Strongly frustrated	Weakly balanced	Weakly frustrated
Triad creation	+ 2.26	- 0.3	+ 0.02	- 1.99
Triad persistence	+ 2.99	- 0.3	+ 0.42	- 3.11
Triad dissolution	+ 8.5	+ 4.82	- 2.69	- 10.64

Table 6. Summary results of the differences in the proportions of triad type changes between combined-weighted and unweighted sovereign alliances conditional on there being a change, including though the deletion of edges.

From	To strongly balanced	To strongly frustrated	To weakly balanced	To weakly frustrated	To nonexistent
Strongly balanced	-	+ 2.7	+ 0.45	- 0.01	- 3.14
Strongly frustrated	+ 5.03	-	- 1.03	- 0.11	- 3.89
Weakly balanced	+ 0.69	+ 2.02	-	- 0.42	- 2.3
Weakly frustrated	+ 0.01	+ 0.05	+ 4.15	-	- 4.21

for the *differences* in change rates between the unweighted analysis and the combined weights, for creation, persistence, and dissolution (Table 5), and for conditional changes (Table 6), both using the data from sovereign alliances. Although we present these only for the combined weights, the patterns in changes are similar across the individual weights (available in the Supplementary Material) with only specific numerical difference by weight.

Table 5 shows that the total effect of the weightings is small in most places, but there are some interesting observations. First, creation and persistence of strongly frustrated triads have very small negative effects (-0.3%), but the combined weighting produces a 4.82% increase in their dissolution. Second, the persistence effect on both strongly and weakly balanced triads is positive, while the effect on weakly and strongly frustrated triads is negative. Thus, the marginal effect of the weights is an increased adherence to balance theory with respect to the kinds of triads that persist. Third, the effect is strongest for triad dissolution, where strongly balanced triads are 8.5% more likely, and weakly frustrated triads are 10.64% less likely, to be dissolved when the combined weights are applied. The strong changes in triad dissolution, however, are not clearly indicative of a stronger or weaker adherence to balance theory as a result of the weightings.

We next examine our hypothesis that the inclusion of weights increases the predictive power of structural balance theory on triad changes. Table 6 reports the changes from Table 4 by including the combined weighting. The top row reveals that strongly balanced triads are 2.7% *more* likely to become strongly frustrated, and 3.14% *less* likely to become nonexistent. Strongly frustrated triads are 5.03% *more* likely to become strongly balanced, and 3.89% *less* likely to become nonexistent.

The effect on transitions from strongly balanced and strongly frustrated triads is too small to make any strong claims, and the effect on other transitions is even smaller. However, we do note that these effects are not in the directions we would expect from balance theory. So, although the unweighted conditional triad changes did lend support to balance theory, our addition of weights to amplify these changes had little effect, and the overall effect does not strengthen our support for the predictions of structural balance.

5 Coalitions and Polarization

In addition to a tendency toward alleviating systemic frustration, balance theory also predicts a tendency for political entities to cluster into mutually friendly coalitions, separated by negative links. This kind of signed network clustering is often referred to as “polarization,” and received a lot of attention in early studies of balance theory. The term “polarization” takes on many senses (Bramson, Grim *et al.* 2017), and there are various measures for each of those senses (Bramson *et al.* 2016). Although polarization is most often measured on distributions of values on a scale (e.g., of political opinions or beliefs; Esteban and Ray 1994), one can also measure polarization on spatial or networked values (Maoz 2006; Esteban and Ray 2008). Importantly, one must distinguish between polarization of attributes across a network and polarization in the network structure itself.

Most social and political analyses are concerned with the effect of network structure on the polarization of social, political, and economic attributes, but structural balance has a long tradition of analyzing structural polarization directly (Cartwright and Harary 1956; Hart 1974). The earliest literature focused on determining whether a network could be split into two “polarized” groups. Latter work also examined how well networks were split into (possibly more than two) polarized groups (Kulakowski 2007; Doreian and Mrvar 2009). We are interested in the latter; specifically, we are interested in using the EVE standings data to determine how well balance theory’s predictions about group formation describe the observed system of political relations in this virtual world.

As discussed in Section 2, the alliances in EVE form unofficial coalitions: political super-entities that are not part of the game’s mechanics, yet are widely recognized by players. Coalition dynamics include tacit internal no-conflict pacts, and joint strategizing among member alliances. We use player-reported, daily coalition data from Chuggi and Sky (2017) in the form of unofficial, yet widely used, coalition maps to assess how well the network structure of actual alliance standings corresponds to these player-reported coalitions. Although data that are collected and reported by players have several limitations (discussed in detail in the Supplementary Material), these player-reported maps are the best record of alliances’ membership in coalitions, and we can use them to perform a preliminary test of the hypothesis that structural balance guides the formation of polarized coalitions.

5.1 Detecting Coalitions

We use the same standings data here as were used in the triad analysis, but now we are using them to assess polarization patterns in the network. Figure 7 displays the adjacency matrices of standings using the same intuitive color scheme the players see in the game. On the left, it is among the 38 sovereign alliances known to be involved in coalitions on February 4, 2015, and on the right, the 80 sovereign alliances in coalitions on April 17, 2016. By sorting the rows by (Canberra) similarity, we can easily discern a block structure of friendly links that is typical of network community structure and reminiscent of polarization in classical balance theory networks. To formalize the detection of coalitions, we apply two categories of methods: (1) network community structure methods and (2) vector distance-based measures. For all methods, we apply them to the directed standings, the reverse directed standings, and the unweighted symmetric matrix.

- 5.1.1 *Network Community Detection.* Community detection algorithms are sophisticated techniques from network theory, specifically designed to find clusters of nodes that are densely connected internally, and sparsely connected externally. In the past decade, several research groups have turned to the challenge of detecting communities in signed social networks (Doreian and Mrvar 2009; Traag and Bruggeman 2009; Anchuri and Magdon-Ismail 2012; Amelio and Pizzuti 2013; Chen *et al.* 2014; Esmailian and Jalili 2015), but these methods are not yet well-tested or easily

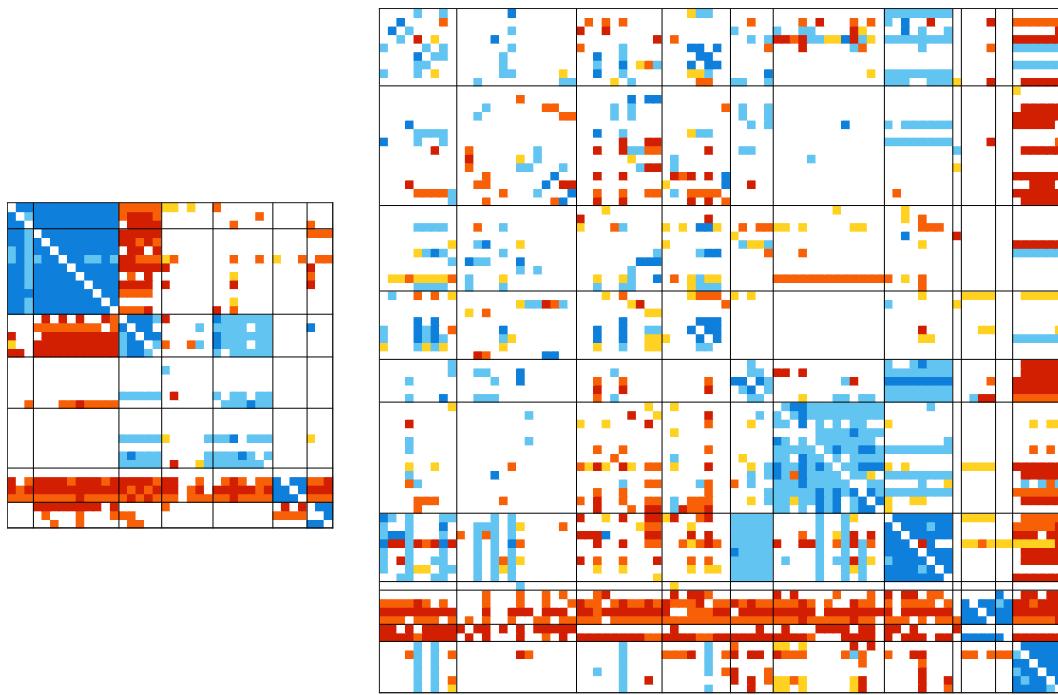


Figure 7. Plots of the adjacency matrices of alliance standings on February 4, 2015 (left) and April 17, 2016 (right) between the sovereign alliances included in Chuggi and Sky’s (2017) alliance and coalition maps. Blocks of dark and light blue (standings of 10 and 5, respectively) represent clusters of alliances that are densely interconnected with positive links. Negative links are represented by red (-10) and orange (-5), and neutral relationships by yellow (0) cells. White space indicates unset standings. The black mesh lines indicate coalitions discovered through clustering by Hamming distance on out-edges, one of the measures we compare to the player-reported coalition data.

accessible. One simple and obvious method to detect network communities is based solely on the positive links (Yang, Cheung, and Liu 2007). Although this technique ignores the repulsive force that negative links should have between communities, it suffices for our purposes, because, in practice, alliances with no set positive standings act hostile toward each other by default. We include six different algorithms for identifying community structure in the networks (Hierarchical, Centrality, Vertex Moving, Modularity, Spectral, and Clique Percolation; Wolfram Research 2019).

- 5.1.2 *Measures of Distance.* Vector-based methods take each row (or column) as a vector of values, and apply a standard distance metric to perform pairwise tests of similarity. We applied the following list of distance measures here: Canberra, Euclidean, Normalized Squared Euclidean, Squared Euclidean, Cosine, Manhattan, Bray Curtis, Damerau–Levenshtein, Hamming, Correlation, and Chessboard (Wolfram Research 2019).
- 5.1.3 *Measure of Accuracy.* After using a vector distance or network community method to partition the alliances into proposed coalitions, we need to determine how accurately they matched our best reference for the “real” coalitions. Naturally, a perfect match (getting “full points”) occurs when each member of a discovered coalition is a member of that coalition in the player-reported data as well. For each discovered coalition, we find the distribution of real coalitions for those members. We then identify the real coalition with plurality within the discovered coalition, and use it as the matching real coalition. For each member of a discovered coalition that is not in the same real coalition as the plurality, we reduce the accuracy by one point. The points are then normalized by the number of alliances to create a percent accuracy score. More details are available in the Supplementary Material.

Table 7. The top five ranked accuracy methods used to identify the coalitions reported by Chuggi and Sky (2017) using the rows of the directed weighted standings matrix (i.e., the out-edges).

Distance measure	First day out-edges	Last day out-edges	Mean
Hierarchical community*	0.895	0.675	0.785
Vertex moving community*	0.816	0.688	0.752
Centrality community*	0.816	0.675	0.745
Spectral community*	0.816	0.662	0.739
Modularity community*	0.816	0.662	0.739

Note: Network-based community detection algorithms are marked with an asterisk.

5.1.4 Results of Detecting Coalitions. The mean accuracy among the best five methods for the out-edges, in-edges, and symmetric edges are 0.752, 0.765, and 0.766, respectively. These very similar scores belie the fact that the five best-performing methods for each category are actually quite different. When analyzing out-edges (the values actually set by alliance leaders), the network community detection methods are the most accurate (see Table 7; full tables available in the Supplementary Material). The best method, hierarchical community detection, only mismatched 4 of the 38 alliances.

When clustering on the in-edges, we find that the vector-based measures outperform the network measures. Furthermore, the best measure (Hamming Distance) achieves scores of 0.895 and 0.675 for the first and last days, respectively; this is the same as Hierarchical Community Detection on the out-edges. We find that the vector-based measures also perform better on the symmetric data, yet the ranking of the measures is again very different. For example, Hamming Distance drops from #1 to #13, and Damerau–Levenshtein Distance jumps from #15 to #1. Although the performance of any given measure may vary widely among the three data representations, the overall performance of the suite of measures is consistent (making it difficult to choose a single best method).

One form of this consistency is that nearly every measure on all three representations performs better on the first day than on the last day. Recall Figure 7, showing the adjacency matrices; not only are there many fewer alliances on the first day, but they are also more clearly organized in a block structure. Some alliances form strong coalitions, and these are almost universally discovered. Other coalitions are weakly bound and more fluid. Upon analysis, some of the actual coalition members seem out of place and perhaps about to leave; thus, some inaccuracy can be attributed to performing a static analysis on a changing, dynamic situation.

Using these techniques, we find a match between predicted and player-reported coalitions of roughly 70%–80%. This result indicates that the unofficial, organically player-created coalitions in the game correspond to the emergent metastructures predicted by balance theory reasonably well. It is rare to see such a clear demonstration of a theory of political organization.

5.2 Polarization Among Coalitions

We now move from the accuracy of the coalitions found to the analysis of the link valences among coalitions. Figures 8 and 9 show network diagrams, both based on the standings on the first day of our data. The nodes represent the alliances, and the edge colors again reflect the directed standings between the alliances. The clusters, shown in Figure 8, reflect the real coalitions, while in Figure 9, the clusters represent the communities discovered by hierarchical partitioning based on positive edges.

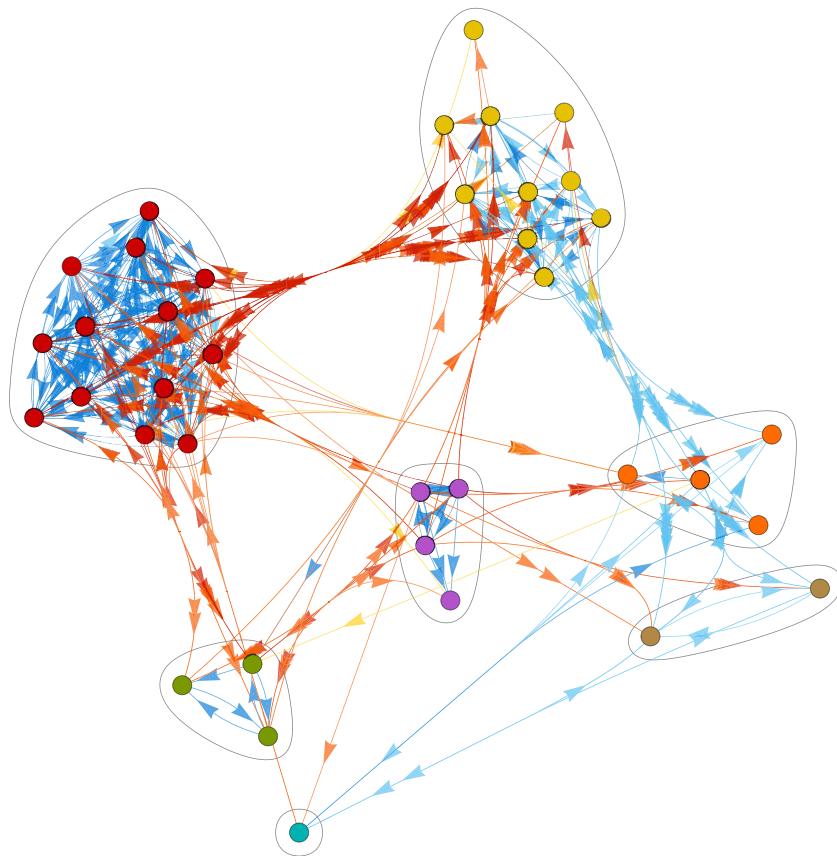


Figure 8. A network diagram of the alliance standings within each coalition on February 4, 2015 according to Chuggi and Sky (2017). Alliance nodes are grouped by their collation, and the edges are colored by their directed standings as in Figure 7. Additional visualizations are available in the Supplementary Material.

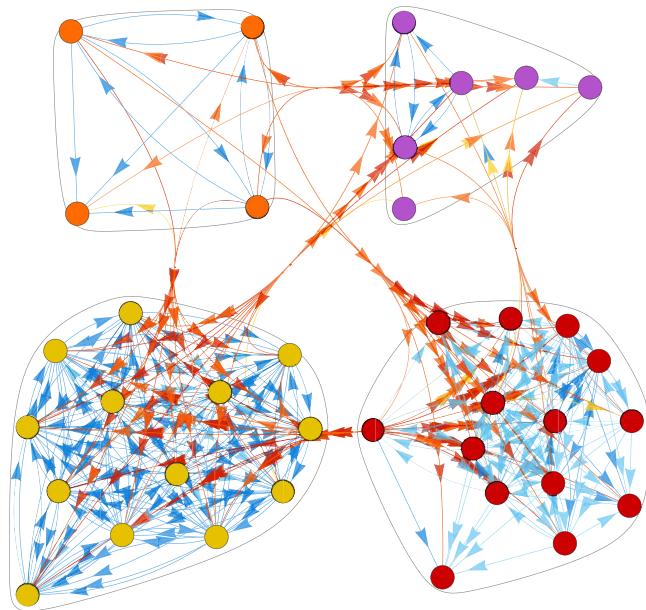


Figure 9. A network diagram of the coalitions discovered by hierarchical clustering on the positive edges, based on the standings data on February 4, 2015.

The first thing to notice are several small coalitions at the top of Figure 8, joined solely by positive edges. These would naturally be considered as one single cluster by any network clustering algorithm. Two of these small coalitions are also mostly positively connected to the cluster on

the bottom right. But with very few exceptions, the remaining intercoalition edges are negative, and the remaining intracoalition edges are positive. So, although some of the coalitions could be fused, by and large, the coalitions created by players conform to the predictions of balance theory.

That said, we can also consider balance theory to claim that social ties will be partitionable (but not necessarily *de facto* partitioned) into groups of internally positive and externally negative relations. For this, we look at the results of hierarchical clustering in Figure 9, where we discover four clusters with only four negative edges within any of the clusters, and two positive edges between clusters (out of a total of 543 edges). So both the actual and the discovered coalitions are highly polarized, as predicted by balance theory.

Although we cannot make any bold claims based on this preliminary analysis, the results here are sufficiently strong to merit deeper study. Because the coalitions are not part of the game mechanics, there are no official rules guiding their formation or behavior. There is no official need for political organizations above the alliance level to exist. Demonstrating that not only do coalitions emerge, but that the emergent coalitions are highly polarized, is a big win for balance theory.

6 Conclusions

Our analysis of political relationships between alliances in the virtual world of *EVE* revealed mixed support for balance theory. The persistence of strongly negative triads in Section 4.1 goes against balance theory, while the fact that these strongly negative triads consistently make up the smallest proportion of triads conforms with what we would expect based on the theory. The analysis of contingent behaviors in Section 4.2 showed that strongly frustrated triads are preferentially avoided, but that they are still tolerated more than expected. Section 5 makes the case that the mere existence of coalitions is already in support of balance theory. The clear polarization of the empirical and discovered coalitions provides the best support.

Our findings also highlight the importance of considering the system as multipolar rather than bipolar, that is, splitting the frustrated and balanced triads into weak and strong versions of each, and tracking the four distinct structures. By doing so, we find that the prevalence of mutual antagonism (triple-negative = weakly frustrated triads) is much greater than predicted by traditional balance theory. In the game, as in real politics, there are more than two factions vying for power. And although unfriendly factions may temporarily team up to defeat a stronger mutual enemy, long-term three-way animosity seems a natural occurrence in large-scale political networks.

The effects of weighting the triads by membership, sovereignty, and distance is smaller than expected, implying triads of all types are somewhat evenly distributed over alliance sizes and locations. Furthermore, the combined weight effects are smaller than the individual weights, implying that at least one of the three weights is very low when another is high. One exception is weighting by sovereignty, which has an appreciable effect for large alliances, and a much smaller effect for sovereign alliances. This supports our belief that sovereignty itself is important for political relations in *EVE* because of how it affects access to resources. We expected the weighted analyses to reveal a small subset of major players (e.g., large and well-managed dominating alliances), adhering more strongly to balance theory than the average participant. Although the effects of the weights were in the right direction, further investigation is required to explain the deviations of alliance behavior from the predictions of balance theory.

6.1 Future Work

One of our broader goals is to explore whether data such as these can provide evidence for generalized law-like features of sociopolitical systems. Other work based on *EVE* data aims to

demonstrate that economic principles apply, and in fact apply more cleanly, in the partially idealized reality of virtual worlds (Hoefman *et al.* 2019). Balance theory is one candidate for a sociopolitical theory that could hold generally. Going forward, we will extend the analysis to other sociopolitical theories, and other aspects of EVE. Potentially, and most excitingly, the cleaner and simpler data from virtual worlds such as EVE may provide the inspiration and fuel to develop novel sociopolitical theories.

In the current work, we use the actual directed edge standing weights (i.e., 0, 5, or 10) in the standing-weighted frustration, but otherwise use symmetric positive versus negative valences to define triad types. This coincides with the binary conflict conditions faced by most players (either shoot, or do not shoot), but setting a standing to 5 instead of 10, or setting it to zero (neutral) instead of leaving it unset, is intentional, and presumably meaningful signals about the relationship. Doing so requires a technique that does not yet exist for balance theory (although Belaza *et al.* (2017) and Belaza *et al.* (2019) account for neutral ties and degeneracies in triad composition); therefore, we are developing new methods for this research thrust.

Although balance theory is a popular and well-supported model, it is not the only model for analyzing tensions in social/political relations. Axelrod and Bennett (1993) offer a *Landscape Theory* that utilizes features of states to partition them into coalitions. Although not explicitly a network approach, the factors considered there (trade relations, distance, etc.) could be given a signed multigraph representation and analyzed for structural polarization.

One can also explore the change and spread of policies/properties across networks through econometric methods on dyadic data. For example, Neumayer and Plümper (2010) examine spatial effects in bilateral investment treaties, and Sommerer and Tallberg (2019) perform a similar analysis for the diffusion of participatory governance. Although these dyadic models may offer improvements over monadic models (Simmons and Elkins 2004), to understand the dynamics of signed international relations, one must at least analyze triadic relationships (Maoz *et al.* 2007). Further examining the statistical properties of the triad changes through something like SIENA (Snijders, Van de Bunt, and Steglich 2010; Snijders 2017) may yield deeper insight into the structural and temporal dependencies driving the dynamics. Although statistical models looking at dyadic, triadic, and even larger relationships face serious theoretical and methodological issues compared to network analytical methods (Cranmer and Desmarais 2016), given the richness of the EVE dataset, we are interested in exploring alternatives to balance theory for explaining the political dynamics among alliances in future work.

Overall, our results bolster the relevance of balance theory for understanding a wide range of human behaviors—including video game politics. They also demonstrate the usefulness of data from virtual worlds for evaluating social theories. People are still social beings, even when interacting in virtual worlds, as long as the incentives structures are sufficiently realistic.

Acknowledgments

This research was supported by the Research Foundation Flanders (FWO) under Grant Number G018115N and Bijzonder Onderzoeksfonds BOF2452014000402.

Data Availability Statement

The replication materials for this paper can be found at Bramson *et al.* (2020) at Harvard Dataverse at <https://doi.org/10.7910/DVN/S8M39Y>.

Supplementary Material

For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2021.1>.

Bibliography

- Abell, P. 1968. "Structural Balance in Dynamic Structures." *Sociology* 2(3):333–352.
- Amelio, A., and C. Pizzuti. 2013. "Community Mining in Signed Networks: A Multiobjective Approach." In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 95–99. New York: ACM.
- Anchuri, P., and M. Magdon-Ismail. 2012. "Communities and Balance in Signed Networks: A Spectral Approach." In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 235–242. Washington, DC: IEEE.
- Antal, T., P. L. Krapivsky, and S. Redner. 2006. "Social Balance on Networks: The Dynamics of Friendship and Enmity." *Physica D: Nonlinear Phenomena* 224(1):130–136.
- Axelrod, R., and D. S. Bennett. 1993. "Landscape Theory of Aggregation." *British Journal of Political Science* 23(2):211–233.
- BBCNews. 2014. "Eve Online Virtual War 'Costs \$300,000' in Damage," January 29. <http://www.bbc.com/news/technology-25944837>.
- Belaza, A. M., K. Hoefman, J. Ryckebusch, A. Bramson, M. van den Heuvel, and K. Schoors. 2017. "Statistical Physics of Balance Theory." *PLoS One* 12(8):12–15.
- Belaza, A. M. et al. 2019. "Social Stability and Extended Social Balance—Quantifying the Role of Inactive Links in Social Networks." *Physica A: Statistical Mechanics and Its Applications* 518:270–284.
- Bramson, A., K. Hoefman, K. Schoors, and J. Ryckebusch. 2020. "Replication Data for: Diplomatic Relations in a Virtual World." doi: [10.7910/DVN/S8M39Y](https://doi.org/10.7910/DVN/S8M39Y), Harvard Dataverse, V1.
- Bramson, A., K. Hoefman, M. van den Heuvel, B. Vandermarliere, and K. Schoors. 2017. "Measuring Propagation with Temporal Webs." In *Temporal Network Epidemiology*, 57–104. Springer: Singapore.
- Bramson, A. et al. 2016. "Disambiguation of Social Polarization Concepts and Measures." *The Journal of Mathematical Sociology* 40(2):80–111.
- Bramson, A. et al. 2017. "Understanding Polarization: Meanings, Measures, and Model Evaluation." *Philosophy of Science* 84(1):115–159.
- Cartwright, D., and F. Harary. 1956. "Structural Balance: A Generalization of Heider's Theory." *Psychological Review* 63(5):277.
- CCP. 2019. "Eve Online." <http://www.eveonline.com>.
- Chen, Y., X. Wang, B. Yuan, and B. Tang. 2014. "Overlapping Community Detection in Networks with Positive and Negative Links." *Journal of Statistical Mechanics: Theory and Experiment* 2014(3):P03021.
- Chuggi, and M. Sky. 2017. "Eve Null-Sec Coalition Influence Maps," July 11. <http://eve-files.com/>.
- CorrelatesOfWar. 2019. "The Correlates of War Project." <http://correlatesofwar.org>.
- Cranmer, S. J., and B. A. Desmarais. 2016. "A Critique of Dyadic Design." *International Studies Quarterly* 60(2):355–362.
- Davis, J. A. 1967. "Clustering and Structural Balance in Graphs." *Human Relations* 20:181–187.
- Doreian, P., and A. Mrvar. 2009. "Partitioning Signed Social Networks." *Social Networks* 31(1):1–11. doi: [10.1016/j.socnet.2008.08.001](https://doi.org/10.1016/j.socnet.2008.08.001).
- DuBois, T., J. Golbeck, and A. Srinivasan. 2011. "Predicting Trust and Distrust in Social Networks." In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom)*, 418–424. Boston, MA: IEEE.
- Economist, T. 2015. "Friends and Foes: Rifts in the Middle East." <http://espresso.economist.com/f59633e6085d53d7aea1c1bf83570f79>.
- Esmailian, P., and M. Jalili. 2015. "Community Detection in Signed Networks: The Role of Negative Ties in Different Scales." *Scientific Reports* 5:14339.
- Esteban, J., and D. Ray. 2008. "Polarization, Fractionalization and Conflict." *Journal of Peace Research* 45(2):163–182.
- Esteban, J.-M., and D. Ray. 1994. "On the Measurement of Polarization." *Econometrica: Journal of the Econometric Society*: 819–851.
- EVE-history.net. 2011. "BRUCE—EVE History, March." Accessed June 26, 2017. <http://eve-history.net/wiki/index.php/BRUCE>.
- Eve-Offline. 2018. "Eve-Online Status Monitor," January. <http://eve-offline.net/?server=tranquility>.
- EveWho. 2017. "Eve Who," June 23. <http://evewho.com/alli/>.
- Facchetti, G., G. Iacono, and C. Altafini. 2011. "Computing Global Structural Balance in Large-Scale Signed Social Networks." *PNAS* 108(52):20953–20958.
- Gibler, D. M., and A. Braithwaite. 2013. "Dangerous Neighbours, Regional Territorial Conflict and the Democratic Peace." *British Journal of Political Science* 43(4):877–887.
- Goh, A. 2018. "How to Build a Robust Game Economy," September 3. <http://hackernoon.com/how-to-build-a-robust-game-economy-lessons-from-one-of-the-worlds-longest-running-mmos-426f8fd94f6d>.
- Harary, F. 1959. "On the Measurement of Structural Balance." *Behavioral Science* 4(4):306–323.
- Hart, J. 1974. "Symmetry and Polarization in the European International System, 1870–1879: A Methodological Study." *Journal of Peace Research* 11(3):229–244.

- Heider, F. 1946. "Attitudes and Cognitive Organization." *Journal of Psychology* 21:107–122.
- Hoefman, K., A. Bramson, K. Schoors, and J. Ryckebusch. 2019. "The Impact of Functional and Social Value on the Price of Goods." *PLoS One* 13(11):e0207075.
- Hummon, N. P., and P. Doreian. 2003. "Some Dynamics of Social Balance Processes: Bringing Heider Back into Balance Theory." *Social Networks* 25(1):17–49.
- Kulakowski, K. 2007. "Some Recent Attempts to Simulate the Heider Balance Problem." *Computing in Science & Engineering* 9(4):80–85.
- Leskovec, J., D. Huttenlocher, and J. Kleinberg. 2010. "Signed Networks and in Social and Media." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1361–1370. New York: ACM.
- Maoz, Z. 2006. "Network Polarization, Network Interdependence, and International Conflict, 1816–2002." *Journal of Peace Research* 43(4):391–411.
- Maoz, Z., L. G. Terris, R. D. Kuperman, and I. Talmud. 2007. "What is the Enemy of My Enemy? Causes and Consequences of Imbalanced International Relations, 1816–2001." *The Journal of Politics* 69(1):100–115.
- McDonald, H. B., and R. Rosecrance. 1985. "Alliance and Structural Balance in the International System: A Reinterpretation." *Journal of Conflict Resolution* 29(1):57–82.
- Mildenberger, C. 2013. *Economics and Social Conflict: Evil Actions and Evil Social Institutions in Virtual Worlds*. London: Springer.
- Neumayer, E., and T. Plümper. 2010. "Spatial Effects in Dyadic Data." *International Organization* 64(1):145–166.
- Simmons, B. A., and Z. Elkins. 2004. "The Globalization of Liberalization: Policy Diffusion in the International Political Economy." *American Political Science Review* 98(1):171–189.
- Snijders, T. A. 2017. "Stochastic Actor-Oriented Models for Network Dynamics." *Annual Review of Statistics and Its Application* 4(1):343–363. doi:[10.1146/annurev-statistics-060116-054035](https://doi.org/10.1146/annurev-statistics-060116-054035).
- Snijders, T. A., G. G. Van de Bunt, and C. E. Steglich. 2010. "Introduction to Stochastic Actor-Based Models for Network Dynamics." *Social Networks* 32(1):44–60.
- Sommerer, T., and J. Tallberg. 2019. "Diffusion Across International Organizations: Connectivity and Convergence." *International Organization* 73(2):399–433.
- Squire, K. 2008. "Open-Ended Video Games: A Model for Developing Learning for the Interactive Age." In *The Ecology of Games: Connecting Youth, Games, and Learning*, 167–198. Cambridge, MA: MIT Press.
- Szell, M., R. Lambiotte, and S. Thurner. 2010. "Multirelational Organization of Large-Scale Social Networks in an Online World." *PNAS* 107(31):13636–13641.
- Traag, V. A., and J. Bruggeman. 2009. "Community Detection in Networks with Positive and Negative Links." *Physical Review E* 80(3):036115.
- Waltz, K. N. 1964. "The Stability of a Bipolar World." *Daedalus* 93(3):881–909.
- Waltz, K. N. 1979. *Theory of International Politics*. Long Grove, IL: Waveland Press.
- Wolfram Research, I. 2019. "Mathematica." Version 10.0. Champaign, IL, 2019. <http://www.wolfram.com/mathematica>.
- Yang, B., W. Cheung, and J. Liu. 2007. "Community Mining from Signed Social Networks." *IEEE Transactions on Knowledge and Data Engineering* 19(10):1333–1348.