

## **Subcoalition Cluster Analysis: A New Method for Measuring Political Conflict in Organizations**

Scott C. Ganz, Georgia Tech School of Public Policy and AEI

Daniel Schiff, Georgia Tech School of Public Policy

### **Abstract**

Behavioral theories of organizational decision making emphasize that organizations are political coalitions. Despite considerable recent qualitative research in management and organizational theory on the role of politics in decision making and managing organizational change, quantitative research in this area has stalled. The reason for the lack of progress is not theoretical, but rather methodological; researchers lack empirical tools for understanding basic processes of coalition formation, evolution, and conflict in organizations. We introduce a novel method for modeling politics in organizations that builds on the model of intra-organizational conflict in March (1962), which we call “subcoalition cluster analysis” (sCCA). The main contribution of sCCA is that it identifies subcoalitions with consistent preferences that are in conflict without placing additional restrictions on the structure of individual preferences. We apply sCCA to two cases, Wikipedia and the Baseball Writers’ Association of America and show how leadership would benefit from conceiving of their membership as competing subcoalitions instead of individuals with idiosyncratic preference disagreement. Finally, we compare the performance of sCCA to Principal Component Analysis (PCA) and  $k$ -means clustering and demonstrate that sCCA does a better job identifying latent structure in the data when the organization consists of more subcoalitions, when individual preferences are not perfectly aligned with those of their subcoalition, and when observations are missing.

## 1. Introduction

Behavioral theories of organizational decision making emphasize that organizations are political coalitions (Cyert and March 1963; Pfeffer 1981). The benefits associated with specialization among and collaboration between units with different expertise, roles, and experiences often come at the cost of disagreement about desired outcomes (Heath and Staudenmayer 2000). Behavioral models of organizations that analogize organizational decision making to individual decision making implicitly deny that many important organizational decisions emerge through a process of contestation and compromise. Yet, despite the considerable attention paid to the role of politics in decision making and organizational change in recent qualitative research on organizations (e.g., Kellogg 2009; Morrill 1991; Zbaracki and Bergen 2010), quantitative research in this area has stalled among management scholars (Gavetti, Levinthal, and Ocasio 2007; Gibbons 2003). The reason for the lack of progress is not theoretical, but rather methodological; management scholars lack empirical tools to identify and measure basic processes of coalition formation, evolution, and conflict in organizations.

We introduce a novel method for modeling politics in organizations that builds on the model of intra-organizational conflict in March (1962), which we call “subcoalition cluster analysis” (sCCA). March (1962) identifies a set of necessary conditions for political analysis of decision making in an organization: (1) there are groups with consistent preferences over outcomes; (2) the groups’ preferences conflict. Political analysis thus requires identification of groups of individuals in organizations with preferences capable of being aggregated consistently, which March defines as “subcoalitions” or “subsystems,” determining the consistent preference ordering for each subcoalition, and identifying the extent to which the subcoalitions’ preferences disagree. The main contribution of sCCA is that it identifies subcoalitions with consistent preferences that are in conflict without placing additional restrictions on the structure of individual preferences. Existing clustering methods such as factor analysis and principal components analysis combined with spatial clustering routines allow the analyst to identify groups of individuals with *similar* preferences, but not necessarily groups of individuals with *consistent* preferences (Jakulin et al. 2009; Van Gunten, Martin, and Teplitskiy 2016). Voting and other preference aggregation methods reveal a collective consistent set of preferences over outcomes, but offer no insight into the potential existence of competing subcoalitions inside of organizations (Saari 2001; Ossadnik, Schinke, and Kaspar 2016).

Quantitative methods for measuring legislative conflict in political science such as NOMINATE that arrange individuals on a unidimensional political spectrum do permit the identification of coalitions (e.g., political parties), coalition preferences (e.g., party platforms), and the extent of conflict (e.g., distance between party medians), but also place an unreasonable “single-peaked” profile restriction on the

preferences of individuals in the organization (Austen-Smith and Banks 1999; Poole and Rosenthal 2001; Saari and Valognes 1999). These methods rely heavily on empirical research demonstrating that elected representatives in the U.S. can be faithfully characterized as having unidimensional spatial preferences (Poole and Rosenthal 1987; Poole, Rosenthal, and Koford 1991). While this assumption is overly restrictive for describing coalition politics in many organizations, the widespread success of this empirical research program in legislative studies should be a source of inspiration for management scholars (Bonica 2014; Martin and Quinn 2002; Poole and Rosenthal 2001). What management scholars lack, and what is proposed here, is an analogous methodology flexible enough to permit analysis of collective decision making in organizations where the single-peaked assumption is unreasonable.

In our paper, we first describe sCCA, emphasizing how it differs from prior clustering and preference aggregation routines. Then, we apply sCCA to two empirical contexts: Wikipedia and the Baseball Writers' Association of America (BBWAA). We show how the failure to conceive of Wikipedia editors as two subcoalitions led some to misinterpret the causes for the decline in the number of active editors on the platform and underestimate the backlash to the introduction of editing tools intended for novice editors. Then, we show how the emergence of subcoalitions surrounding questions related to the performance-enhancing drug use of Hall of Fame candidates became self-reinforcing through changes to the structure of social influence. Finally, we examine the robustness of sCCA and compare it to the performance of principal components analysis (PCA) and  $k$ -means clustering, which is a method frequently used to identify clusters of preferences in organizations.

## **2. Introducing sCCA**

### **2.1 Consistent Subcoalition Preferences**

March (1962) requires that a subcoalition has a “consistent preference ordering,” i.e., “an ordering [of states] such that for any realizable subset of possible states of the system there exists one state that is at least as good as any other state in the subset” (p. 663). A subcoalition thus can assign ordinal ranks to all potential outcomes, where outcomes with higher ranks are preferred to those with lower ranks. The major contribution of sCCA is that it takes a list of individual preferences over outcomes and sorts the individuals into groups with internally-consistent, externally-conflicting preferences.

There are two types of algorithms commonly used to construct a collective ranking out of a list of individual preferences. The first and most common is a voting system. Individual decision makers allocate points among potential outcomes, which are then ranked based on the number of points they receive. Voting thus determines a consistent preference ordering in a bottom-up fashion. Some voting systems are simple (e.g., everyone can allocate one point among the potential outcomes and the outcomes are ranked according to the tally of points received) and some are quite complex (e.g., voters can allocate

many points; voters can have different numbers of points to allot; the set of outcomes can be sequentially culled or expanded across several stages, etc.). A major drawback of a voting system in the context of identifying subcoalitions is that, if presented with two potential hierarchies of outcomes, there is not a straightforward way of determining which one an individual voter prefers.

An alternative method for finding collective preferences involves selecting the consistent preference ordering that maximizes the utility of individuals polled (Nurmi 2004; Saari and Merlin 2000). Finding the maximum-utility preference ordering requires an algorithm that sorts among all possible consistent preference orderings, assigns a score to each, and identifies the best one. In contrast to voting, this algorithm, sometimes called a distance-based method, takes a top-down approach. The advantage of distance-based methods is that they allow the analyst to identify the best consistent preference ordering as well as the extent of collective and individual disagreement with it. This feature permits the identification of stable subcoalitions within organizations.

## 2.2 Distance-based Collective Preference Profiles

The major issue with many distance-based methods of preference aggregation is computational. The most studied distance-based method in the literature on social choice, called the Kemeny rule, involves finding the collective preference ordering that minimizes the sum of individual violations of binary comparisons of outcomes. For example, imagine a hiring committee choosing between hiring [A]lice, [B]ob, and [C]harlotte. If the collective preference is  $A$  is better than  $B$  is better than  $C$  (i.e.,  $A \succ B \succ C$ ) and there is a committee member who prefers  $C$  to  $A$  ( $C \succ A$ ) and  $B$  to  $A$  ( $B \succ A$ ), then that member contributes two units of disutility to the collective. Finding the consistent preference ordering that satisfies the Kemeny rule, however is NP-hard (Bartholdi, Tovey, and Trick 1989; Conitzer, Davenport, and Kalagnanam 2006).

Gupte et al. (2011) identifies another distance-based method where the best consistent preference ordering can be identified in polynomial time. This method places greater weight on larger preference violations. For example, if the hiring committee again prefers  $A$  to  $B$  to  $C$  while a committee member prefers  $C$  to  $A$ , that is a bigger violation of the collective preference than if the same committee member preferred  $B$  to  $A$ . Disutility in Gupte et al. (2011) is calculated in the following way: Let  $r(u)$  indicate the ordinal rank of alternative  $u$  and  $r(v)$  indicate the rank of alternative  $v$ . The disutility associated with an individual binary comparison of alternatives,  $v \succ u$ , is  $\max(r(u) - r(v) + 1, 0)$ . Thus, for a group ranking of  $r(A) = 2$ ,  $r(B) = 1$ , and  $r(C) = 0$ , and an individual member with  $C \succ A$  and  $B \succ A$ , the member contributes 5, i.e.,  $(2+1) + (1+1)$ , to collective disutility. Finding the preference ordering for a group of individuals involves minimizing collective disutility, which is the solution to a linear program.

This method, therefore, allows the analyst to efficiently find a collective group preference ranking *and* a measure of individual disutility with that collective preference ranking.

### 2.3 Assigning Individuals to Coalitions

We permit organizations to be characterized by subcoalitions with competing preference rankings and assign individuals to subcoalitions based on the extent to which their preferences and the subcoalition's preferences align. Subcoalitions and their associated preference rankings are identified in order to split the organization into groups with minimal intragroup discord and maximal intergroup disagreement.

Returning to the example of the hiring committee, imagine there are two subcoalitions with different preferences over candidates. The first coalition prefers candidates with shorter names. As such,  $Bob \succ Alice \succ Charlotte$  and  $r(B)=2$ ,  $r(A)=1$ , and  $r(C)=0$ . The second prefers candidates whose names end in vowels, so  $A=C \succ B$  or  $r(A)=r(C)=1$  and  $r(B)=0$ . A member who prefers  $C \succ A$  and  $B \succ A$  would contribute two units of disutility to the first coalition, i.e.,  $(1+1)+(0)$ , and three to the second coalition, i.e.,  $(0+1)+(1+1)$ , so we would assign the member to the second subcoalition.

We require that partitions of individuals into subcoalitions satisfy the Pareto criterion: no individual would be better off being a member of a different subcoalition than the one to which they are assigned. Through an iterative process of finding the minimum-disutility consistent preference ranking for each subcoalition, and then assigning members to the subcoalition to which they would contribute the least disutility, the algorithm finds a stable sorting of individuals into subcoalitions. As this MM algorithm finds a local minimum, not a global one, we initialize the search for minimum-disutility subcoalitions with different initial conditions seeded using a close variant of the k-means++ algorithm and choose the one with the lowest collective disutility (Arthur and Vassilvitskii 2007).

### 2.4 sCCA Output and Inference

sCCA calculates a series of indicators relevant to the study of conflict in organizational decision making. For a given number of subcoalitions, it returns the preference ordering of outcomes for each subcoalition, the allocation of individuals across subcoalitions, and the collective level of disutility. In addition, it returns the individual level of disutility with each of the subcoalitions and the extent to which each subcoalition satisfies the goals of its members.

We also propose a set of heuristic methods to identify whether there is subcoalition conflict in an organization and, if there is, how many subcoalitions exist. The first, the "elbow method," involves identifying a discontinuity in the slope of the function that maps the number of subcoalitions to overall disutility. A more robust approach compares the observed relationship between the number of

subcoalitions and overall disutility to the expected relationship if there were no subcoalitions in the organization. In this method, based on the “gap statistic” proposed in Tibshirani, Walter, and Hastie (2001), we create multiple synthetic datasets in which the aggregate preferences over outcomes for the organization remains the same, but the preferences for each individual are randomized. Then, we compare the observed relationship between subcoalition count and disutility to the average across the synthetic datasets.

An alternative approach examines the stability of the identified subcoalitions with respect to random perturbations in the organization’s membership (von Luxburg 2010). In this method, we use the bootstrap to sample individuals with replacement and examine the extent to which the subcoalitions identified in each sample align with the subcoalitions in the observed data. If individuals in the bootstrap samples are consistently assigned to the same subcoalitions, then it provides additional confidence that there are meaningful subcoalitions in the organization. In contrast, if many individuals who are identified in the same subcoalition in the observed data are assigned to different ones in the bootstrap samples, then it is less likely that the subcoalitions in the observed data represent meaningful conflict. Our stability score equals the average share of individuals who are identified in the same subcoalition as in the observed data across bootstrap samples.

### **3. Applying sCCA: Wikipedia and the Baseball Writers’ Association of America**

Next, we apply sCCA to two cases: Wikipedia and the BBWAA. Both cases represent examples of organizations trying to understand the values, preferences, and goals of their membership at critical decision points when the existing routines that facilitated internal cooperation appeared to be breaking down. In the case of Wikipedia, the problem was declining editor participation. For the BBWAA, the problem was the growing difficulty in finding consensus candidates for the National Baseball Hall of Fame.

#### **3.1 Wikipedia**

Following rapid growth in editor participation from 2003 to 2007, the number of new editors active on Wikipedia suddenly began to decline (Suh et al. 2009). Wikipedia responded with new tools and infrastructure designed to make editing easier, but with little impact. In his 2009 “State of the Wiki” address, founder Jimmy Wales argued that “if active contributors continue to decrease, there may not be a large enough cohort to ‘look after’ Wikipedia” (Wales 2009). Concerns that a smaller, more homogeneous editor population would jeopardize the future health of the Wikipedia project prompted a focus on editor recruitment and retention as a part of the 2011-2015 strategic plan. As part of the plan,

Wikipedia designed a series of surveys in order to understand the experience of editors and recommend strategic changes.

Inside Wikipedia, there were competing theories about the causes of the weakening participation. The first was that editing articles of Wikipedia was too technically demanding. By creating an easier user interface for editors, more people would be attracted to the community. The second was a culture hostile to new editors. Rather than welcoming new participants and encouraging them to learn the norms of the platform, experienced editors would delete the edits of new editors that did not conform with Wikipedia's guidelines without providing constructive feedback. The editor survey offered the potential to get anonymous feedback from participants in order to prioritize changes to the platform and community standards.

**3.1.1 Data.** We analyze a question on the 2012 Wikipedia Editor Survey in which respondents were asked to identify “the most important problems that have affected you personally, making it harder for you to edit.” Editors were given a list of ten potential problems and had the opportunity to select up to three they found most relevant to their experience. In Wikipedia's topline analysis, approximately two-fifths of the over 17 thousand participants surveyed criticized “editors who feel like they ‘own’ specific articles and don't want others to contribute to them.” Thirty-two percent felt that there were “too many rules and policies.” Between twenty and thirty percent responded affirmatively to four of the remaining eight options. The results supported the theory that Wikipedia's editorial culture was its biggest problem, but that the other concerns were valid as well.

**3.1.2 Results.** We use sCCA to analyze internal disagreement over Wikipedia's most important problems. We restrict the data set to editors who contribute primarily to English Wikipedia and who indicate that they at least sometimes participate as editors. For each respondent, we create a set of binary comparisons of the ten candidate problems. We assume that each editor finds the problems they voted for more serious than the ones they did not vote for. After excluding editors who selected none of the problems, we are left with 1,416 individuals in our dataset

We present the ranks of problems conditional on the number of subcoalitions in the data in Figure 1. When analyzed as a single group, our data reflect the results of Wikipedia's topline analysis. The biggest problem is editors who “feel like they own specific articles.” Tied are “too many rules and policies,” the editing interface being “hard to use,” editors not being “fun to work with,” “lack of support from other editors,” “lack of access to research materials,” and “criticism of you and your work.” Lowest ranked were problems related to the software, “warning messages on your talk page,” and “harassment” by other editors.

However, when we analyze Wikipedia as two subcoalitions, we identify substantial internal conflict. One subcoalition (“B”), which includes 690 editors, identifies editorial ownership of articles,

editors not being fun to work with, and criticism of one’s work as the most serious problems. The second subcoalition (“A”), which includes 726 editors, instead points to the number of rules, difficulties using the editing interface, and lack of access to research materials as the most important problems. Two of the problems that were most important to the latter group were least important to the former. The top line results thus obscure a division among editors about the major problems with the platform. Subcoalition B identifies cultural problems as the most important problems for Wikipedia to address. Subcoalition A identifies technical problems as more important. This divide is also captured by the difference between the average disutility of editors with the preferences of their own subcoalition – 8.3 and 8.8 for subcoalitions A and B, respectively – against average disutility of editors with the preferences of the other subcoalition – 21.6 and 23.8 for subcoalitions A and B, respectively.

When we use sCCA to identify more subcoalitions, we observe more moderate disagreement across the groups. When we allocate the editors among three groups, the third groups preferences are a mix of the two identified previously. The third group’s top problem is a lack of access to research materials, but their second tier of problems includes both issues with other editors and difficulties using the editing platform. For four groups, similarly, one group continues to emphasize cultural issues, one emphasizes technical ones, and the other two are a blend.


<b>Aggregate Group Preferences</b>	<b>Rank</b>
Editors who feel like they own specific articles and don't want others to contribute to them	3
Too many rules and policies The editing interface is hard to use Editors who are not fun to work with Lack of support from other editors Lack of access to research materials like scholarly articles and books Criticism of you and your work, including reverts and deletions	2
The software for doing maintenance tasks (such as patrolling and fighting vandalism or spam) is too hard to use Warning messages on your talk page Harassment by other editors	1

<b>2-Group Top Preferences: Subcoalition A</b>	<b>2-Group Top Preferences: Subcoalition B</b>
Too many rules and policies The editing interface is hard to use Lack of access to research materials like scholarly articles and books	Editors who feel like they own specific articles and don't want others to contribute to them Editors who are not fun to work with Criticism of you and your work, including reverts and deletions
<b>2-Group Bottom Preferences: Subcoalition A</b>	<b>2-Group Bottom Preferences: Subcoalition B</b>
Harassment by other editors	The editing interface is hard to use The software for doing maintenance tasks (such as patrolling and fighting vandalism or spam) is too hard to use Lack of access to research materials like scholarly articles and books

3-Group Top Preferences: Subcoalition A	3-Group Top Preferences: Subcoalition B	3-Group Top Preferences: Subcoalition C
Too many rules and policies The editing interface is hard to use	Editors who feel like they own specific articles and don't want others to contribute to them Editors who are not fun to work with Criticism of you and your work, including reverts and deletions	Lack of access to research materials like scholarly articles and books
3-Group Bottom Preferences: Subcoalition A	3-Group Bottom Preferences: Subcoalition B	3-Group Bottom Preferences: Subcoalition C
Lack of support from other editors Lack of access to research materials like scholarly articles and books	The editing interface is hard to use	Harassment by other editors

**Figure 1. Wikipedia: Aggregate Group and Subcoalition Preferences**

Next, we use the three heuristic methods to identify the number of subcoalitions in the editor population. The left panel of Figure 2 reports aggregate disutility levels conditional on subcoalition count for one through five groups. The “elbow” that appears at two groups suggests that a meaningful second subcoalition may be present. The center panel reports the difference between the aggregate disutility level in the observed data and aggregate disutility when individual preferences are randomized in 100 synthetic datasets, along with 95 percent confidence intervals. The mean disutility difference between observed clusters for a subgroup count of two is significantly different from the mean difference for the combined group. The right panel reports the stability score for 100 simulated datasets in which individuals are sampled with replacement. The score remains as high as 0.9 for two subgroups, dropping greatly if a third is added. All three of the methods therefore agree that the editors are well-characterized by two subcoalitions.



**Figure 2. Wikipedia: Disutility, Gap Statistic, and Stability Score by Number of Subcoalitions**

We next examine what types of editors fall in each subcoalition and how the subcoalitions differ in terms of organizational goals. We use logistic regression to predict membership in subcoalition B conditioned on the amount of time spent on Wikipedia, editorial experience, age, gender, and country of origin. We find that editors who spend more than an hour in the last week contributing to Wikipedia are 16 percent more likely to be a member of Subcoalition B ( $\beta=0.067$ ,  $se=0.012$ ). Further, each additional year of editing experience makes an editor three percent more likely to be a member of subcoalition B ( $\beta=0.11$ ,  $se=0.022$ ). Both effects are significant at the 0.01 confidence level. Gender, age, and country-of-origin are not significant predictors of subcoalition membership. Experienced editors and avid users are more likely to be a member of the subcoalition who believe other editors are the most significant problem with Wikipedia. Inexperienced and less frequent users are more likely to point to the problems associated with the editing platform.

This same schism exists when using subgroup membership to predict editors' values related to resource allocation. On the same survey, editors were asked to allocate a \$100 donation among different priorities for the Wikimedia foundation. On average, Subcoalition B allocated three more dollars than subcoalition A to technical upgrades to the platform and features directed at experienced editors. Subcoalition A allocated three more dollars than subcoalition B to features directed at new editors and features for readers. Both of these differences are significant at the 0.01 confidence level as well.

These results reflect the tension at the heart of Wikipedia's efforts to stem its declining editor base (Simonite 2013; Halfaker et al. 2013). According to the survey, the biggest problems for new editors were primarily related to ease of use. However, the more powerful subcoalition of avid users believed that the biggest problems at Wikipedia related instead to editorial culture. Difficulty using the interface and lack of access to research materials were the least important problems reported by avid users.

This conflict between the subcoalitions erupted when English Wikipedia introduced its new Visual Editor in 2013 in order to make it easier for new editors on the site. Existing editors complained that the bug-marred rollout was disrespectful to committed contributors, a waste of donated money, and reflected a lack of commitment to Wikipedia's core mission. In response, editors "staged a rebellion" by adopting user-written code that allowed them to circumvent the visual editing platform, leading Wikipedia to grudgingly change the Visual Editor interface from opt-out to opt-in (Orlowski 2013; Sampson 2013). Had leadership been aware that the rationale for this major change to the editing interface was in conflict with the experience of its most avid users, they would have known they needed to be more careful during the Visual Editor's introduction.

### **3.2 Baseball Writers Association of America (BBWAA)**

The most important public function of the BBWAA is determining which retired baseball players are inducted into the National Baseball Hall of Fame. Beginning in 2013, however, the election process became embroiled in controversy because none of the candidates received the seventy-five percent of votes required for induction. For the first time in a half-century, the National Baseball Hall of Fame had a summer induction ceremony with no living inductees. The BBWAA considers Hall of Fame voting to be “the ultimate privilege” for its members (Chappell 2014). Yet rumors swirled that, if the BBWAA continued to be unable to serve its purpose of selecting candidates for the Hall of Fame, the Hall of Fame would begin to explore other options for determining who should be inducted (Keri 2013).

Prior to the election, baseball writers acknowledged that 2013’s crop of new candidates were particularly controversial. Barry Bonds and Roger Clemens, two players who are among the most distinguished in the game’s history but whose candidacies were shrouded in rumors of performance-enhancing drug (PED) use, were submitted to the baseball writers for consideration for the first time. Absent the PED question, both Bonds and Clemens would be near-consensus Hall of Fame picks. While other marginal Hall of Fame candidates suspected of PED use had been considered by the voters in prior years, this was the first time the writers had to judge players with such extraordinary on-the-field performance who were also suspected of using PEDs. Despite many baseball writers acknowledging a crowded, controversial ballot with many qualified newcomers, few suspected that no candidates would be elected.

*New York Times* columnist Nate Silver hypothesized that the tenor of the debate had changed during the 2013 voting: “Instead of the typical friendly arguments about how a player’s lifetime accomplishments might be weighed against how dominant he was in his best seasons, or how to compare players at different positions, the writers are now spending most of their time arguing about who used steroids and when, and how this should affect Hall of Fame consideration” (Silver 2013). Disagreements over players prior to 2013 tended to be idiosyncratic, as each voter applied slightly different criteria for measuring greatness. The 2013 disagreement, instead, felt more like competing factions within the BBWAA: the writers who felt that PED rumors disqualified a player from Hall of Fame induction against those willing to weigh possible PED use against on-the-field performance.

Following the public announcement of the 2013 election results, baseball writers began to publicly complain that the voting system had become dysfunctional in its current form. Some argued that the problem lay in the ambiguous voting criteria, or the 10-vote cap on the number of candidates each voter could select, or an aging voter pool who had grown out of touch with the modern game (Caple 2012; Keri 2013; Silver 2013). The BBWAA responded with a series of changes to the voting procedures. In 2014, the maximum amount of time that a player could remain on the Hall of Fame ballot was shortened from 15 years to 10 years in order to decrease the number of players on future ballots. In 2015,

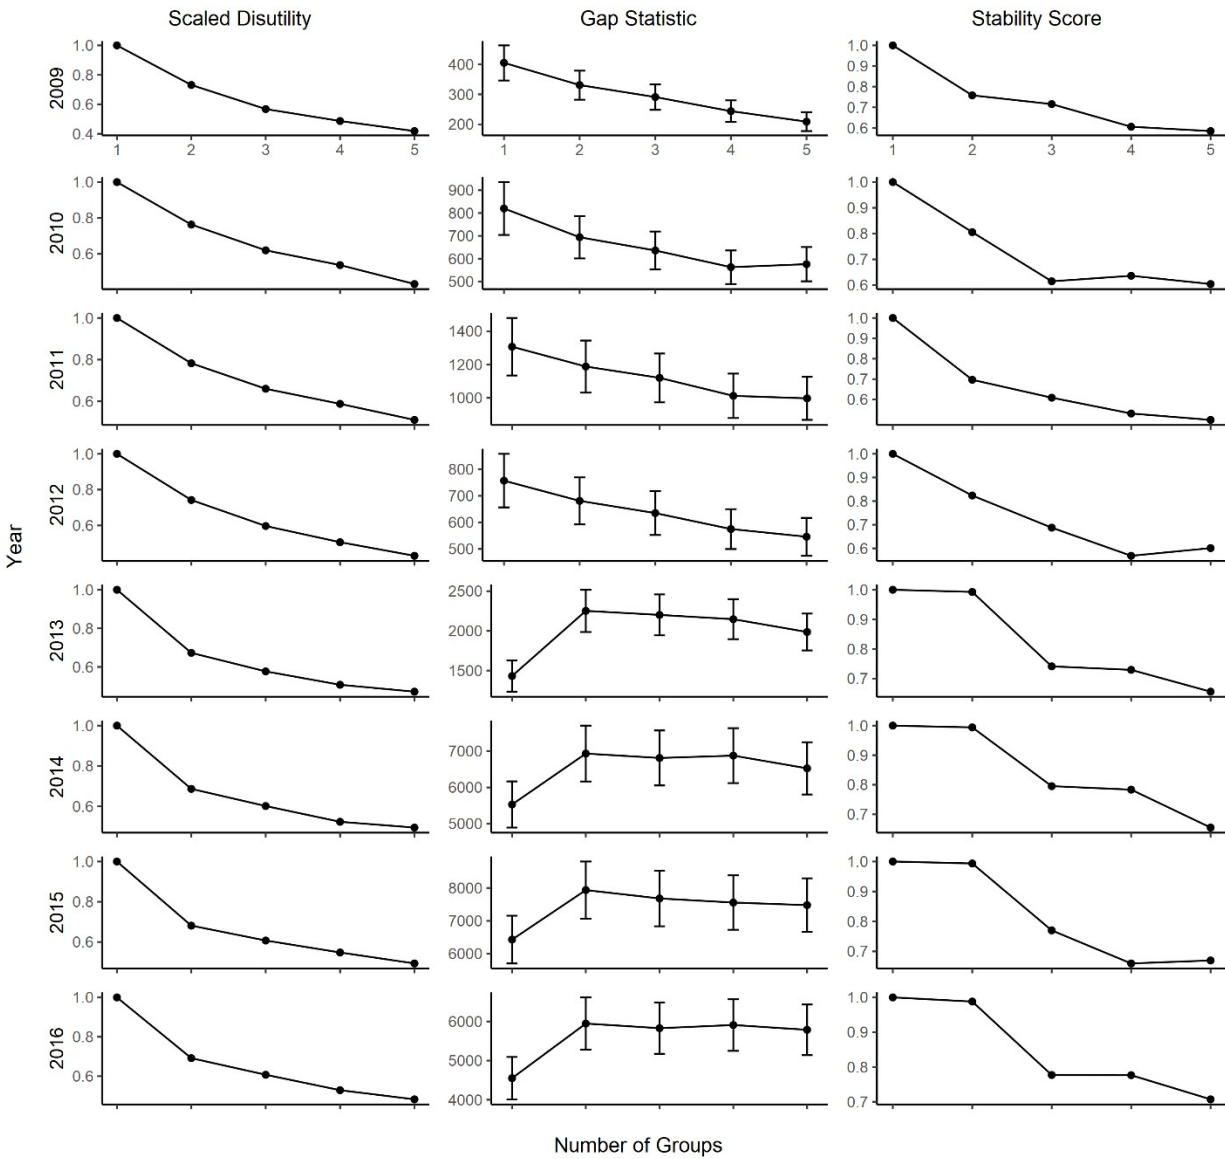
the BBWAA also decided to limit the voter pool to members who have actively covered baseball in the past ten years.

**3.2.1 Data.** Many baseball writers voluntarily make their Hall of Fame ballots public. Beginning with the 2009 class of candidates, a Twitter user @leokitty began tracking these public votes. In 2013, Ryan Thibodeaux took over the ballot tracking effort in a public spreadsheet posted on his website. In addition, starting in 2013, the BBWAA began to publish the votes of willing writers. Over time, the number of public ballots has increased, from 60 in 2009, to 168 in 2013, to 312 in 2017.

We use these public ballots to study the nature and evolution of the conflict among the baseball writers over Hall of Fame candidates. From each writer's ballot, we create a series of binary comparisons between players. We assume that writers prefer players for whom they voted to those for whom they did not.

**3.2.2 Results.** We explore whether the nature of disagreement over which candidates deserve induction into the Hall of Fame and whether it changed in 2013, the first year that Bonds and Clemens were on the Hall of Fame ballot. Then, we examine whether subcoalition conflict in the BBWAA has changed how voters use other writers' prior votes to update their beliefs about active candidates and inform their future voting behavior.

Disutility, Gap Statistic, and Stability Score by Number of Groups  
Hall of Fame Dataset: 2009-2016



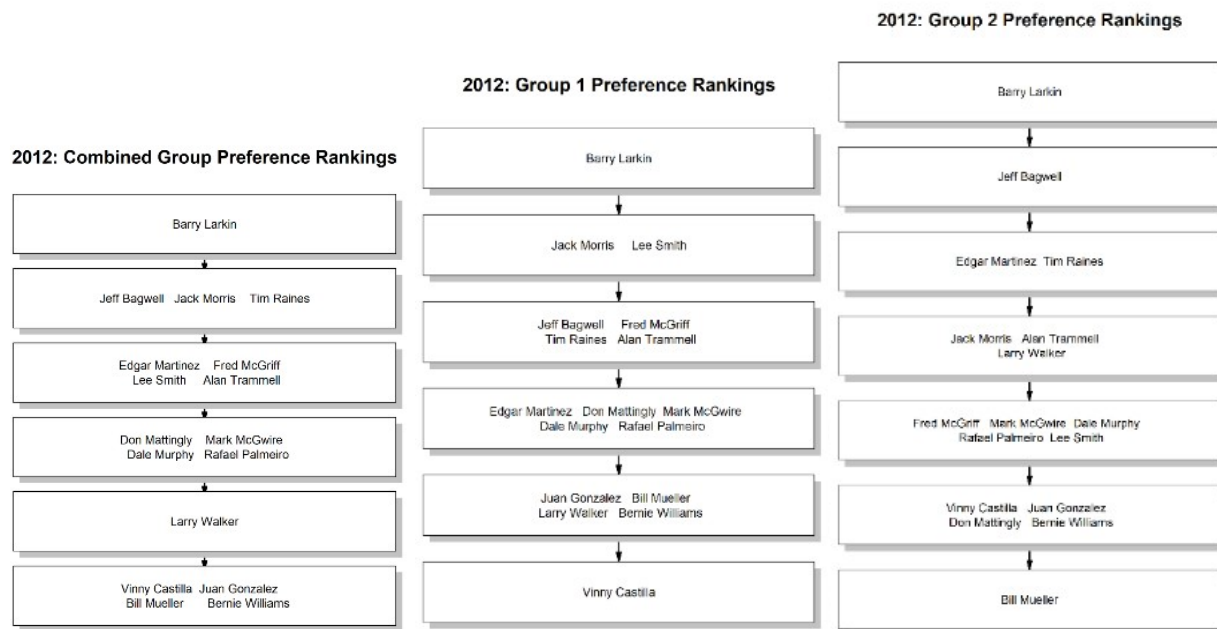
**Figure 3. BBWAA: Disutility, Gap Statistic, and Stability by Number of Subcoalitions**

We begin by examining the aggregate level of disutility conditional on the number of subcoalitions for each year from 2009 through 2016, displayed in the first column of Figure 3. Because there are different numbers of writers and candidates in each year, we scale disutility so that it equals 1 when the writers are treated as a single group. The longitudinal data identify a clear elbow beginning with the 2013 voting that did not exist in the earlier years. From 2009 through 2012, the relationship between the number of subcoalitions and overall agony is roughly linear. From 2013 through the present, it is

highly nonlinear, with a kink at two groups. This aligns with the hypothesis that there is one subcoalition from 2009 through 2012, but two distinct subcoalitions from 2013 on.

The second column of Figure 3 reports the difference between the observed aggregate disutility level and the disutility in synthetic datasets with randomized preferences, which also indicates a clear shift beginning in 2013. Prior to 2013, the difference between the observed and randomized data by subgroup count is relatively linear and monotonically decreasing. After 2013, the mean difference in disutility between the observed and synthetic data for two groups is significantly different from the mean difference for the aggregate group.

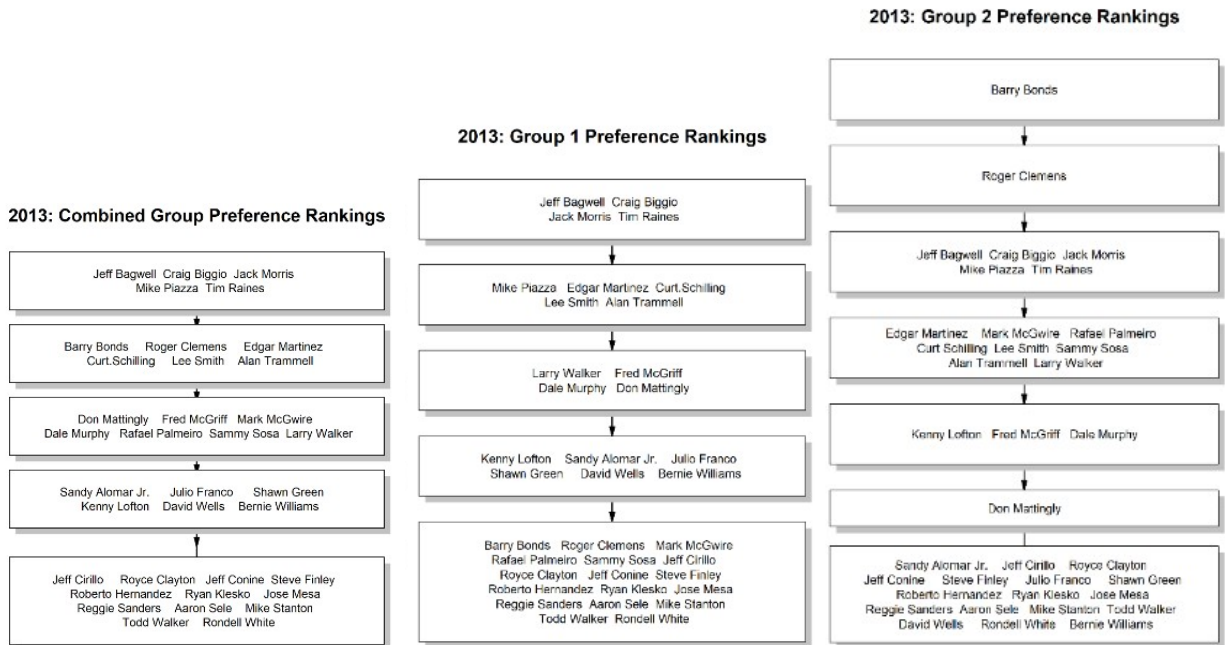
The third column of Figure 3 reports the stability score. Again, the pattern changes after 2013. Prior to 2013, many voters who are identified as a part of the same subcoalition in the observed data are part of different subcoalitions in bootstrap samples. After 2013, individuals in the bootstrap samples are placed into the same subcoalition as they are in the observed data nearly 100 percent of the time. Once again, the three methods agree that the BBWAA is well-characterized as competing subcoalitions.



**Figure 4. BBWAA: 2012 Combined and Subcoalition Preferences**

We next focus on the transition from 2012 through 2013 to better understand the split in the electorate. In Figure 4, we display the ranks of each candidate if the electorate is treated as one subcoalition, and then the ranks for each subcoalition when the electorate is divided into two subcoalitions. In 2012, Lee Smith and Edgar Martinez are the primary sources of disagreement between the two subcoalitions. Suspected PED-users Mark McGwire and Rafael Palmeiro are ranked behind seven

other candidates in both groups, indicating that the PED issue was yet not a major source of disagreement among the voters.



**Figure 5. BBWAA: 2013 Combined and Subcoalition Preferences**

In 2013, in contrast, the differences in ranks between the two subcoalitions is much more significant. For the PED-insensitive subcoalition, Barry Bonds has no candidates ranked above him and Roger Clemens just has Bonds above him. For the PED-sensitive subcoalition, both Bonds and Clemens have 45 candidates ranked above them. Also of note is the significant disagreement that emerges among more marginal candidates widely perceived as non-PED users, like Don Mattingly and Fred McGriff. In the PED-insensitive coalition, Mattingly and McGriff have 18 and 15 candidates ranked above them, respectively. For the PED-sensitive coalition, both have just 9 ranked above them. The ranks for 2013 are displayed in Figure 5.

Next, we use sCCA to examine whether changes to patterns of social influence are contributing to the ossification of subcoalition boundaries and difficulty reaching consensus on candidates. Baseball writers frequently cite each other's analysis and prior voting behavior when publicly justifying their voting choices, creating a bandwagon effect for some candidates. Although some candidates are elected in their first year on the ballot, most face a slower process in which writers gradually become convinced of their value following many years of debate, balloting, and re-evaluation. This multiyear conversation about whether a player is worthy of induction has tended to generate coordination among the electorate.

As a result, every player but one who has received 50 percent of the vote from the writers has eventually been inducted.

However, if baseball writers' evaluations of candidates are influenced more strongly by writers in their own subcoalition than by writers outside their subcoalition, then the social influence of the out-group writers would be weakened and the cross-subcoalition divergence in evaluation would be self-reinforcing. In contrast, if baseball writers are influenced uniformly by members of their own subcoalition as by members of the opposing one, the divergence in player evaluation would be more likely to weaken over time. The data thus provide an opportunity to apply the theory of "associative diffusion" proposed in Goldberg and Stein (2018).

We examine the dynamics of social influence among the writers using a series of regression analyses. These models estimate whether a writer,  $i$ , voted for a candidate,  $j$ , in year,  $t$ , where a vote by a writer for a candidate in a year,  $v_{ijt}$ , is evaluated as a 1 and the absence of a vote is evaluated as a 0. In all years, we divide the writers into two subcoalitions. The inclusion of a lagged dependent variable complicates estimation of a discrete choice model (Honore and Kyriazidou 2000), so we estimate linear probability models of the following form:

$$v_{i,j,t+1} = \delta v_{i,j,t} + \beta_1 E(v_{i,j,t} | \text{in-subcoalition}) + \beta_2 E(v_{i,t,t} | \text{out of subcoalition}) + \lambda_t + \epsilon_{i,j,t}$$

$\beta_1$  measures the influence of in-subcoalition writers on a writer's voting decision in the subsequent year.

$\beta_2$  measures the influence of other writers on a writer's voting decision in the subsequent year.  $\delta$

measures the persistence of a writer's voting decision.  $\lambda_t$  is a set of year fixed effects. Then, we examine whether the pattern of influence changes post-2013 using the following model:

$$v_{i,j,t+1} = \delta v_{i,j,t} + \beta_1 E(v_{i,j,t} | \text{in}) + \beta_2 E(v_{i,t,t} | \text{out}) + \beta_3 E(v_{i,j,t} | \text{in}) \times \text{Post 2013} + \beta_4 E(v_{i,t,t} | \text{out}) \times \text{Post 2013} + \lambda_t + \epsilon_{i,j,t}$$

$\beta_3$  and  $\beta_4$  measure the change in in- and out-subcoalition influence, respectively, after 2013. Two-way cluster robust standard errors are reported for all regressions, where errors are clustered on the level of the writer and candidate.

The results, which are presented in Table 1, are consistent with a self-reinforcing split among the writers after 2013. Model 1 indicates that writers are more influenced by in-subcoalition writers than by out-subcoalition writers, because  $\beta_1$  exceeds  $\beta_2$ . However, Model 2, demonstrates that this difference is isolated to the post-2013 period. Prior to 2013, writers are about equally influenced by in- and out-subcoalition writers, consistent with the absence of meaningful subcoalitions in the data. After 2013, however, they are significantly less influenced by out-subcoalition writers and marginally more influenced by in-subcoalition writers. The difference between  $\beta_3$  and  $\beta_4$  is significant at the 0.01-level,

which indicates that the change in influence from in-subcoalition writers is different than the change in influence from out-subcoalition writers.

	(1)	(2)
	$v_{i,j,t+1}$	$v_{i,j,t+1}$
$v_{i,j,t}$	0.634*** (0.0178)	0.634*** (0.0178)
$\beta_1$	0.315*** (0.0278)	0.283*** (0.0367)
$\beta_2$	0.127*** (0.0265)	0.242*** (0.0319)
$\beta_3$		0.0342 (0.0426)
$\beta_4$		-0.135*** (0.0359)
Observations	19,813	19,813
R-squared	0.584	0.585

Robust standard errors in parentheses

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

**Table 1. BBWAA: Logistic Regression Results**

sCCA thus points to a change in the social structure of Hall of Fame voting following the 2013 election. Prior to 2013, conflict over candidates was idiosyncratic and writers learned equally from those who they had more similar preferences and those who had more different ones. After 2013, not only did the writers split into two subcoalitions, one sensitive to PED use and one insensitive to it, but also this split appears self-reinforcing, because writers' future voting decisions are no longer as strongly influenced by voters outside of their subcoalition as by voters inside their subcoalition.

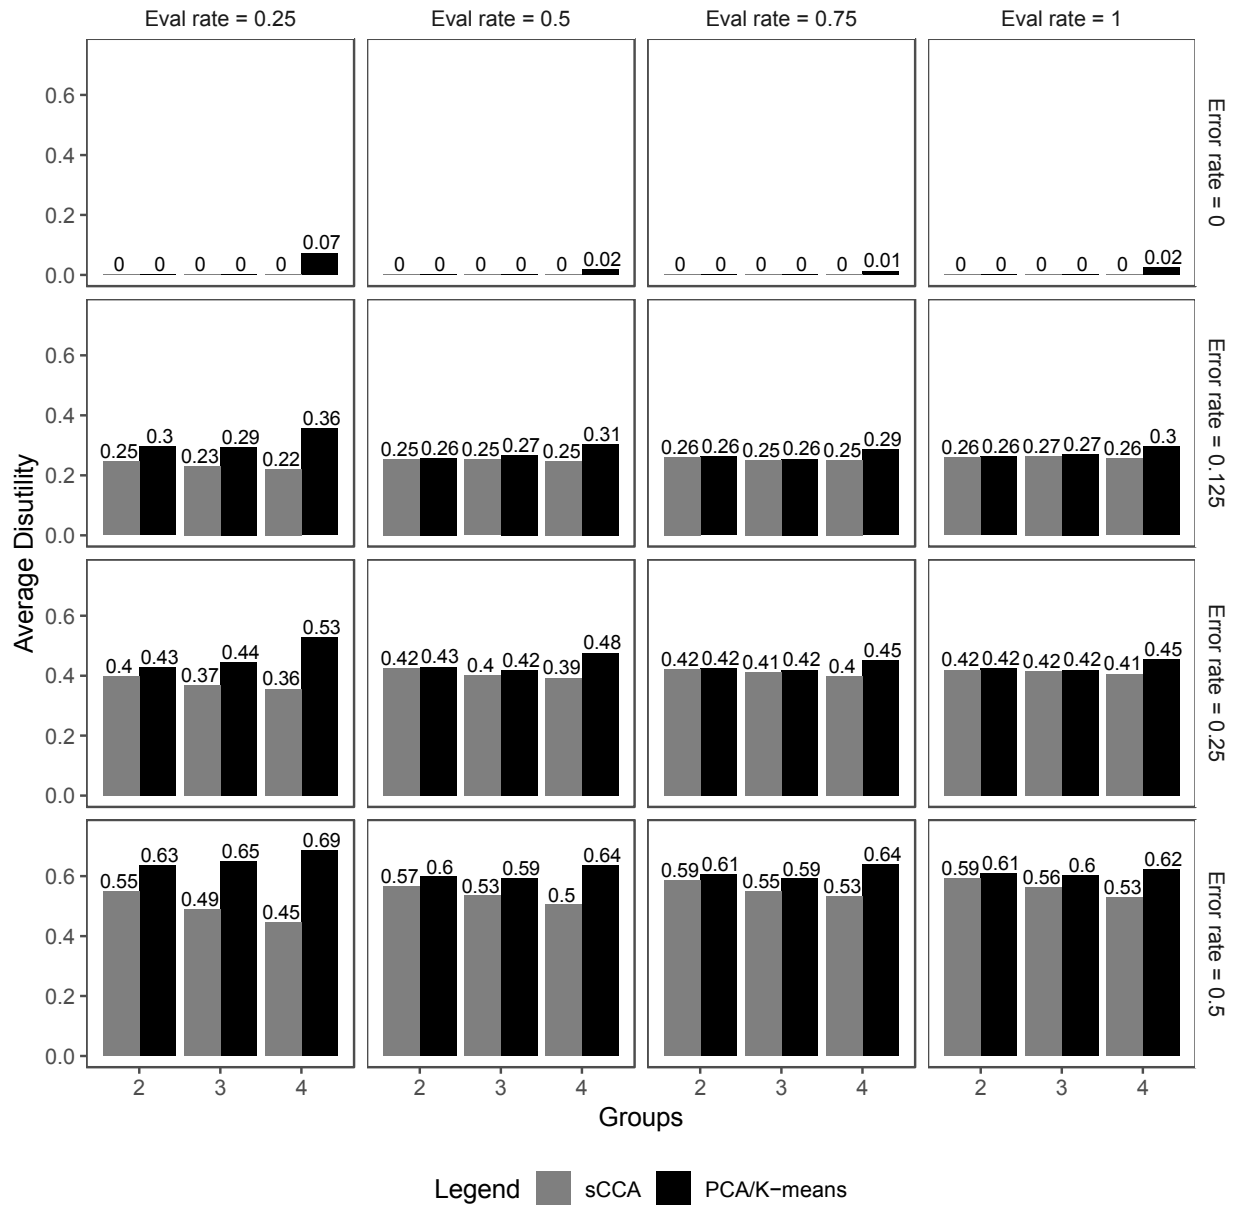
#### 4. Robustness Checks and Comparisons with Prior Methods

Next, we examine the robustness of sCCA to the presence of many subcoalitions, heterogeneous preferences within subcoalitions and to missing data. We use Principal Components Analysis followed by  $k$ -means clustering in order to benchmark the performance of sCCA, because PCA/ $k$ -means is one commonly-used methodology for identifying groups of individuals with similar preferences in sociology and political science.

There are two major differences between sCCA and PCA/ $k$ -means. First, sCCA requires that the identified subcoalitions have consistent preferences. PCA/ $k$ -means, in contrast, clusters individuals based on the similarity of their preferences. Second, sCCA is more flexible in terms of data requirements. In sCCA, individuals are matched to subcoalitions solely based on the disutility that arises from the preferences they report. PCA/ $k$ -means, in contrast, requires that individuals be placed in a metric space prior to the clustering step. As such, missing data on individual preferences need to be imputed. For the PCA/ $k$ -means analysis, we use the regularized iterative PCA algorithm that is a part of R's "missMDA" package version 1.14.

All of our robustness tests include 100 simulated datasets with 96 individuals choosing among ten alternatives. We perform analyses on data with two to four subcoalitions, where each subcoalition has a randomly assigned set of consistent and strict preferences over outcomes. The proportion of all possible pairs of outcomes evaluated by each individual range from  $\frac{1}{4}$  to 1. We also randomly induce disagreement between individuals and their assigned subcoalition with regards to the rank of a particular alternative. This rate of disagreement ranges from 0 to  $\frac{1}{2}$ . We present results from 48 ( $3 \times 4 \times 4$ ) different cases (2, 3, and 4 subcoalitions; evaluation rates of  $\frac{1}{4}$ ,  $\frac{1}{2}$ ,  $\frac{3}{4}$ , and 1; rates of disagreement of 0,  $\frac{1}{8}$ ,  $\frac{1}{4}$ , and  $\frac{1}{2}$ ).

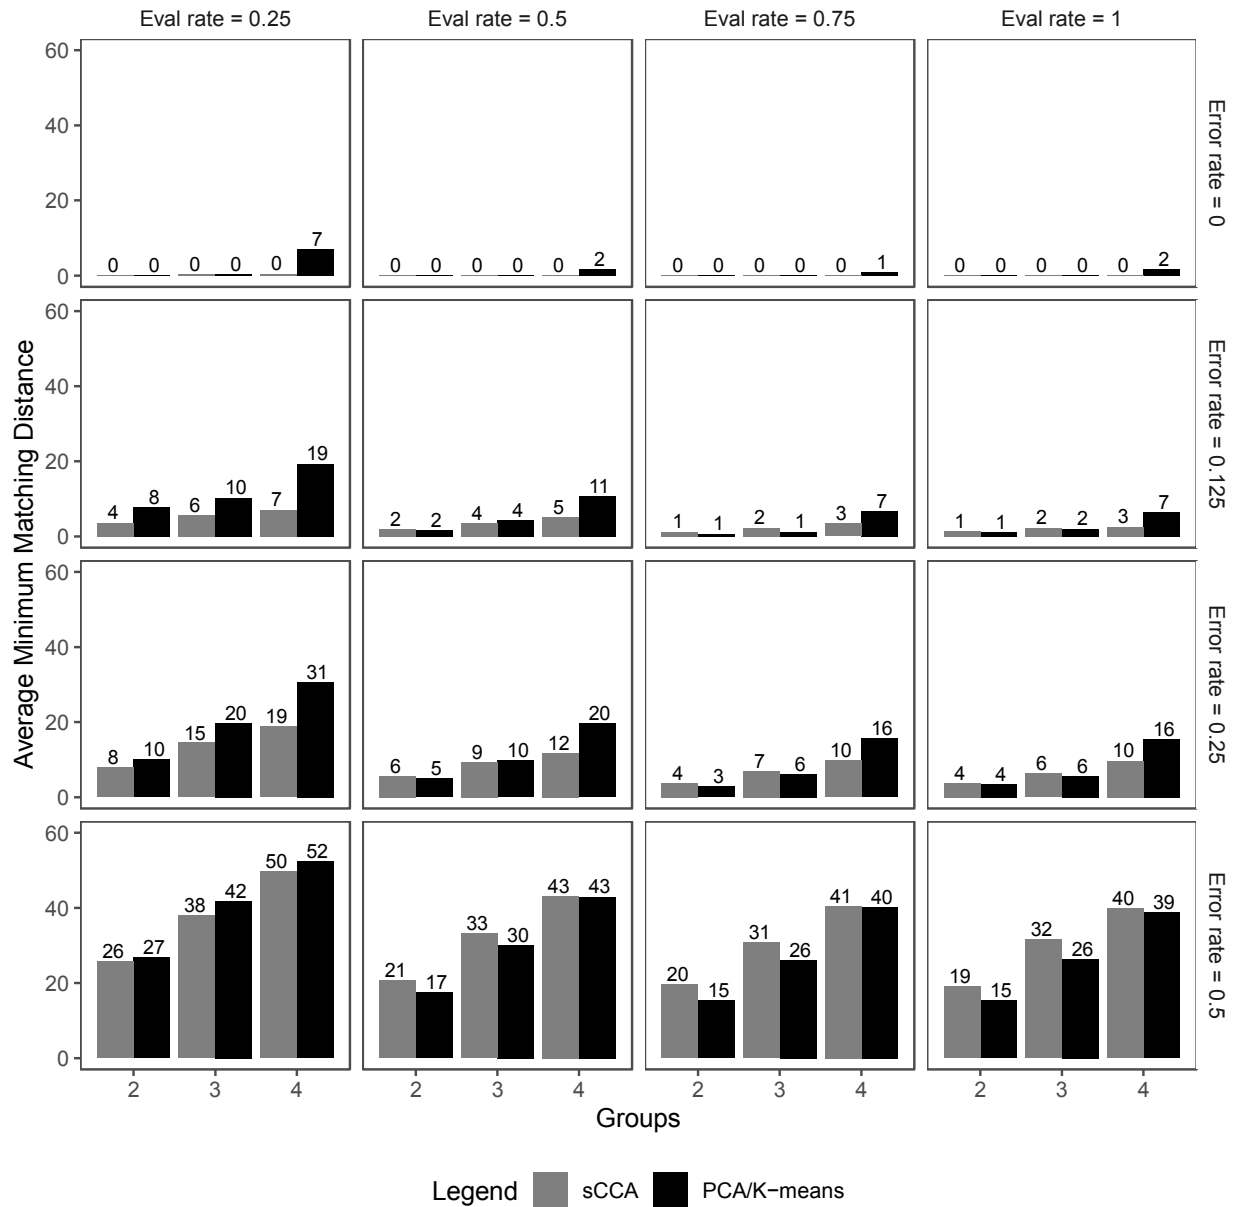
Figure 6 presents the average disutility for each individual in the 100 simulated datasets. The black bars with dots indicate average disutility for sCCA and gray bars indicate average disutility for PCA/ $k$ -means clustering. Average disutility under sCCA is increasing in the level of disagreement and increasing in the number of groups. In all cases, sCCA identifies subcoalitions that produce less disutility than PCA/ $k$ -means clustering, and the divergence between sCCA and PCA/ $k$ -means clustering is increasing with respect to the number of groups, increasing with the extent of disagreement of individuals with their subcoalition, and decreasing with the evaluation rate.



**Figure 6. sCCA vs. PCA/k-means Clustering: Average Disutility**

Figure 7 presents the average minimum matching distance between the true subcoalitions and those identified in the 100 simulated datasets. The average minimum matching distance is 0 when individual preferences are entirely aligned with subcoalition preferences, and is increasing in the extent of disagreement of individuals with their subcoalition. The average minimum matching distance is also increasing in the number of groups and constant with respect to the evaluation rate. sCCA outperforms PCA/k-means clustering at low rates of evaluation and at modest rates of disagreement. At high rates of disagreement and high evaluation rates, however PCA/k-means clustering marginally outperforms sCCA.


This is unsurprising, because PCA/ $k$ -means is not restricted to consistent subcoalition preferences and, as such, does a good job identifying individuals with similar preferences when the ability to aggregate individual preferences consistently is more limited.



**Figure 7. sCCA vs. PCA/ $k$ -means Clustering: Average Minimum Matching Distance**

As sCCA always identifies Pareto optimal subcoalitions with consistent preferences over outcomes, we next examine whether PCA/ $k$ -means identifies groups that (1) are stable under the Pareto criterion and (2) have preferences capable of being consistently aggregated. The first test asks whether

there is an individual who is assigned to a group who would prefer to be in a different group (using the same disutility measure as we use sCCA). If there is, then the PCA/ $k$ -means grouping of individuals fails the Pareto criterion. The second test examines the presence of intransitive cycles when binary preferences over outcomes are defined by Condorcet voting. For example, an intransitive cycle occurs if a majority of group members prefer  $A$  to  $B$ , a majority prefer  $B$  to  $C$ , and a majority prefer  $C$  to  $A$ . We use the probability of a Condorcet cycle to quantify the ease of constructing consistent group preferences for PCA/ $k$ -means. Figure 8 demonstrates that PCA/ $k$ -means clustering fails to identify stable subcoalitions with consistent preferences when individual disagreement with their subcoalition is high, when there are more than two groups in the data, and when the evaluation rate is low. This figure thus reinforces the major difference in performance between sCCA and PCA/ $k$ -means clustering. When seeking to identify subcoalitions of individuals with consistent preferences, sCCA consistently outperforms PCA/ $k$ -means clustering. In most cases, sCCA does at least as good of a job of identifying the latent structure in the data and, when it does not, it is because PCA/ $k$ -means clustering fails to identify consistent basic units.



**Figure 8. sCCA vs. PCA/*k*-means Clustering: Probability of Intransitive Cycle**

## 5. Discussion and Conclusion

Frustration among organizational researchers on the lack of quantitative research on the role of politics and conflict in organizational decision making is as old as the original theory of the “business firm as a political coalition” (March 1962). We argue that progress analyzing organizations from a political perspective has been hampered by a lack of empirical tools to a greater extent than by a lack of theoretical frameworks. Specifically, the tools that we have for identifying groups in conflict in

organizations are either too unstructured, i.e., PCA/ $k$ -means clustering, or too restrictive, i.e., NOMINATE, to be appropriately applied to many organizational contexts. We propose a new method called sCCA that identifies subcoalitions in organizations with consistent preferences over outcomes. sCCA thus permits the type of political analysis proposed in March (1962).

We apply sCCA to two empirical contexts. Wikipedia sought to learn about its editors in order to implement organizational changes to stem the decline in editor participation. The topline results of the survey indicated that editors felt the biggest problem was other editors, which supported the view that Wikipedia's culture required reform. However, sCCA shows that Wikipedia's editors were composed of two subcoalitions, one with experienced, active users and one with newer, less frequent users. The former group identified editorial behavior as the biggest problem with Wikipedia, but the latter instead pointed to the technical complexity of the editing platform. This divergence of opinion was also reflected in the recommendations about how to use Wikipedia's resources. The former group recommended spending more on technical upgrades to the platform and features for experienced users. The latter recommended spending more on features for new users and for readers. This split between two subcoalitions erupted with the introduction of the Visual Editor on English Wikipedia in 2013, when experienced editors rebelled against the new editor interface.

In the case of the BBWAA, we demonstrate that the nature of conflict over candidates for the National Baseball Hall of Fame changed in the 2013 election. Prior to 2013, conflict over candidates was idiosyncratic, but the issue of PED use split the writers into two subcoalitions. We use sCCA to demonstrate not only the presence of subcoalition-based conflict, but also that the conflict was self-reinforcing. Prior to 2013, writers were influenced by the entire BBWAA when deciding whether to vote for candidates that had been on the ballot in previous years. However, once the split occurred in 2013, the social influence of writers in the competing subcoalition weakened. The bifurcation of the writers thus undermined the bandwagon effect that the BBWAA had relied upon to get candidates above the supermajority threshold in prior years.

Finally, we show that sCCA is robust to many groups, missing data, and imperfect alignment of subcoalition and individual preferences to a greater extent than competing methods. When seeking to identify subcoalition in organizations, sCCA outperforms PCA and  $k$ -means clustering, which is a common way of identifying clusters of individuals in organizations. In addition, we believe that sCCA is more transparent and easier-to-use than PCA and similar group-detection methods that rely on matrix decomposition. sCCA output is straightforwardly derived from the output of a linear program that seeks to minimize an objective function. Researchers need no background in multivariate analysis in order to understand how sCCA works.

We think that sCCA can be a springboard to future studies analyzing the presence, emergence, and evolution of subcoalition-based conflict and decision making in organizations. In addition to identifying and measuring political conflict in organizational decision making, sCCA can be a valuable tool for organizational researchers to better understand power dynamics in organizations, as in the Wikipedia case, and social influence in organizations, as in the BBWAA case. It took nearly a half century from the formalization of the median voter theorem in Black (1948) until the introduction of NOMINATE scores in Poole and Rosenthal (1985), which have since become the workhorse for the empirical study of the U.S. legislature. We hope that sCCA, by similarly tying method to model, will become the new standard for analyzing political decision making in organizations.

## References

- Arthur, David, and Sergei Vassilvitskii. 2007. "K-Means++: The Advantages of Careful Seeding." In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1027–35. Society for Industrial and Applied Mathematics.
- Austen-Smith, David, and Jeffrey S. Banks. 1999. *Positive Political Theory I: Collective Preference*. Michigan Studies in Political Analysis. Ann Arbor: University of Michigan Press.
- Bartholdi, J., C. A. Tovey, and M. A. Trick. 1989. "Voting Schemes for Which It Can Be Difficult to Tell Who Won the Election." *Social Choice and Welfare* 6 (2): 157–65.  
<https://doi.org/10.1007/BF00303169>.
- Black, Duncan. 1948. "On the Rationale of Group Decision-Making." *Journal of Political Economy* 56 (1): 23–34.
- Bonica, Adam. 2014. "Mapping the Ideological Marketplace." *American Journal of Political Science* 58 (2): 367–86. <https://doi.org/10.1111/ajps.12062>.
- Caple, Jim. 2012. "Partisan Deadlock in the Hall of Fame." *ESPN.Com* (blog). December 26, 2012. [http://www.espn.com/mlb/story/\\_/id/8777793](http://www.espn.com/mlb/story/_/id/8777793).
- Chappell, Bill. 2014. "Sportswriter Who Let Readers Fill Out Hall Of Fame Ballot Is Banned." *NPR.Org* (blog). January 9, 2014.  
<https://www.npr.org/sections/thetwo-way/2014/01/09/261229164/sportswriter-who-let-readers-fill-out-hall-of-fame-ballot-is-banned>.
- Conitzer, Vincent, Andrew Davenport, and Jayant Kalagnanam. 2006. "Improved Bounds for Computing Kemeny Rankings." In *AAAI*, 6:620–26.
- Cyert, Richard M., and James G. March. 1963. *A Behavioral Theory of the Firm*. 2nd ed. Cambridge: Blackwell.

- Gavetti, Giovanni, Daniel Levinthal, and William Ocasio. 2007. "Perspective--Neo-Carnegie: The Carnegie School's Past, Present, and Reconstructing for the Future." *Organization Science* 18 (3): 523–36.
- Gibbons, Robert. 2003. "Team Theory, Garbage Cans, and Real Organizations: Some History and Prospects of Economic Research on Decision-Making in Organizations." *Industrial and Corporate Change* 12 (4): 753–87.
- Goldberg, Amir, and Sarah K. Stein. 2018. "Beyond Social Contagion: Associative Diffusion and the Emergence of Cultural Variation." *American Sociological Review* 83 (5): 897–932. <https://doi.org/10.1177/0003122418797576>.
- Gupte, Mangesh, Pravin Shankar, Jing Li, S. Muthukrishnan, and Liviu Iftode. 2011. "Finding Hierarchy in Directed Online Social Networks." In *Proceedings of the 20th International Conference on World Wide Web*, 557–66. Hyderabad, India: ACM.
- Halfaker, Aaron, R. Stuart Geiger, Jonathan T. Morgan, and John Riedl. 2013. "The Rise and Decline of an Open Collaboration System: How Wikipedia's Reaction to Popularity Is Causing Its Decline." *American Behavioral Scientist* 57 (5): 664–88. <https://doi.org/10.1177/0002764212469365>.
- Heath, Chip, and Nancy Staudenmayer. 2000. "Coordination Neglect: How Lay Theories of Organizing Complicate Coordination in Organizations." *Research in Organizational Behavior* 22 (January): 153–91. [https://doi.org/10.1016/S0191-3085\(00\)22005-4](https://doi.org/10.1016/S0191-3085(00)22005-4).
- Honore, Bo E., and Ekaterini Kyriazidou. 2000. "Panel Data Discrete Choice Models with Lagged Dependent Variables." *Econometrica* 68 (4): 839–74. <https://doi.org/10.1111/1468-0262.00139>.
- Jakulin, Aleks, Wray Buntine, Timothy M. La Pira, and Holly Brasher. 2009. "Analyzing the U.S. Senate in 2003: Similarities, Clusters, and Blocs." *Political Analysis* 17 (3): 291–310. <https://doi.org/10.1093/pan/mpp006>.
- Kellogg, Katherine C. 2009. "Operating Room: Relational Spaces and Microinstitutional Change in Surgery." *American Journal of Sociology* 115 (3): 657–711. <https://doi.org/10.1086/603535>.
- Keri, Jonah. 2013. "The Fallacy of the Baseball Hall of Fame." *Grantland* (blog). January 9, 2013. <http://grantland.com/the-triangle/the-fallacy-of-the-baseball-hall-of-fame/>.
- Luxburg, Ulrike von. 2010. "Clustering Stability: An Overview." *Foundations and Trends® in Machine Learning* 2 (3): 235–74. <https://doi.org/10.1561/22000000008>.
- March, James G. 1962. "The Business Firm as a Political Coalition." *Journal of Politics* 24 (4): 662–78.
- Martin, Andrew D., and Kevin M. Quinn. 2002. "Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953–1999." *Political Analysis* 10 (2): 134–53. <https://doi.org/10.1093/pan/10.2.134>.

- Morrill, Calvin. 1991. "Conflict Management and Organizational Change." *American Journal of Sociology* 97 (3): 585–621.
- Nurmi, Hannu. 2004. "A Comparison of Some Distance-Based Choice Rules in Ranking Environments." *Theory and Decision* 57 (1): 5–24. <https://doi.org/10.1007/s11238-004-3671-9>.
- Orlowski, Andrew. 2013. "Revolt Peasants Force Wikipedia to Cut 'n' paste Visual Editor into the Bin." *The Register* (blog). September 25, 2013. [https://www.theregister.co.uk/2013/09/25/wikipedia\\_peasants\\_revolt/](https://www.theregister.co.uk/2013/09/25/wikipedia_peasants_revolt/).
- Ossadnik, Wolfgang, Stefanie Schinke, and Ralf H. Kaspar. 2016. "Group Aggregation Techniques for Analytic Hierarchy Process and Analytic Network Process: A Comparative Analysis." *Group Decision and Negotiation* 25 (2): 421–57. <https://doi.org/10.1007/s10726-015-9448-4>.
- Pfeffer, Jeffrey. 1981. *Power in Organizations*. Cambridge: Ballinger Pub. Co.
- Poole, Keith T., and Howard Rosenthal. 1985. "A Spatial Model for Legislative Roll Call Analysis." *American Journal of Political Science* 29 (2): 357–84. <https://doi.org/10.2307/2111172>.
- . 1987. "Analysis of Congressional Coalition Patterns: A Unidimensional Spatial Model." *Legislative Studies Quarterly* 12 (1): 55–75. <https://doi.org/10.2307/440045>.
- . 2001. "D-Nominate after 10 Years: A Comparative Update to Congress: A Political-Economic History of Roll-Call Voting." *Legislative Studies Quarterly* 26 (1): 5–29. <https://doi.org/10.2307/440401>.
- Poole, Keith T., Howard Rosenthal, and Kenneth Koford. 1991. "On Dimensionalizing Roll Call Votes in the U.S. Congress." *The American Political Science Review* 85 (3): 955. <https://doi.org/10.2307/1963858>.
- Saari, Donald G. 2001. *Decisions and Elections: Explaining the Unexpected*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511606076>.
- Saari, Donald G., and Vincent R. Merlin. 2000. "A Geometric Examination of Kemeny's Rule." *Social Choice and Welfare* 17 (3): 403–38. <https://doi.org/10.1007/s003550050171>.
- Saari, Donald G., and Fabrice Valognes. 1999. "The Geometry of Black's Single Peakedness and Related Conditions." *Journal of Mathematical Economics* 32 (4): 429–56. [https://doi.org/10.1016/S0304-4068\(98\)00062-7](https://doi.org/10.1016/S0304-4068(98)00062-7).
- Sampson, Tim. 2013. "Will Wikipedia's Pretty New Editing Software Solve Its Recruitment Crisis?" *The Daily Dot* (blog). July 4, 2013. <https://www.dailydot.com/business/wikipedia-visual-editor-wysiwyg/>.
- Silver, Nate. 2013. "In Cooperstown, a Crowded Waiting Room." *FiveThirtyEight* (blog). January 12, 2013. <https://fivethirtyeight.blogs.nytimes.com/2013/01/12/in-cooperstown-a-crowded-waiting-room/>.

- Simonite, Tom. 2013. "The Decline of Wikipedia." *MIT Technology Review*, October 22, 2013.  
<https://www.technologyreview.com/s/520446/the-decline-of-wikipedia/>.
- Suh, Bongwon, Gregorio Convertino, Ed H. Chi, and Peter Pirolli. 2009. "The Singularity Is Not near: Slowing Growth of Wikipedia." In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration - WikiSym '09*, 1. Orlando, Florida: ACM Press.  
<https://doi.org/10.1145/1641309.1641322>.
- Tibshirani, Robert, Guenther Walther, and Trevor Hastie. 2001. "Estimating the Number of Clusters in a Data Set via the Gap Statistic." *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 63 (2): 411–23.
- Van Gunten, Tod, John Martin, and Misha Teplitskiy. 2016. "Consensus, Polarization, and Alignment in the Economics Profession." *Sociological Science* 3: 1028–52. <https://doi.org/10.15195/v3.a45>.
- Wales, Jimmy. 2009. "State of the Wiki." Wikimania, August 27.  
[https://commons.wikimedia.org/wiki/File:State\\_of\\_the\\_Wiki.pdf](https://commons.wikimedia.org/wiki/File:State_of_the_Wiki.pdf).
- Zbaracki, Mark J., and Mark Bergen. 2010. "When Truces Collapse: A Longitudinal Study of Price-Adjustment Routines." *Organization Science* 21 (5): 955–72.