



Same people, different group: Social structures are a central component of group concepts

Alexander Noyes^{a,*}, Frank C. Keil^b, Yarrow Dunham^b, Katherine Ritchie^c

^a New York University, New York, NY, United States

^b Yale University, New Haven, CT, United States

^c University of California – Irvine, Irvine, CA, United States

ARTICLE INFO

Keywords:

Social cognition
Social groups
Concepts
Numerical identity
Individual concepts

ABSTRACT

We examine whether people conceptualize organized groups as having at least two parts: In addition to members (e.g., Alice), they also have social structures (i.e., roles and relations). If groups have members and social structures, then numerically distinct groups can have the same members if they differ in their structures. In Studies 1–4, participants numerically distinguished groups that had the same members when they had different structures. Participants numerically distinguished even when groups had the same function—the same people playing chess together Monday and Tuesday can be numerically distinct groups. In Study 4, we compare clubs to tables, and find that participants numerically distinguish tables by their structures too (i.e., the configuration of their parts) even when they have the same parts (which can be disassembled and then reassembled with ease). In Study 5, we find that participants rate groups as existing in space and time like concrete objects, suggesting that participants represent groups as at least partially concrete, such that groups have at least two parts (their structures and their members). Finally, in Study 6, we show that people will judge the same person as exemplary with respect to one group but condemnable with respect to another—even when those groups have the same members.

1. Introduction

Human societies are filled with organized groups, such as clubs, government agencies, tribes, and academic departments. Organized groups (also known as dynamic, task, or institutional groups) are structured wholes. For example, a basketball team—like a dog or a chair—has connected parts that work together to achieve the properties of the whole. A dog can run, a chair can support weight, and a basketball team can win games. In the case of a basketball team, there are two kinds of parts. There are the members. There are also the roles and the relations between them: For example, leadership (e.g., manager, head coach, assistant coach, players) and player roles (e.g., point guard, forward). We ask whether structure is a central component of group concepts. We ask this within the context of numerical identity judgments. If organized groups are structured wholes, then a group is not identical to its members. If people can occupy roles within two structures, then two groups can overlap in members if they possess different structures.

1.1. Theoretical rationale

We ask whether people conceptualize groups as structured wholes on the premise that the best empirical and philosophical description of groups will incorporate social structure as a central component. The idea that groups are structured wholes—or social systems—is advanced by theorizing in social psychology (e.g., Baumeister, Ainsworth, and Vohs, 2016; Steiner, 1974), social ontology (e.g., Haslanger, 2000, 2016; Ritchie, 2020), and sociology (e.g., Porpora, 2013). The common thread is that structural explanation is needed to understand groups and the people in them (Haslanger, 2016; Steiner, 1974). Thus, despite variation in explanatory goals and methods, disciplines that aim to understand groups often converge on the centrality of social structure.

We can see that groups are not their members by considering a hypothetical. Amy, Brad, and Lisa are the current members of a chess team Rookies. Amy is the current captain. *Having the role of captain* cannot be reduced to facts about her because she might not have been elected; therefore, the role of captain cannot be reduced to facts about Amy either. Amy, Brad, and Lisa are members because the former captain,

* Corresponding author at: New York University, New York, NY, United States.

E-mail address: alexander.noyes@nyu.edu (A. Noyes).

<https://doi.org/10.1016/j.cognition.2023.105567>

Received 20 October 2022; Received in revised form 14 July 2023; Accepted 15 July 2023

Available online 4 August 2023

0010-0277/© 2023 Elsevier B.V. All rights reserved.

Megan, accepted their applications. The property *being the current members of Rookies* cannot be reduced to facts about Amy, Brad, and Lisa because Megan might have decided otherwise. Therefore, Rookies cannot be reduced to facts about Amy, Brad, and Lisa either. These patterns of decisions are constrained by norms. The norm that there is a captain, and the norm that captains select members and members select captains. These norms constitute the group's structure (Gilbert, 1992, 2013; Thomasson, 2016; Ritchie, 2020), which minimally encompasses the roles (e.g., captain, member) and the relations between them. Structure is also tightly linked with proper function. When people perform their roles, they realize the proper function of the group (e.g., to play chess) and the proper function of the group constrains the norms that constitute the roles.

The empirical project of characterizing group concepts is distinct from the project of characterizing groups. That is, there are open questions about group concepts that are not answered by prior work on groups. A distinction between groups and group concepts is blurrier than distinctions in other domains. It is easy to distinguish biology from the cognitive and developmental psychology of biological concepts; indeed, people's concepts are often error-ridden, such as the pervasive failure to understand how natural selection works (Coley, Arenson, Coley, and Tanner, 2017; Gelman & Rhodes, 2012; Shtulman and Schulz, 2008). By their nature, groups are interrelated with our concepts. If groups depend on people's minds, then how can the reality of groups deviate from the concepts people hold when making those decisions? Although the distinction is blurrier, group concepts can deviate from empirical and philosophical descriptions. Consider a well-documented deviation: The categories *man* and *woman* depend on facts about our societies (Haslinger, 2000), and yet people often take *man* and *woman* to depend on facts about human biology (Gelman and Taylor, 2000). Indeed, the error in people's gender concepts may play a causal role in sustaining the social reality of gender. Groups depend on people's intentional behavior but they may not require veridical concepts. Instead, groups and group concepts causally interact through "looping effects," in which groups affect group concepts and group concepts affect groups (Hacking, 1995). Therefore, the empirical project of characterizing groups and group concepts are mutually informative but distinct.

1.2. Group concepts

There is also a literature on group concepts. The best studied topic is social categories (e.g., Haslam, Rothschild, & Ernst, 2000). There is also work on organized groups too, though. One program of research examines entitativity (Hamilton and Sherman, 1996; Hamilton, Sherman, and Lickel, 1998), which is the extent to which a group is perceived as an entity (e.g., the Senate) rather than a mere collection (e.g., people at a bus stop). Lickel et al. (2000) found that people distinguish intimacy groups, organized groups, social categories, and loose associations, which they uncovered both by having people sort groups into categories, and by looking at clusters in feature ratings (e.g., size, stability). The distinction between organized groups and social categories aligns with theorizing about groups (e.g., Rabbie and Horwitz, 1988; Ritchie, 2013; Ritchie, 2015). People rated organized groups higher in entitativity than social categories (Lickel et al., 2000), which is consistent with organized groups being structured wholes instead of collections of people.

The idea that groups have interdependent members appears present early in development too. For example, children expect members of groups to cooperate (Misch, Paulus, & Dunham, 2021). Children even moralize interdependence, believing it is immoral for members of groups to harm each other (Chalik and Rhodes, 2020; Rhodes and Chalik, 2013). Indeed, children believe that the interdependence between members partly constitutes the group. Young children believe that a person is a member of a group when she and others recognize her membership (Noyes and Dunham, 2017). Young children see this interdependence as interpersonal (i.e., they are friends with each other), whereas children age 6 and older see this interdependence as held

together by rules, suggesting they conceptualize groups as social structures (Noyes & Dunham, 2020).

There is also a large body of work investigating how people represent coalitions (e.g., Kurzban, Tooby, and Cosmides, 2001). This work suggests that people identify coalitions by attending to patterns of alliances. A recent account develops this view in detail. Pietraszewski (2022) argues that observers represent groups—in the context of conflict—as roles in four triadic relations. In a relation called "generalization," if A attacks B, then A attacks C. In a relation called "alliance," if A attacks B, then C attacks B. In a relation called "displacement," if A attacks B, then B attacks C. In a relation called "defense," if A attacks B, then C attacks A. Pietraszewski (2022) refers to these as group-constitutive roles. If two people play the role B and C in generalization, A and C in alliance, A and C in displacement, and B and C in defense, then they are in a group with respect to this conflict. Although we and Pietraszewski connect group concepts to roles in a constitutive way, our accounts have different theoretical aims; indeed, our accounts are orthogonal to each other. Pietraszewski (2022) does not characterize how people represent a group *per se*; instead, he aims to characterize how people represent the structure of individual behavior during conflict when people belong to groups. To see this, consider how his account can (intentionally) encompass groups of different kinds: People should expect conflict to conform to the four triadic relations when a family comes into conflict with a multinational company, or when plane passengers come into conflict with a terrorist group. Yet, people surely differ in their concepts of families, multinational companies, plane passengers, and terrorist groups. Therefore, people can vary in how they conceptualize the nature of a group but still expect the group to conform to the four triadic relations in conflict. Our account, on the other hand, is specific to organized groups and neutral with respect to conflict. Organized groups are structured wholes in and out of conflict, which they may never participate in. Social structures are primarily a means of organizing intragroup cooperation.

1.3. Numerical identity

We focus on how people judge the synchronic identity of organized groups. Synchronic identity refers to the identity that exists between an object and itself at a moment in time. Diachronic identity refers to the identity that exists between an object and itself over time. To use an analogy: The question we seek to answer about groups is less akin to asking if a clay sculpture survives losing its arm, and more akin to asking whether a clay sculpture has both clay and form as parts. One possibility is that people take group A to be identical to group B at a time if there is a one-to-one correspondence between their members. We argue against this. We argue that structure is an important part of identity judgments too. Groups A and B are numerically distinct if they have different structures. We do not argue that a group's structure is static. To continue the analogy, even if people represent a clay sculpture as identical to its whole form—arms, legs, head, torso—they might still perceive the sculpture as surviving the loss of its arms.

1.4. Empirical approach

In Studies 1–4, we investigated numerical identity judgments about groups at a time. We relied on diachronic scenarios to aid comprehension—for example, we told participants about two original groups that combined their membership roster. However, we asked people only to consider how many groups there were after the series of events; we did not ask people to consider whether the original group still existed.

In Study 1, we examined scenarios in which a pair of academic departments, government agencies, tribes, or clubs ended up with complete overlap in membership. In Studies 2 and 3, we compared a case where two groups merged their structures to a case where two groups merged only their membership roster. In Study 4, we compared a case where three people filled roles in two social structures to a case where

three pieces (of wood or plastic) filled roles in two physical structures. In Studies 1–4, we expected participants to indicate that the same people can be two groups at the same time, showing that groups are not identical to their members. We expected participants to say two groups become a single group when their structures merged (Study 2 and 3). We expected participants to also say that three pieces (of wood or plastic) can be two tables, consistent with the idea that clubs and tables are both structured wholes (i.e., entities) rather than mere collections of parts.

In Studies 5 and 6, we investigated two additional features of group concepts. Our account is that structure is one component of groups and that people are another component. A social structure, like the design in a blueprint, is an abstract entity. In the same way wood and nails can realize a design, people and their behaviors can realize a social structure. Therefore, our account predicts that people will conceptualize groups as at least partially concrete. Study 5 investigates this by asking people to judge whether groups have physical extent and spatiotemporal location and comparing those ratings to concrete and abstract entities. We also propose that structure is normative. Therefore, social structure should be sufficient to produce evaluations (positive or negative) of a person, such that their behaviors can be exemplary or condemnable in group-specific ways. Study 6 investigates this by seeing whether people are willing to reward a person in the context of one group but punish them in the context of another group. Together, these two studies provide more detail into our account.

2. Study 1

We describe scenarios in which two pre-existing groups end up with the same members. We asked participants whether there were two groups or one. We predicted that participants would indicate the presence of two groups even when they completely overlap in membership.

2.1. Method

2.1.1. Participants

We recruited 100 Amazon participants. We would have 93% power to detect a difference from chance if participants reason there are two groups at least 67% of the time, which we considered the minimum effect consistent with our account. After exclusions, we had 97 participants. We did not collect demographic information for this or any of the studies. Yale University's Institutional Review Board.

2.1.2. Stimuli and measures

There were four total vignettes. Each vignette corresponded to a type of group: Interest groups (i.e., a chess team and flower-arranging club), government agencies, tribal nations, and academic departments. In each vignette, two groups that initially have different members end up having the same members, brought about by circumstantial processes relevant to the type of group.

Club. Ben, Jonathan, and Margaret are the current members of an amateur chess-playing club called Rookies. They elect Margaret as President and Ben as vice-President. Danielle and Natalie are the only current members of an evening flower-arranging club called Late Bloomers. They decide Danielle will be the leader. One day, Margaret and Danielle meet each other and learn of the other's club. Danielle asks Ben, Jonathan, and Margaret to join the flower-arranging club, and Margaret asks Danielle and Natalie to join her chess-playing club. They all agree. Therefore, the chess-playing club and the flower-arranging club currently have the same exact members.

Tribe. Aeduia and Menape are two tribes. Each is governed by a council of elders and they each have distinct traditions and customs. Aeduia initiates members through a ritual cleansing; Menape welcomes members when a family 'adopts' them. One custom they share with several other regional tribes is the practice of exogamy:

Members can only marry members of other tribes. Aeduia and Menape recognize children of those marriages as having dual tribal membership. Because of food scarcity, other regional tribes migrate eastward leaving Aeduia and Menape to practice exogamy only with each other. After a particularly bad famine, many members die. The only members left are all children of Aeduia-Menape marriages and each has dual tribal membership. Therefore, Aeduia and Menape currently have the exact same members.

Government agency. The Massachusetts Fishing Agency (MASSFA) regulates fishing licenses, sets guidelines for catch and release, and decides policy for commercial fishing companies. The Massachusetts Bay & River Agency (MBRA) establishes guidelines for public and private water usage in Massachusetts. After severe budget deficits, State legislators initiate many layoffs and reductions of worker time from full- to part-time. They also decide to appoint the same members to MASSFA and MBRA, to reduce the total budget for salaries by half. Therefore, MASSFA and MBRA currently have the same exact members.

Academic Department. At Prarieumont University, the Art History Department specializes in French and Francophone art and is rated 8th in the county. Many professors in the Art History department are cross-appointed in the French Department and teach courses cross-listed for the French major. After students complain about the lack of diversity of faculty and staff, the departments coordinate a diversity cluster hire, hiring five new professors that are cross-appointed in both departments. That same year, several older professors retire, and two other professors move to new universities. As a consequence, all current members are dual-appointed. Therefore, the Art History Department and the French Department currently have the same exact faculty and staff.

Test question. Participants were then asked: *How many Xs [clubs, agencies, tribes, departments] are there?* There were two options: one or two.

2.1.3. Design and procedure

Participants were randomly assigned to read and respond to two of the four possible vignettes in randomized order.

2.2. Results

We examined whether responses were significantly higher than chance in a logistic generalized estimation equation (GEE). Participants were significantly above chance, $b = 2.55$, robust standard error = 0.27, $p < .001$. Participants indicated there were two groups 94% of the time for the government agency, 92% of the time for the academic department, 100% of the time for the interest groups, and 85% of the time for the tribe.

2.3. Discussion

Participants indicated that two groups of the same type could have complete overlap in membership. These results indicate that participants conceptualize groups as distinct from a collection of particular people. The major limitation is that we did not examine a contrasting case; in Study 2 we contrast a club with a table to show that when participants indicate two groups, it is not merely a response bias or task demand.

3. Study 2

We contrast a case where two groups merge their structures and members with a case where groups merge only their members. All else being equal, structure should modify participants' numerical judgments, so we should observe a significant difference between these conditions, and participants should indicate two groups when there are two

structures.

3.1. Method

3.1.1. Participants

We recruited 200 participants from Amazon Mechanical Turk. After exclusions, we had 191 participants.

3.1.2. Design and procedure

Participants read about two groups: a chess club and a flower-arranging club. The descriptions were accompanied by diagrams of the relations between the roles (Fig. 1). The description of the chess club was: “There is an amateur chess club called ‘Rookies.’ Members meet Mondays to play chess. Below shows how the group is organized: There is a chair, treasurer, and the other members at large. One of the rules they have is that members elect the chair and treasurer.” The description of the flower-arranging club was: “There is a flower arranging club that meets Tuesday evenings called Late Bloomers. Below shows how the group is organized. There is a President, Vice President, and General members. One of the rules they have is that the general members vote on new meeting times.”

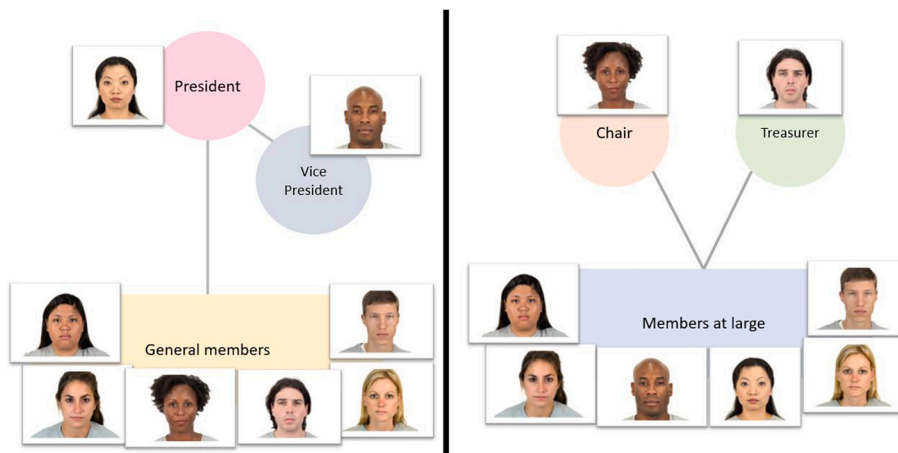
Participants then learned about the current members and then saw a simplified diagram of the group with a picture of the member next to their role. Participants were then randomly assigned to the Merge condition or the non-Merge condition. In the Merge condition, participants

read a story in which the groups merge their social structures; in the non-Merge condition, their social structure remain intact.

Non-merge: One day Angela, the chair of Rookies, meets Vivian, the president of Late Bloomers. Angela is excited because she also loves arranging flowers, and Vivian is excited because she also loves playing chess! Angela and Vivian discuss how fun it would be for everyone to play chess together on Mondays and arrange flowers together on Tuesday. Vivian tells the others about playing chess on Mondays; they are excited and want to play chess too – Angela admits them as members-at-large. Angela tells the others about arranging flowers on Tuesdays; they are excited and want to arrange flowers too – Vivian admits them as general members. Below is a diagram of how the members fit into the roles now. So now, the same exact people play chess on Mondays and arrange flowers on Tuesdays.

Merge: One day Angela, the chair of Rookies, meets Vivian, the president of Late Bloomers. Angela is excited because she also loves arranging flowers, and Vivian is excited because she also loves playing chess! Angela and Vivian discuss how fun it would be for everyone to play chess together on Mondays and arrange flowers together on Tuesday. Everyone discusses; they decide that Angela will enter the role of President and Vivian will enter the role of Co-President. They decide it’ll be helpful to have a Moderator, since

Panel A: Non-merge condition.



Panel B: Merge condition:

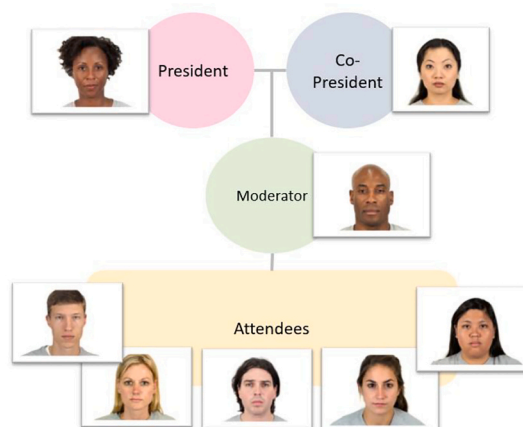


Fig. 1. Participants saw Panel A in the non-merge condition. Participants saw panel B in the merge condition.

meetings will be larger now: Brad agrees to be the moderator. Everyone else will take on the role of attendees: Attendees play chess on Monday and arrange flowers on Tuesdays. Below is a diagram of how the members fit into the roles now.

Finally, participants judged whether there were two groups now or still one, using the same measure as in Study 2.

3.2. Results

Pre-registered analysis. We used the Agresti–Coull method to compute confidence intervals for the two conditions. Participants indicated the existence of two groups when the structures merged only 16% of the time, 95%CI: [0.10: 0.25]. Participants indicated the existence of two groups when the structures did not merge 89% of the time, 95%CI: [0.80, 0.93]. Neither confidence interval includes zero and they do not overlap (Fig. 2).

Additional analysis. We also analyzed the results with a logistic generalized linear model (GLM). There was a significant effect of condition, $b = 3.72$, $SE = 0.43$, $p < .001$. The structure merge condition was significantly below chance, $b = -1.67$, $SE = 0.28$, $p < .001$, and the non-merge was significantly above chance, $b = 2.04$, $SE = 0.32$, $p < .001$.

3.3. Discussion

Participants relied on structure to make identity judgments—they indicated one group when there was one structure and two groups when there were two structures. One worry is that the conditions also varied in the number of photographs. This is the most natural way to depict the structures but it also introduces confounding pragmatic cues, and it is possible to avoid it by merely showing the same diagram twice. A second worry is that the groups also differed in function, which could contribute to why participants indicated two groups.

4. Study 3

We replicate Study 2 by presenting two chess clubs coming to have the same structure and members or only the same members. We use two photographs in both conditions.

4.1. Method

4.1.1. Participants

We recruited 100 participants from Amazon Mechanical Turk. 101

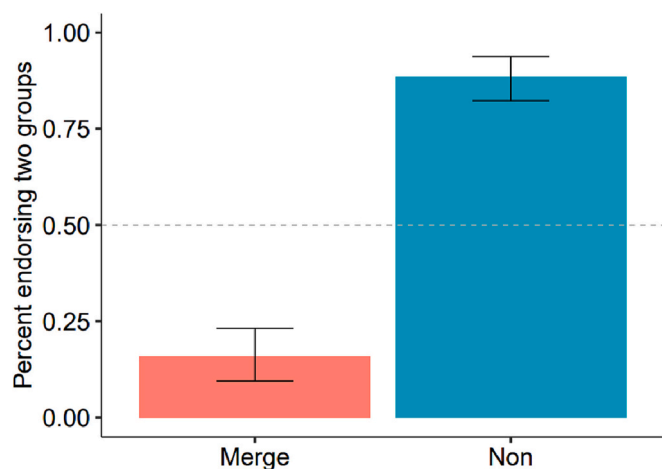


Fig. 2. Study 2. Comparisons of how often participants said there were still two groups as a function of whether the structures merged (merge) or merely the membership rosters (non-merge). Error bars are 95% confidence intervals. The dotted line is the midpoint of the scale.

participants submitted a finished survey.

4.1.2. Design and procedure

There were two major changes from Study 2: (1) All references to the flower club were replaced by references to a second chess club. (2) The merged structure was presented twice, so participants also saw two photographs. We also added and revised text for clarity.

4.2. Results

We followed the pre-registered analyses of Study 2. We used the Agresti–Coull method to compute confidence intervals for the two conditions. Participants indicated the existence of two groups when the structures merged only 27% of the time, 95%CI: [0.17: 0.40]. Participants indicated the existence of two groups when the structures did not merge 88% of the time, 95%CI: [0.75, 0.95]. Neither confidence interval includes zero and they do not overlap (Fig. 3).

Additional analysis. We also analyzed the results with a logistic generalized linear model (GLM). There was a significant effect of condition, $b = 2.97$, $SE = 0.54$, $p < .001$. The structure merge condition was significantly below chance, $b = -1.00$, $SE = 0.31$, $p = .001$, and the non-merge was significantly above chance, $b = 1.97$, $SE = 0.44$, $p < .001$.

4.3. Discussion

We replicated Study 2. We can see that, all else being equal, structure is sufficient for participants to numerically distinguish groups. Participants readily endorsed the existence of two groups even when they had the same members engaging in the same activity. As before, we find that participants indicate one group when the groups had the same structure.

5. Study 4

Although we have focused on groups, there are also structural accounts of artifacts too (Fine, 1999; Koslicki, 2008). In both cases, there are physical parts (e.g., people, pieces of wood) that occupy functional roles (e.g., president, table top). Unlike groups, it is rare for the physical parts of an artifact to be regularly disassembled and reassembled into another functional structure. Therefore, people may be less likely to represent a collection of parts as two tables “at the same time,” because there is never a duration of time in which the parts are switching between structures the way people effortlessly move in and out of roles. However, if we recreate this situation with tables, then participants




Fig. 3. Study 3. Comparisons of how often participants said there were still two groups as a function of whether the structures merged (merge) or merely the membership rosters (non-merge). Error bars are 95% confidence intervals. The dotted line is the midpoint of the scale.

should draw the same conclusion. If week after week, three pieces of wood and plastic are assembled into one structure on Monday and one structure on Tuesday, then participants should indicate the existence of two tables.

5.1. Method

5.1.1. Participants

We recruited 100 participants from Amazon Mechanical Turk. 101 participants submitted a finished survey. This design was within-subject, so participants saw both conditions.

5.1.2. Design and procedure

Participants read about three people who have roles in two social structures or three pieces of wood and plastic that have roles in two physical structures. In each case, they made a numerical judgment—they decided whether there was one or two. Because the parts of the table are only ever one table at a time, numerical judgments might deviate from identity judgments. Therefore, we included a second question about identity that did not rely on counting how many tables or clubs there were. We assigned a name to the club or table as depicted in one picture, then asked participants if the second picture depicted the same club or table. Because this question deviated from the design of Studies 2–3, we always included it second after the numerical judgment. Thus, this question was included merely to prevent false negatives. Here is the script for the club and table conditions:

Club. Here is what you will consider: Can two clubs have the same people as members? These people go to Room 101 on Monday and Tuesday. [Picture of people]. On Monday, they are parts of the structure depicted below. Megan is President, Brad is Vice President, and Liam is an Apprentice. [Picture of structure one]. On Tuesday, they are parts of the structure depicted below. Megan is a Recruit,

Brad is Treasurer, and Liam is Chair. [Picture of structure two] Do the same people play a role in both structures? [Picture of both structures] [Participants input yes or no]. Yes. These people here. [Picture of people] Do these three people fill the same roles on Monday and Tuesday? [Picture of people] No. The people fill very different roles on Monday and Tuesday. [Picture of both structures] [Participants input yes or no]. To recap: There is one set of people and two structures. How many clubs are there? Let's name the club on the left Dax. [Picture of structure one] Is the thing on the right Dax? [Picture of structure two].

Table. Here is what you will consider: Can two tables have the same pieces as parts? These pieces go to Room 101 on Monday and Tuesday [Picture of parts]. On Monday, they are parts of the structure depicted below. The acrylic circle is the tabletop. The flat maple board is the left leg. The red oak dowel is the right leg. [Picture of structure one]. On Tuesday, they are parts of the structure depicted below. The acrylic circle is the foot. The flat maple board is the tabletop. The red oak dowel is the stand. [Picture of structure two]. Do the same pieces play a role in both structures? [Picture of both structures] [Participants input yes or no]. Yes. These pieces here. [Picture of parts] Do these three pieces fill the same roles on Monday and Tuesday? [Picture of parts] [Participant input yes or no]. No. The pieces fill very different roles on Monday and Tuesday. [Picture of both structures] To recap: There is one set of pieces and two structures. How many tables are there? Let's name the table on the left Dax. [Picture of structure one] Is the thing on the right Dax? [Picture of structure two].

To avoid the pragmatic cues of showing one or two photographs, we did not show photographs on the page with the test question, and we rotated between showing one photograph (of the parts disassembled) or two photographs (of the structures) during vignette and comprehension check questions. See photographs in Fig. 4.





Fig. 4. Photographs of a collection of three parts and the two structures they can assemble into. The diagrams for the groups were the same as the non-merge conditions in Studies 2–3, except only three people were depicted.

5.2. Results

We used a multi-level logistic model to see whether there were any effects of condition (club versus table) or question (count versus name). In a model with condition and participant as random effect, there was no significant difference, $b = 0.32$, $SE = 0.28$, $p = .267$ (Fig. 5). In a model with condition and question, there was no significant interaction, $b = -0.98$, $SE = 0.58$, $p = .090$. Participants always indicated two clubs or tables, including the direct numerical judgment and the direct name-based question, $b = 1.85$, $SE = 0.28$, $p < .001$.

5.3. Discussion

Study 4 demonstrates that people incorporate structure into their concepts of clubs and tables. The cases we presented had relatively simple structures. Therefore, we expect the results to generalize broadly, as many other groups and artifacts have far more intricate structures. This raises a possibility of a generalized structural account, in which people represent complex individuals—like organized groups, physical artifacts, and animals—as structured wholes. Specifically, if an individual has functional parts—e.g., coach, top, head—that stand in relation to each other, and are physically instantiated—e.g., by people, wood, flesh, then people may understand the structure as partially constituting the individual and thus distinguishing it from other individuals.

Although we find similar reasoning about clubs and tables, there are likely important, qualitative differences. Study 4 casts doubt on an account in which people exclusively represent groups as having structures as parts; it does not erase the differences between physical structures and social structures and between collections of people and combinations of physical components. Physical structures require parts to occupy spatiotemporal locations relative to other part—for example, in Study 4, a board had to either stand perpendicular to the floor or parallel to the floor. It is impossible for a physical part to have more than one orientation and location at the same time, so parts cannot combine into two physical structures at the same time. On the other hand, we tend to think of people as retaining their obligations and privileges even when they are not performed or exercised, so a person can occupy two roles at the same time. If participants share this intuition, then they would affirm the possibility of coincident clubs (i.e., club with all the same members) coexisting even for arbitrarily small intervals of time, whereas they would deny this possibility for coincident tables.

One may also wonder whether groups are less concrete than physical

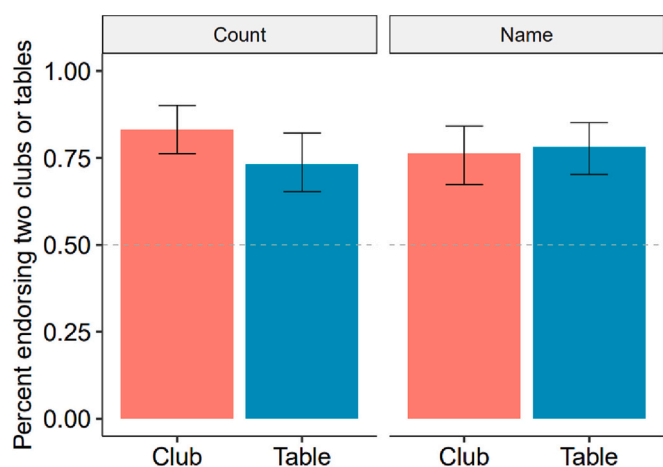


Fig. 5. Study 4. Comparisons of how often participants said there were two clubs or two tables when the same people or parts occupied roles in two structures. The count question is a direction numerical judgment as in Studies 2–3, whereas the name assesses identity by asking participants if the name of the entity in one picture is the name of the entity in the other.

artifacts. For example, participants might represent groups as having one part (their social structures) and physical artifacts as having two parts (their physical structure and the parts instantiating it). If this were true, then participants would represent groups as wholly abstract without any concrete parts. In Study 5, we test whether participants represent groups as wholly abstract or as (at least partially) concrete to determine whether groups are identical to their social structures or whether they have concrete parts too (i.e., their members).

6. Study 5

Are groups concrete, like tables and rocks, or abstract, like numbers and mathematical sets? If groups were identical to their social structures—in the sense of systems of roles—they would be abstract. For example, people might think about chess this way: People can play chess and own chess boards but it does not make sense to ask about the current location of chess. We propose that group concepts have two parts, though: A social structure and a collection of people who occupy that social structure. In this way, organized groups are like other complex wholes, including physical artifacts and animals: They have physically-instantiated functional parts that relate to each other in a structured way. To measure concreteness, we asked participants to evaluate the appropriateness of questions about location and size. We compared organized groups to concrete entities (e.g., dogs, bales of hay) and abstract entities (e.g., justice, the number zero).

6.1. Method

6.1.1. Participants

We anticipated large effect sizes for the comparisons to the midpoint. We estimated that we could achieve 95% power with 22/23 participants to detect large effects ($d = 0.8$); we doubled the sample size to ensure this number in every cell and rounded to 50 to account for noise in recruitment on mTurk. Fifty participants were included in analyses.

6.1.2. Stimuli and measures

There were three domains of entities: There were group items, which included “a soccer team”, “a reading club”, “a non-profit company”, “an academic department”, “a government agency”, and “a tribe”. There were the concrete control items: “a chair”, “a telescope”, “a hammer”, “a dog”, “a salamander”, “an elephant”, “a bale of hay”, “a bowl of soup”, “my luggage”. Finally, there were the abstract control items: “the number five”, “the number forty-seven”, “the number zero”, “justice”, “freedom”, “truth”, “an opinion”, “a delay”, “an ability”.

Participants evaluated the appropriateness of four questions: “where is X located, currently?”, “when is the last time you were in the same building as X?”, “will X fit in this conference room?”, and “how big is X?”. Participants rated these questions on a 1–6 Likert scale; specifically, they were asked “Does this question apply to this: X?” and then selected “not at all”, “mostly not”, “slightly not”, “slightly”, “mostly”, “completely”.

6.1.3. Design and procedure

Participants first read a simple description of the task: “In this study, you will decide if a question is applicable or not. For example, ‘how accurate is it?’ applies to a theory but not to an apple.” Then, participants responded to the four questions in randomized order (current location, past location, fit, size) for 12 of the 24 total stimuli.

6.2. Results

We used multi-level models and included participant, item, and question as random intercepts. Participants rated both groups, $b = 2.70$, $SE = 0.21$, $p < .001$, and concrete control items, $b = 3.08$, $SE = 0.21$, $p < .001$, as more concrete than abstract control items (Fig. 6). There was no significant difference between groups and the concrete control items, $b = -0.38$, $SE = 0.21$, $p = .081$.

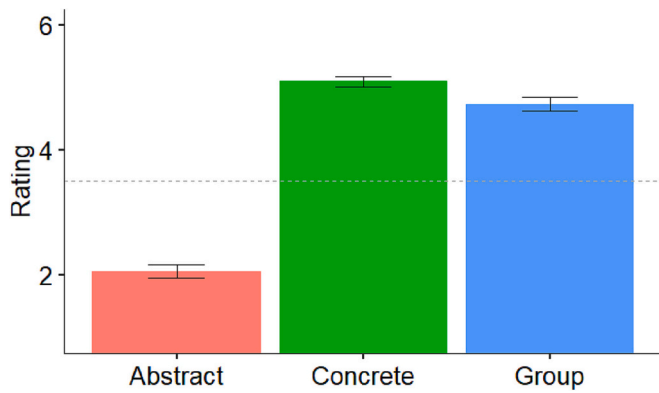


Fig. 6. Study 3. Ratings of whether spatiotemporal location and extent questions apply to abstract entities, concrete entities, or groups. Error bars are 95% confidence intervals. The dotted line is the midpoint of the scale.

There was variation among the items (Table 1). For example, “how big is it” sounded natural for numbers (though less natural than groups and concrete objects) and even somewhat natural for other abstract objects like opinions. Participants rated questions about fitting as less natural for all items, and sounded least natural for complex groups compared to simple groups. This provides some evidence that there are senses in which groups are abstract, especially complex groups. However, for every question, groups were always higher than all the abstract items—at minimum, approximately 1 scale point higher than “how big is it” about numbers. The most diagnostic question—defined as the largest difference score between concrete items and abstract items—was the question about spatiotemporal location. On this question, there was no difference between simple and complex groups, and groups were far closer to concrete objects than abstract objects.

6.3. Discussion

Participants rated groups as more concrete than abstract items and above the midpoint of the scale. There was variation by item and question (Table 1). However, the most diagnostic question—spatiotemporal location—had groups as concrete rather than abstract. Insofar as groups are in space and time with nonzero size and duration, then groups are not wholly abstract. It is unlikely that people represent groups as objects in the colloquial sense, especially in the sense implicitly represented in visual working memory. Groups lack a single physical boundary, and so can be on both sides of a solid wall, which violates the coherence principle that drives object perception. Because they lack coherence, there are reasons to question their concreteness relative to less controversial cases, such as a sphere of pure

Table 1
Average ratings by question and item.

	Where is it?	Could it fit?	How big is it?	Average
<i>Concrete</i>				
Artifacts	5.59	5.08	5.68	5.45
Animals	5.38	4.47	5.76	5.21
Other Concrete	5.34	4.66	5.62	5.21
Average	5.44	4.74	5.69	5.29
<i>Abstract</i>				
Numbers	2.38	1.81	4.36	2.85
Ideals	1.86	1.32	1.90	1.69
Other Abstract	1.78	1.37	3.13	2.09
Average	2.00	1.50	3.13	2.21
Simple Group	5.16	4.72	5.32	5.07
Complex Group	5.23	3.31	5.31	4.62

Note. Simple group = team, club, tribe; complex group = agency, department, non-profit.

iron. Thus, we cannot claim that the pattern of results generalizes to every dimension of concreteness. We conclude that participants distinguish organized groups from wholly abstract entities because concrete objects—people—are one part of groups; whether their concrete parts assemble into a concrete whole is a trickier issue that our results cannot answer. Nevertheless, Studies 4 and 5 suggest people represent groups as having at least two parts (their social structure and their members), which excludes an account in which people represent groups as identical to social structures.

7. Study 6

Social structures are normative. Because roles assign obligations to people, social structures should be sufficient to support evaluations. All else being equal, a person can be exemplary with respect to one structure but condemnable with respect to another. For example, an ambitious workaholic might be exemplary as a professor but lackluster as a journal editor, committee member, or spouse. If groups have the same members, the same people might even reach different conclusions about a person. In Study 6, we examined a scenario like this, presenting two groups with the same members—indeed, even the same leader—and a group member that was exemplary with respect to one group but not the other. We asked whether the leader of both groups should reward or punish the target people when acting as the leader of each group. We predicted that participants would make group-specific judgments even though the evaluator and target were the same for each group.

7.1. Method

7.1.1. Participants

We recruited 100 participants from Amazon Mechanical Turk. A higher-than-expected number of returned surveys appeared to be automated respondents based on nonsensical write-in responses to the exit survey (We asked participants to briefly summarize the task; a nonsensical write-in might include “I saw my friend yesterday and I eat salad with my friend”). Forty-seven participants were retained; fortunately, this still gave us over 95% power to detect the pre-registered, predicted effect size.

7.1.2. Stimuli and materials

The vignettes built on those used in Study 1. They included a pair of interest clubs, government agencies, tribal nations, and an academic department. In the vignette, participants learned that the two groups had the same members and the same leader. Participants then read about a character who upheld (and went above and beyond) the norms for group A but violated the norms of group B. Here is an example:

There is a flower arranging club that meets Tuesday evenings called ‘Late Bloomers’. There is an amateur chess club called ‘Rookies.’ Members meet Mondays to play chess. Amy, Mark, Lisa, and Brad are the current members of the flower-arranging club Late Bloomers. Lisa is the President of Late Bloomers. Amy, Mark, Lisa, and Brad also happen to be the current members of the chess club Rookies. Lisa is the Captain of Rookies. Brad keeps missing flower-arranging practice. He never reads about flower arranging in his free time and seems disengaged when he does attend practice. Brad attends every chess practice. He studies during his spare time and brings in strategies to share with others that enrich their practices.

There were four questions in two blocks. In one block, the evaluator was acting as leader of group A; in the other block, the evaluator was acting as leader of group B. In each block, there was a reward judgment and a punishment judgment, which was made on a 1–6 Likert scale: Not at all, mostly no, slightly no, slightly yes, mostly yes, completely. The rewards were always an award or honor (e.g., employee of the month: “Do you think Brad should receive the MVP award?”); the punishment was always revocation of group membership (e.g., “Do you think she

should ask Brad to leave Rookies?”).

7.1.3. Design and procedure

Participants were randomly assigned to respond to 2 out of 4 vignettes. Each vignette was accompanied by two sets of questions (a reward and punishment decision per group) presented in randomized order. This created a two-by-two design: Judgment type (reward or punishment) and the individual's behavior (good or bad) relative to the group the decider was acting as leader of.

7.2. Results

There was a significant interaction, $b = 3.54$, $SE = 0.31$, $p < .001$, between the judgment being considered (reward vs. punish) and the group the decider was representing (Fig. 7). Participants indicated that a leader should punish an individual only when acting as leader of the group whose norms the individual was violating, $b = -1.44$, $SE = 0.24$, $p < .001$, and should reward an individual only when acting as leader of the group whose norms the individual was upholding, $b = 2.10$, $SE = 0.23$, $p < .001$.

7.3. Discussion

We see that evaluations of a person are shaped by social structure, such that evaluations occur with respect to a group. Because both groups had the same members, and the evaluator was always the leader, the differences in evaluations cannot be explained by interpersonal relationships. The evaluations cannot be explained by a summative judgment of the individual behavior aggregated across roles because the same behavioral history was rewarded or punished solely as a function of the leader's role. Yet, the punishment decision was substantial: Removal from the group. In other words, participants thought it was consistent for a leader to both reward and ostracize the same individual for their role-specific behaviors.

We claim that social structure is an important part of group concepts and that we can observe this in evaluations. We do not claim that social structure is the only source of evaluations. People could apply universal moral principles—though this is not necessarily true (Chalik and Rhodes, 2020; Melnikoff and Bailey, 2018). Indeed, rule violations may be dual character; they are the combination of technical language and deeper value-based considerations that can incorporate moral reasoning (Struchiner and Hannikainen, 2023). Groups interact with other

institutions and can be contained within them. Illegality can motivate expulsion from many groups even if the illegality is not directly relevant. Finally, interpersonal behaviors matter. A person who is generally unkind and mean could trigger negative evaluations that generalize to all contexts in which the person's behavioral history is known. We find context-specific evaluations precisely because the behaviors in question were not immoral, illegal, or interpersonally mean—the behaviors were right or wrong in virtue of a social structure. Therefore, this study shows that social structures—by being normative—are a sufficient (but not necessary) grounds for evaluation.

8. General discussion

Six studies support the following conclusion: People represent organized groups as having at least two parts—members and structure. A difference in structure is sufficient to numerically distinguish groups, so complete overlap in membership is insufficient to identify a group. Therefore, we can conclude that social structure is an important part of how people represent an organized group and should be incorporated in models of group concepts. Although we have focused on structure as a central part of groups, we cannot claim that structure would complete these models. Even the combination of structure and membership may be insufficient to identify a group because there may be additional parts. Indeed, we cannot claim that structure is the most central part of group concepts without a full description of the other parts of a group. Nevertheless, our results suggest that social structure is indispensable for a complete theory of organized groups.

8.1. Proper function

A candidate third part of groups is proper function—the purpose of the group. People create groups because they intend to cooperate in a joint activity, such as playing baseball, regulating fishing in public lakes and rivers, and advancing Art History education and research. Indeed, the purpose of a group is an important part of how we distinguish kinds of groups—e.g., “is it a baseball team or a basketball team?” Neither of these considerations require function to be part of the identity of a group. For example, we distinguish people by their functions (“is she a lawyer or a doctor?”) and objects can afford multiple uses (e.g., repurposing a hammer for a paperweight); yet, people presumably believe that a person survives a career change and a hammer survives its many uses. Still, function appears to be an important component of group concepts, which raise the possibility that at least the proper function—the purpose—of a group could be part of its numerical identity (cf. Noyes & Keil, 2020).

Therefore, as an empirical matter, differences in function may be sufficient to distinguish numerical identity. Study 2 speaks against this possibility because participants identified a group by structure even when there were two functions. However, the pragmatic effect of showing one photograph could have made the difference. In a supplement, we report two studies that held the number of photographs constant. In both studies, we find (i) a significant condition difference; (ii) participants indicate two groups when there are two structures (iii) but no difference from chance when the groups differ in their function. Therefore, although Study 3 shows that structure is sufficient to distinguish groups, the supplemental studies raise the possibility that function enters identity judgments (though not as strongly as structure). Because participants did not systematically report two groups, the results are ambiguous. Indeed, participants may have wondered whether the groups truly had the same structure. That is, function could be a cue to structure rather than a third part of identity judgments. For example, imagine we provided two identical diagrams of a basketball team and a baseball team depicting the relationships between the roles manager, head coach, assistant coach, and players. No one would suppose that the leadership diagrams exhaust the structure of the group because the nature of basketball and baseball necessitate additional levels of



Fig. 7. Study 5. Judgments of whether an individual should be punished (left) or rewarded (right) based on whether the individual's behavior was good (red) or bad (blue) relative to the group the reward or punishment was administered by. Error bars are 95% confidence intervals. The dotted line is the midpoint of the scale. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

structure—within “player,” there are many additional roles, such as pitchers (for baseball teams) and point guards (for basketball teams). Indeed, the supplemental studies relied on leadership structures; in combination with two photographs, participants may have inferred additional lawyers of structure within non-leadership roles. Therefore, we need further empirical work to understand whether and how function enters identity judgments.

As a conceptual matter, it is unclear whether proper function should even be distinguished from structure—a consideration that holds for artifacts too. We can use a hammer in an indefinitely many ways—to drive nails into wood, break ice, weigh down papers, scratch our backs, and so on. However, there is also the proper function of the hammer. When a person sat down and designed the ball pen hammer, the design was meant to achieve a purpose—to shape metal materials. This purpose explains why ball pen hammers have the structure they do, whereas their incidental functions do not. Likewise, a baseball team can star in a movie, clean up litter, or speak to a first-grade class. However, playing baseball is their proper function and playing baseball (and not other incidental functions) explains their structure. It is not an accident that major league baseball teams are as similar in structure as they are in proper function, and it is not an accident that major league baseball teams are as different in structure from academic departments as they are in proper function. Therefore, although it may be intelligible to distinguish structure from incidental function, it may not be intelligible to distinguish structure from proper function. We think the best avenue for future work is to empirically investigate the connection between structure and proper function in group concepts; if they are entangled, it may conceptually incoherent for groups to merge their structures but not their proper functions (and vice versa).

8.2. Normativity

Study 6 suggests that people take the roles a person plays to be relevant when making judgments about reward and punishment. In a supplement, we also find that people distinguish between betraying a group and betraying its current members. These findings suggest that people reason about normative issues, at least in some cases, by centering roles and relationships. Major contemporary ethical theories—particularly consequentialism and deontology—are largely impartial and universal in ways that do not make roles or relationships central. Our findings suggest that relational ethical theories, like care ethics or role ethics, are important ways of understanding at least how we in fact reason when considering people as members of groups (Evans and Smith, 2018; Gilligan, 1982; Held, 2006). This is not to say that these theories tell us how we *ought* to reason about normativity or morality, but our findings suggest that relational theories play an important part in judgments of praise and blame. Even when the same people are involved, judgments about praise and blame are made relative to roles a person plays (see Earp, McLoughlin, Monrad, Clark, and Crockett, 2021 for similar finding). However, Study 6 reveals only that social structures are sufficient for evaluations—a behavior that is otherwise amoral, legal, and mundane can become exemplary or condemnable in the context of a role. Future work can examine how people integrate multiple sources of evaluations, which could provide information about how roles are prioritized in complex evaluations.

8.3. Social ontology

We investigated how people represent groups and found that people include social structure as a part of the identity of a group. Philosophers working in social metaphysics and social ontology have examined questions about what groups are, how much change they can undergo, and how they are differentiated from one another (Epstein, 2019; Hindriks, 2013; Ritchie, 2015, 2020; Tuomela, 2013). While our findings here are not about groups themselves, representations matter when considering the nature of social entities. For example, the rights and

obligations that are afforded with citizenship in the United States and the function of dollar bill relate to collectively accepted decisions. This marks a contrast with, for instance, the best theories of biological species or subatomic particles. Whether we represent biological species as having underlying genetic essences is simply irrelevant to whether species really have underlying essences (Gelman, 2003, 2004; Leslie, 2013). While the nature of social entities is certainly not always obvious (Khalidi, 2015; Thomasson, 2009), representation is an important feature in how we construct the social world. So, while our findings do not directly answer questions about the nature of social groups, along with a theory of how representation and social reality connect, our results could support a view on which groups are structured entities (Ritchie, 2015, 2020).

8.4. Diachronic identity

We asked participants to count the number of groups. We did not ask them to judge the persistence of a group over time. We know from prior work that people perceive groups as surviving change (Noyes & Keil, 2020). Consider a group in which half the members meet Monday and half meet Wednesday. People say the group survives when the division is merely practical. However, if the division is ideological (e.g., a pro-life and pro-choice meeting), people perceive the creation of two new groups. One interpretation is that people perceive a practical fracture as preserving the shared norms that bind people into a common structure, whereas the ideological fracture severs these binding norms. Indeed, real-life schisms occur this way (Sani & Reicher, 1998, 2000). When two parties see the other as acting contrary to the group's purpose, the value-laden disagreement can prompt group fission, creating one or more new groups.

Our theory about synchronic identity motivates a couple predictions about diachronic identity. First, if people and structures are parts of groups, then a group cannot survive the permanent loss of membership—the structure must be realized by people. Second, if structure is a central component of groups, then fission events that cleave structure create new groups, and can end identity. For example, consider the fracturing of Harvard's biology department: In 1967, biochemistry and molecular biology established a new department; in 1971, the remaining faculty voted to form two committees (one for evolutionary and organismic biology, and one for cellular and developmental biology), which acquired increasingly distinct administrative infrastructures until their formal division in 1982. People should count two departments by 1967 and three by 1982, and they should see the original biology department as ending by 1982. In general, we expect diachronic judgments will combine domain-specific intuitions about groups with domain-general intuitions about causal continuation (Rips, Blok, and Newman, 2006). That is, people look for the closest causal continuation of an original object, and perceive the end of identity when there is no one continuation. This may explain why Harvard's biology department appears to lose an area in 1967 but dissolve in 1982. The event in 1967 was uneven—a third of the department leaves, leaving the biology department to persist as the remaining two-thirds. The event in 1982 is even—two halves go separate ways, leaving the original biology department with no identifiable continuator.

Future work can build on our structural account of synchronic identity and existing work on diachronic identity (Noyes & Keil, 2020) to ask new questions about diachronic identity. What is the threshold of structural change a group can withstand? Are functional changes to structure more disruptive than mere bureaucratic changes to structure? Are illegitimate changes more disruptive than legitimate ones? We think these are interesting questions meriting future work. We think diachronic identity will involve intuitions about the mechanisms through which a structural replicates itself over time (e.g., admitting new members), and that people will be generous about structural changes insofar as they occur gradually. We think structural changes that disrupt the proper function of the group will be more disruptive.

Prior work demonstrates that groups can withstand drifting from their proper function (Noyes & Keil, 2020). However, this work involved mission drift rather than a complete overhaul of proper function (e.g., a shift from baseball team to basketball team). Complete overhauls in proper function entail corresponding overhauls in structure, which may terminate group identity.

9. Conclusion

Our studies support the conclusion that structure is an important part of group concepts and is sufficient to numerically distinguish groups. We have provided little detail into how people represent structure, though. For example, what do people believe makes a person a member of a role? Is it a status assigned by formal rules—e.g., Lisa received a formal offer and signed the necessary paperwork? Is it an informal recognition—e.g., everyone recognizes Lisa as having the job? Is it a practical matter—e.g., Lisa shows up and does the job? Indeed, being in a role may be a bundle of many such factors, each entailing different consequences. How do people view the relationship between structure and proper function? People may understand them as deeply connected. For example, if a structure is represented as a solution to a coordination problem (e.g., how to coordinate education and research), then the purpose of the group is to enact its structure. Do people distinguish layers of structure? For example, how do people represent the relationship between leadership (e.g., coach, player) and player roles (e.g., point guard, forward)? Do people prioritize layers of structure—e.g., player roles over leadership; faculty over administrators? Group concepts may encompass a hierarchy of how central roles and other components of structure are to groups, and that hierarchy may be proportional to how essential they are to the proper function of the group. Although baseball teams and academic departments would be dramatically less effective without their coaches and administrators, one may conclude there is not baseball team or academic department without players and faculty. This could have powerful effects on group concepts. For example, when people hear “the Art History department is meeting in Room 101,” they may picture and expect people proportional to how central they view their roles. Although we cannot answer these questions, the present investigation shows that they are worth asking because structure is an important component of group concepts.

Author's Note

This study's design, hypotheses, and analysis plan was preregistered. Study materials, data files, and pre-registrations available at: <https://osf.io/wmnqd/>

CRediT authorship contribution statement

Alexander Noyes: Conceptualization, Methodology, Data curation, Writing – original draft, Writing – review & editing. **Frank C. Keil:** Writing – review & editing. **Yarrow Dunham:** Writing – review & editing. **Katherine Ritchie:** Conceptualization, Methodology, Writing – review & editing.

Data availability

Data will be made available on request.

References

- Baumeister, R. F., Ainsworth, S. E., & Vohs, K. D. (2016). Are groups more or less than the sum of their members? The moderating role of individual identification. *Behavioral and Brain Sciences*, 39.
- Chalik, L., & Rhodes, M. (2020). Groups as moral boundaries: A developmental perspective. *Advances in Child Development and Behavior*, 58, 63–93.

- Coley, J. D., Arenson, M., Coley, J. D., & Tanner, K. D. (2017). Intuitive biological thought: Developmental changes and effects of biology education in late adolescence. *Cognitive Psychology*, 92, 1–21.
- Earp, B. D., McLoughlin, K. L., Monrad, J. T., Clark, M. S., & Crockett, M. J. (2021). How social relationships shape moral wrongness judgments. *Nature Communications*, 12(1), 1–13.
- Epstein, B. (2019). What are social groups? Their metaphysics and how to classify them. *Synthese*, 196(12), 4899–4932.
- Evans, J., & Smith, M. (2018). Toward a role ethical theory of right action. *Ethical Theory and Moral Practice*, 21, 599–614.
- Fine, K. (1999). Things and their parts. *Midwest Studies in Philosophy*, 23, 61–74.
- Gelman, S. A. (2003). The essential child: Origins of essentialism in everyday thought. In *Oxford Cognitive Development*.
- Gelman, S. A. (2004). Psychological essentialism in children. *Trends in Cognitive Sciences*, 8, 404–409.
- Gelman, S. A., & Rhodes, M. (2012). Two-thousand years of stasis: How psychological essentialism impedes evolutionary understanding. In K. R. Rosengren, S. Brem, E. M. Evans, & G. Sinatra (Eds.), *Evolution challenges: Integrating research and practice in teaching and learning about evolution* (pp. 3–21). New York, NY: Oxford University Press.
- Gelman, S. A., & Taylor, M. G. (2000). Gender essentialism in cognitive development. In P. H. Miller, & E. Kofsky Scholnick (Eds.), *Toward a Feminist Developmental Psychology* (pp. 169–190). Taylor & Francis/Routledge.
- Gilbert, M. (1992). *On Social Facts*. Princeton University Press.
- Gilbert, M. (2013). Considerations on joint commitment. *Joint Commitment*, 37–57.
- Gilligan, C. (1982). *In a Different Voice: Psychological Theory and Women's Development*. Cambridge, MA: Harvard University Press.
- Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition: A Multidisciplinary Debate* (pp. 351–394). Clarendon Press/Oxford University Press.
- Hamilton, D. L., & Sherman, S. J. (1996). Perceiving persons and groups. *Psychological Review*, 103(2), 336–355.
- Hamilton, D. L., Sherman, S. J., & Lickel, B. (1998). Perceiving social groups: The importance of the entitativity continuum. In C. Sedikides, J. Schopler, & C. A. Insko (Eds.), *Intergroup Cognition and Intergroup Behavior* (pp. 47–74). Lawrence Erlbaum Associates Publishers.
- Haslam, N., Rothschild, L., & Ernst, D. (2000). Essentialist beliefs about social categories. *British Journal of social psychology*, 39(1), 113–127.
- Haslanger, S. (2000). Gender and race: (what) are they? (what) do we want them to be? *Notis*, 34, 31–55.
- Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies*, 173(1), 113–130.
- Held, V. (2006). *The Ethics of Care: Personal, Political, and Global*. New York: Oxford University Press.
- Hindriks, F. (2013). The location problem in social ontology. *Synthese*, 190(3), 413–437.
- Khalidi, M. (2015). Three kinds of social kinds. *Philosophy and Phenomenological Research*, 90, 96–112.
- Koslicki, K. (2008). *The Structure of Objects*. Oxford University Press.
- Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences*, 98(26), 15387–15392.
- Leslie, S. J. (2013). Essence and natural kinds: When science meets preschooler intuition. *Oxford Studies in Epistemology*, 4, 108–165.
- Lickel, B., Hamilton, D. L., Wiecekowska, G., Lewis, A., Sherman, S. J., & Uhles, A. N. (2000). Varieties of groups and the perception of group entitativity. *Journal of Personality and Social Psychology*, 78(2), 223–246.
- Melnikoff, D. E., & Bailey, A. H. (2018). Preferences for moral vs. immoral traits in others are conditional. *Proceedings of the National Academy of Sciences*, 115(4), E592–E600.
- Misch, A., Paulus, M., & Dunham, Y. (2021). Anticipation of future cooperation eliminates minimal ingroup bias in children and adults. *Journal of Experimental Psychology: General*, 150(10), 2036.
- Noyes, A., & Dunham, Y. (2017). Mutual intentions as a causal framework for social groups. *Cognition*, 162, 133–142.
- Noyes, A., & Dunham, Y. (2020). Groups as institutions: The use of constitutive rules to attribute group membership. *Cognition*, 196, 104143.
- Noyes, A., & Keil, F. C. (2020). Collective recognition and function in concepts of institutional social groups. *Journal of Experimental Psychology: General*, 149(7), 1344.
- Pietraszewski, D. (2022). Toward a computational theory of social groups: A finite set of cognitive primitives for representing any and all social groups in the context of conflict. *Behavioral and Brain Sciences*, 45.
- Porpora, D. V. (2013). Four concepts of social structure. In *Critical Realism* (pp. 339–355). Routledge.
- Rabbie, J. M., & Horwitz, M. (1988). Categories versus groups as explanatory concepts in intergroup relations. *European Journal of Social Psychology*, 18(2), 117–123.
- Rhodes, M., & Chalik, L. (2013). Social categories as markers of intrinsic interpersonal obligations. *Psychological Science*, 24(6), 999–1006.
- Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, 113(1), 1–30.
- Ritchie, K. (2013). What are groups? *Philosophical Studies*, 166, 257–272.
- Ritchie, K. (2015). The metaphysics of social groups. *Philosophy Compass*, 10, 310–321.
- Ritchie, K. (2020). Social structures and the ontology of social groups. *Philosophy and Phenomenological Research*, 100, 402–424.
- Sani, F., & Reicher, S. (1998). When consensus fails: An analysis of the schism within the Italian Communist Party (1991). *European Journal of Social Psychology*, 28(4), 623–645.

- Sani, F., & Reicher, S. (2000). Contested identities and schisms in groups: Opposing the ordination of women as priests in the Church of England. *British Journal of Social Psychology, 39*(1), 95–112.
- Shtulman, A., & Schulz, L. (2008). The relation between essentialist beliefs and evolutionary reasoning. *Cognitive Science, 32*(6), 1049–1062.
- Steiner, I. P. (1974). Whatever happened to the group in social psychology? *Journal of Experimental Social Psychology, 10*(1), 94–108.
- Struchiner, N., & Hannikainen, I. R. (2023). Rule is a dual character concept. *Cognition, 230*.
- Thomasson, A. (2009). Social entities. In R. Le Poidevin (Ed.), *The Routledge Companion to Metaphysics*. Routledge.
- Thomasson, A. L. (2016). The ontology of social groups. *Synthese, 196*(12), 4829–4845.
- Tuomela, R. (2013). *Social ontology: Collective intentionality and group agents*. Oxford University Press.