

The Emergence of Specialized Roles Within Groups

Robert L. Goldstone
Indiana University

Edgar Andrade-Lotero
Universidad del Rosario

Robert D. Hawkins
Princeton Neuroscience Institute

Michael E. Roberts
DePauw University

Author Note

Robert L. Goldstone. Department of Psychological and Brain Sciences and Cognitive Science Program, Indiana University, Bloomington, IN.

Correspondence concerning this article should be addressed to Robert Goldstone, Department of Psychological and Brain Science, Indiana University, Bloomington, Indiana. 47405. Email: rgoldsto@indiana.edu

Abstract

Humans routinely form groups to achieve goals that no individual can accomplish alone. Group coordination often brings to mind *synchrony* and *alignment*, where all individuals do the same thing (e.g., driving on the right side of the road, marching in lockstep, or playing musical instruments on a regular beat). Yet effective coordination also typically involves *differentiation*, where specialized roles emerge for different members (e.g., prep stations in a kitchen or positions on an athletic team). Role specialization poses a challenge for computational models of group coordination, which have largely focused on achieving synchrony. Here, we present the CARMI framework, which characterizes role specialization processes in terms of five core features that we hope will help guide future model development: Communication, Adaptation to feedback, Repulsion, Multi-level planning, and Intention modeling. Although there are many paths to role formation, we suggest that roles emerge when each agent in a group dynamically allocates their behavior toward a shared goal to complement what they expect others to do. In other words, coordination concerns beliefs (who will do what) rather than simple actions. We describe three related experimental paradigms —“Group Binary Search”, “Battles of the Exes”, and “Find the Unicorn” — that we have used to study differentiation processes in the lab, each emphasizing different aspects of the CARMI framework.

Keywords: collective intelligence, collective behavior, joint action, division of labor, specialized roles

Word Count: 8,316

The Emergence of Specialized Roles Within Groups

1. Synchronization vs. Specialization

When most people think about group coordination, they imagine group members matching their actions in time, as when an audience spontaneously synchronizes their applause at the end of a concert or fireflies synchronizing their flashes of light (Strogatz, 2003). This kind of synchrony is often taken to be the default route to successful coordination for human groups. Well-functioning teams are “in sync,” with members working together “in unison.” For example, theoretical interest in behavioral matching phenomena — mimicking one another’s body language, vocal features, and actions — stems from the premise that synchrony is an engine of empathy and cooperativity in groups: when a person aligns their behavior with another person’s, the imitated person supposedly likes the imitator more and subsequently engages in more pro-social behavior towards them (Lakin et al., 2003; Sebanz, Bekkering, & Knoblich, 2006; Hove & Risen, 2009). In this way, it seems like groups are drawn toward conformity to function (Asch, 1956). Even in adversarial tasks with explicit incentive not to synchronize with others, people often still end up synchronizing their spatial choices (Frey & Goldstone, 2013) and foraging positions (Goldstone & Ashpole, 2004).

Although synchronization may be one mechanism for facilitating coordination, social life is rife with examples of groups working better when members do *different* things. For example, players on a soccer team assume roles such as defending a particular opponent or feeding the ball to their striker/forward/center (Bowling & McCracken, 2005), and the doctors and nurses in an operating room assume roles such as managing anesthesia or intubation (Helmreich & Schaefer, 1994). At a macro-level, most organizations have a top-down management chart specifying distinct roles and their relationships — for instance, engineers may not be involved in advertising. But at a more micro-level, individuals on a single engineering team may adopt *ad hoc* functional roles within a given meeting (e.g. Dong et al., 2013), or dynamically specialize on sub-tasks (e.g. Valentine & Edmondson, 2015; Dafoulas & Macaulay, 2001). Often, the division of a group into specialized roles is task-dependent and at least partially self-organized rather than being top-down stipulated (Stone & Veloso, 1999; Allen et al., 2015; Wu et al., 2021).

Given the complexities of the modern world, coordination via role specialization may be even more ubiquitous than coordination via synchronization. The tasks facing large groups such as factories, laboratories, orchestras, and governments are too complex for any individual to be able to understand everything that must be accomplished by the group and how to best accomplish it (Henrich, 2016; Laland, 2017). For example, the conductor has an excellent overview of what is required to play each of the instruments in her orchestra, but cannot play them all at a sufficiently high level of expertise. Developing a professional level of play in any single instrument consumes a significant portion of an all-too-short lifespan. Generally speaking, members of modern societies rely on specialized experts to fix their cars, hearts, plumbing, and computers. In fact, most of what we think we know is not contained inside our own skulls, but is distributed across our social network of friends and colleagues (Sloman & Fernbach, 2017; Sloman & Rabb, 2016).

We view role specialization as one distinctive solution to the broader class of division of labor problems. Not all division of labor involves different roles. For example, if multiple cakes must be baked, and each member of the group bakes one cake in parallel, they have successfully divided up the labor without meaningfully adopting different roles. If those cakes needed to be baked in a shared kitchen, under resource constraints, then the dependencies between agents increase and stable roles may emerge. For example, some bakers may allocate themselves to a ‘morning shift’ vs. ‘night shift’ to take turns. Roughly speaking, in the kinds of phenomena we review in this paper, agents are not perfectly exchangeable. Even if they start a task being exchangeable, they are no longer exchangeable as a result of their interactions. Differentiation may be based on which actions should be performed, when and where certain objects should be acted upon, or both (e.g., Raveendran, Puranam & Warglien, 2016). For example, members of a security detail may split up to keep watch over different areas of a venue. In that case, all agents may be doing the same action (visual search) at a broad level but are carrying out the action on complementary objects or spatiotemporal regions (the ‘roof’ role vs. the ‘stage’ role). Alternatively, in the sports, music, and organizational chart examples, team members are specialized for distinct actions (the ‘striker’ role vs. the ‘sweeper’ role) but may be acting on the same objects in the same spaces. Moreover, in both cases, the division of labor can either be imposed top-down or may instead arise through bottom-up self-organization (e.g., Richardson, Harrison, Kallen, Walton, Eiler, Saltzman & Schmidt, 2015). We will primarily focus our review on features of bottom-up self-organization, as these remain less well-understood.

2. The CARMI Framework for Role Specialization

In this section we propose an overarching framework that provides a broad characterization of group-level role specialization in terms of five core features (Durkheim, 1893; Rabb, Fernbach, & Sloman, 2019): Members of a group [C]ommunicate their plans and proposals for coordination, [A]dapt to feedback on how well they are doing at a task, are [R]epulsed by their partners' strategies, develop [M]ulti-level plans to facilitate division of labor, and create internal models of their partners' [I]ntentions and knowledge. In this section, we will walk through each of the features of the CARMI framework making reference to the example illustrated in Figure 1. Suppose a pair of agents are building a restaurant together. This is a remarkably complex task with many inter-locking subtasks. What allows them to effectively collaborate, and how do they divvy up the work?

a. Communication of intentions, plans, and proposals

The most obvious (and effective) way to establish roles is for the agents to explicitly communicate with each other using verbal or non-verbal cues. The conventions provided by a shared language enable efficient access to abstract intentions, plans, and proposals (Hawkins et al., 2020; 2022). For example, Alice (on the left) could simply say, "I'd be happy to handle the bathrooms, since I have a lot of experience with plumbing" and Bob (on the right) could say, "That sounds good; I don't know much about plumbing, so I'll get started on the dining room design." Even if they don't speak the same language, non-verbal communication is still better than nothing. Alice could point to the bathroom area and Bob could give a thumbs up before wandering off to the kitchen. Although classic game theoretic accounts often conclude that non-binding promises ("cheap talk") should have no effect on coordination, groups in which members are allowed to freely communicate to make proposals, assurances, and promises regarding the distribution of roles to resource management are more likely to come up with efficient and fair cooperative schemes (Ostrom, 2006; Ostrom, Walker, & Gardner, 1994; Mirsky et al., 2021). For example, when people can discuss possible rules systems for divvying up a resource into complementary private properties, they collectively harvest that resource much more efficiently than when communication is not allowed

(Janssen, Goldstone, Menczer, & Ostrom, 2008). Furthermore, early effort expended on communication produces dividends for future interactions, as the negotiated roles may continue to be adopted on subsequent rounds when communication is no longer necessary or possible. A relatively small initial investment in explicit coordination continues to reap benefits for the individuals in a group.

Communication is such a powerful mechanism that it tends to overshadow other factors driving role differentiation. But explicit communication is neither necessary nor sufficient for roles to emerge in groups. The remainder of the CARMI framework emphasizes factors driving role differentiation even in the absence of a shared communication protocol. To effectively observe these other processes at work, many studies of role formation intentionally restrict communication among agents. All three case studies we review have this property. We return in the discussion to the question of how explicit communication ought to work in tandem with these other factors of theoretical interest.

b. Adaptation to feedback

In the absence of explicit communication, the most basic ingredient for emergent role differentiation is a degree of sensitivity to feedback about what other individuals are doing and how the group is performing as a whole. Feedback may come in many forms, but whether it is weaker or stronger, it provides some signal which can be used to adjust one's own behavior to improve group performance. For example, returning to the restaurateurs in Figure 1, Bob could wander back into the bathroom where Alice is working and notice that she's not using the same tiles he's been using, or he could receive notice that their current plan violates building code, or that potential customers don't like the color scheme. Each of these observations ought to lead to specific adjustments in their subsequent policy, and potentially in their understanding of different roles.

All computational models of group coordination rely to some extent on feedback cues to drive adaptation and learning over time, but accounts differ concerning what forms of feedback are available and what learning rule should be used to update the policies. For example, consider a classical "win stay/lose shift" strategy (Sutton & Barto, 2018) where agents persist with a given strategy when receiving positive feedback and only switch strategies upon receiving negative feedback. This rule may be contrasted with a more graded version better glossed as "Stay closer to a

strategy as its outcome becomes increasingly good” (see Sloman et al., 2021), allowing agents to flexibly keep parts of their previous strategy while switching others (Campbell, Izquierdo, & Goldstone, 2022). We will examine the consequences of feedback in our case studies below, but for now, we simply emphasize that subtle differences in feedback-handling may lead to significant changes in group dynamics.

c. Repulsion from partner’s strategy

So far, these features seem to characterize all forms of group coordination, including synchronization. What accounts for the distinctive patterns of differentiation-based solutions? The third component, repulsion, aims to capture the descriptive flavor of complementary strategies found in these groups. As agents receive feedback and track other group members’ actions, they tend to be repulsed away from their partners’ strategies when conflicts are detected. They adjust their strategy towards what they perceive to be complementary aspects of the task. In Figure 1, Bob may come in one morning to find Alice working on a task in the kitchen that he is usually responsible for. Rather than inefficiently barreling into this same task, he may dynamically re-allocate himself on the fly to something else that doesn’t get in the way.

In empirical (Setzler & Goldstone, 2020) and computational (Setzler & Goldstone, 2022) analyses of improvising jazz musicians, simultaneously interacting musicians tended to play notes that harmonically complement their partners’ notes, compared to cases when one player lays down their musical track on top of a pre-existing track provided by another musician. Even if there is a prepotent tendency to synchronize actions, as some models predict, the need for complementarity in such environments may sometimes be in tension with this tendency, requiring effortful inhibition. For example, the strong tendency to act “in sync” with somebody with whom one empathizes could ironically lead to coordination difficulties when complementary strategies are required. When a group would benefit from differentiated roles, the natural inclination to engage in behavioral mimicry with in-group members (Chartrand & Lakin, 2013) could actually lead to poorer group outcomes.

d. Multi-level planning

A major challenge for models of group coordination is that role specialization unfolds across multiple time scales, spanning several orders of magnitude. At the long time scale of an entire

lifetime, role specialization is tantamount to the development of expertise. When considering careers, people often sort themselves into training paths that give them expertise in fields that they expect to be in demand and in which they have the aptitude to excel. At the intermediate time scale, people are remarkably adaptable to the composition of their teams and the problems they need to solve. For example, instead of many software engineers grappling with the same idiosyncrasies of a particular software package, it may make sense for a handful of individuals to master the details and become *de facto* experts for the community. Finally, at a short time scale, teams may need to spontaneously organize themselves into *ad hoc* roles for a specific, novel situation that they face (Misyak et al., 2014). For example, doctors in emergency triage units spontaneously assign themselves to patients, and defense teams in basketball dynamically adjust their player-to-player coverage to accommodate picks and lopsided matchups between opposing players.

Importantly, these scales interpenetrate one another. If a person dynamically adopts a particular role in enough individual micro-situations, then they will gradually become more proficient in the role, leading them to adopt it in subsequent situations, eventually leading to genuine expertise at the macro-level. Much like flowing water creating canals that channel subsequent water flows, roles adopted in the short term may become entrenched with repetition. For human groups to be as flexible as they are, agents must have some ability to reason and plan appropriately across these scales. For example, in Figure 1, when designing an individual bathroom fixture, Bob must not only account for what Alice is currently designing in the kitchen, but also what fixtures she is likely to include in the future. In line with the previous component of the framework, we suggest that locally arising conflicts may be instrumental in achieving longer-term coordination.

e. Modeling partners' intentions and knowledge

A final component of our framework is the capacity of agents to develop models of their partners' intentions and goals and to use these models to shape their own downstream behavior (Bratman, 1992). One cognitive mechanism that may support this capacity is agents making inferences about others' likely beliefs, knowledge, and intentions given observations of the actions others' take (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017). These inferences allow for flexibility across different contexts and groups. For example, in heterogeneous groups, inferences about the background knowledge or skills of each agent may support more effective role allocation, as less knowledgeable agents may defer to others who are better-suited for a particular task (Liemhetcharat

& Veloso, 2014). Such inferences may account for the repulsion of strategies that is characteristic of role specialization. Given strong enough social expectations — for example, that Alice has a deep knowledge and expertise in installing countertops and ventilation systems, and that Alice believes Bob is better at bathroom plumbing — Bob may head straight to the bathroom on the site without requiring any feedback or conflict at all.

More broadly, when agents are able to track where knowledge is stored within a group (Wegner, 1987), they are able to access it when needed (Rabb, Fernbach, & Sloman, 2019). Access to others' knowledge is often reliable enough that individuals give themselves credit for knowing something when it is not stored within their own heads but is readily accessible from other people (Sloman & Rabb, 2016). Groups can coordinate on complex tasks when their individuals have good models of what the others are trying to do (Goldstone & Theiner, 2017). Group members benefit by converging on their understanding of the overall task and what the possible actions are. This allows them to develop mutual understandings of the internal models that the players are using, and then choose actions/roles that lead to good individual and group solutions.

3. Experimental Investigations of Role Specialization

Our laboratory has developed several experimental paradigms for exploring the dynamics of self-organized role specialization. We will use our CARMI framework to organize the explanation of the emergence of role specialization for three reviewed paradigms. Communication was restricted, but the tasks vary on multiple dimensions such as the number of members, the number of available actions per member, and the types of available feedback. Our goal has been to develop purposefully simplified, parable-like paradigms in order to distill the dynamics of self-organized specialization, and observe theoretically important factors that could be obfuscated by real-world complexities. The time course of coordination that we study is rather short, over the course of tens of minutes. As such, our paradigms are ill-suited for studying the development of lifelong expertise or the dynamics of larger teams. Instead, the paradigms are useful for understanding how people spontaneously differentiate themselves into stable roles, even when they cannot rely on explicit communication or prior social knowledge. These paradigms thus provide lower bounds on our capacity to self-organize. That is, if the members of our experimental groups are able to form reliable specializations even in conditions that are adverse to their formation, it speaks to a broader capacity

for coordination through specialization. To foreshadow our findings, many of our groups in each of the paradigms do exhibit role specialization, and the extent to which they do is strongly associated with the success of the group at its assigned task. In what follows, we describe three experimental paradigms: Group Binary Search, Battle of the Exes, and Find the Unicorn. We compare the most relevant characteristics of these experimental paradigms in Table 1.

a. Group Binary Search

Our first paradigm was intended to capture the social phenomenon of people adapting their contributions to a group over time so that the group has the right level of total contributions. For example, if a group of friends decides to have a weekly potluck event, early on there may be too little or too much food. Over several weeks, the friends will adjust the amount of food that they prepare so that the total amount of food they bring comes close to the group's requirements. For a more abstract example of pooled contributions, imagine faculty members in a departmental meeting dynamically regulating how much they speak. Too little and too much discussion are both problematic, for reasons of adequately considering important perspectives on an issue and time management, respectively. Faculty members often regulate their own spoken contributions so that the total amount of discussion is appropriate¹. Both examples of dynamically adjusted pooled contributions are impressive because the members may differ markedly in their individual level of contributions, and there is typically no explicit communication between members about who will adjust their contributions.

In order to explore the coordination capacities of groups, we developed a simple round-based group game called "Group Binary Search". In computer science, binary search is a search algorithm used to find an item in a sorted array by repeatedly dividing the search interval in half, reducing the time complexity to $O(\log n)$ compared to the $O(n)$ time complexity of a linear search process that simply increments one item at a time starting with the first item in the array. For the sorted array [2, 5, 8, 10, 15, 16, 18, 19, 27] and the target item 10, the binary search algorithm would first check the item halfway in the array, the 5th item, 15, to see if it equals 10. Seeing that $15 > 10$, the algorithm then finds the item that falls halfway in the remaining possible items, 8, and seeing that $10 > 8$, it then

¹ Some faculty members seem to be trying to maximize how close the discussion comes to filling the entire allotted faculty meeting time rather than adequate coverage of important considerations, leading some to complain, "Whatever issue our department considers, no matter how trivial, somehow ends up filling an entire hour."

probes the only remaining item between 8 and 15. A simpler version of binary search is the familiar game of one person coming up with a random number between 1 and 100. A second person guesses numbers, with the first person providing “you’re too high” or “you’re too low” feedback until the first person correctly guesses the number.

In our group version of this game (Roberts & Goldstone, 2011), a computer server randomly chooses a number between 51 and 100, and without communication, each group member submits a contribution between 0 and 50. The computer compares the sum of participants’ numbers to its selected number, and broadcasts the same directional (e.g., “Too High”) or numeric (e.g., “Too Low by 17”) feedback to all members. During the next round, members adjust their guesses and receive new feedback, and the game continues until the group correctly sums to the computer’s chosen number. A demonstration of this group experiment can be run at <https://pc.cogs.indiana.edu/software-and-simulations/>.

It is immediately evident that the group version of this number guessing game is much harder than the individual version. In fact, the individual version is trivial in the case of numeric feedback. If a tester tells an individual guesser that their guess of 24 was too low by 12, then the guesser immediately knows that the target number is 36 and will guess that on the next round. However, even exact numeric feedback does not allow the group to immediately solve its task. If the group is told that it is collectively too low by 17, this does not help an individual know whether they personally should adjust their guess and by how much. We find that the average number of rounds needed for a group to guess the target number increases as group size increases, echoing modeling work showing similar coordination benefits for small over large groups (Galesic, Barkoczi, & Katsikopoulos, 2018). The extent to which this task is harder for the group than the individual is a good measure of how challenging group coordination is.

Figure 2 shows that both directional and numeric feedback versions of the task are challenging for groups comprised of 2-17 participants. Groups are able to use the greater amount of information in the numeric than directional feedback to reduce the average number of rounds required for solution, but they are still far slower than individuals would be in the numerical feedback version. Groups play the same game multiple times, and over the course of multiple plays, they become more

efficient at finding solutions. Analyses indicate that this is because the members of a group develop different roles.

A good way to think about roles in this game is through the notion of reactivity to feedback supplied by the computer. Some players strongly react to feedback such that if the computer tells all of the players that their sum was too high by 23 then they will considerably reduce their contribution on the next round. Other players react only slightly, others do not react at all, and a few players react in the opposite way as expected (e.g., guessing an even higher number) because they predict that their fellow players will overreact. If we denote by G_r a participant's contribution on round r , we can quantitatively measure the reactivity of a participant's contribution on round r by $(G_r - G_{r-1})$ if the group's sum was lower than the target number on the previous round, and $(G_{r-1} - G_r)$ if the group was too high. Figure 3 shows how groups react to feedback in aggregate. In general, groups react in the appropriate direction to feedback, but not strongly enough. The best fitting line relating the groups' deviation from the target number to their reaction on the subsequent round is flatter than it would ideally be.

This group-level analysis conceals crucial individual differences among the members of a group. Over the course of an experiment, individuals tend to adopt reactivity roles, with some players reacting much more to feedback than others. Over the course of an experiment, the similarity of round-to-round reactivities within a player tends to increase while the similarity of reactivities across players decreases. This pair of results is a good quantitative signature for the essence of role specialization – an agent developing a consistent pattern of behavior which becomes differentiated from their teammates. Questions about the causes and consequences of role specialization can be answered using this metric of self-similarity (consistency) and other-dissimilarity (differentiation).

One influence on role specialization is group size. As group size increased in our experiment, players tended to adopt more consistent and more differentiated reactivities. Role specialization seems to be spontaneously discovered out of the necessity of coping with large groups. Rounds to solution increased markedly as group size increased, and members within large groups realized that they needed to become predictable to their peers while also taking actions that complemented their peers' actions. Consistent with this interpretation, we found that groups with members that engaged in more role specialization performed better at the group-level task, particularly so for large groups.

Additional analyses showed how players used feedback to adjust their reactivities. Overall, players decreased their reactivities as more rounds within a game were played. When the direction of feedback switched from “too high” to “too low” or vice versa, players tended to markedly reduce their reactivities. Incorporating these strategies into an agent-based model led to better fits to the human results compared to baseline models without these strategies. Consistency was incorporated into the agent-based model by having each agent randomly sample its own level of reactivity from a distribution of reactivities, and then slightly adjusting this personalized reactivity based on the round number and feedback. This model that incorporated consistent reactivity roles fit our empirical results better, and solved the task faster, than models without agent consistency.

Broadly speaking, the experiments and modeling for Group Binary Search show the importance of people dynamically assuming roles. By reacting in a consistent way that is differentiated from their partners, players help their group solve its task. Over the course of the experiment, groups improve their performance in large part because of this role specialization process. With respect to the CARMI framework, the empirical results of Group Binary Search illustrate clear Adaptation to feedback, with individual members increasingly Modeling and recognizing others’ likely patterns of behavior and adopting Repulsed reactivity patterns that collectively achieve appropriate collective-level reactivity.

b. Battle of the Exes

It is fairly obvious to the participants in Group Binary Search that they are facing a challenge of coordination. They are given an explicit group-level task — all members of the group share the same goal — and on every round of guesses, they understand that their guess must combine together with others’ guesses exactly right. In a second paradigm we call Battle of the Exes (Hawkins & Goldstone, 2016), we were interested in whether people develop complementary roles even when it is not clear to them that they are in a game of coordination. Participants are given an individual goal of maximizing their own points. However, coordinating with their partner is the only way to achieve a solution that is jointly stable, fair, and efficient — three priorities of social contracts (Binmore, 2013).

The name “Battle of Exes” derives from an imagined situation in which two ex-lovers are no longer speaking with each other and are each trying to avoid the other. Suppose there are only two coffee shops in their town, one with better coffee than the other. Both exes want to go out for coffee during their simultaneously occurring coffee breaks, but if they both pick the same place and run into one another, both will be unhappy. One of the exes could go to the better coffee shop every day, but that would be unfair. Each could choose randomly, but that would end up with exes often seeing each other, which would not be efficient, and would not provide a stable solution in the long run.

In our real-time instantiation of Battle of the Exes, 568 players were randomly paired together to create 284 two-player groups. On each round of play, the two players are given the choice of moving their avatar to one of two circles, one that they can visibly see will give them a small monetary payoff and one that will give them a larger payoff. The only catch is that if both players move to the same circle, then neither player gets anything for that round. For half of the groups, there was a small discrepancy between the prizes (1 cent vs 2 cents), and for the other half, there was a large discrepancy (1 cent versus 4 cents). Also, for half of the groups, each of the players could see the other player’s moment-to-moment position as they moved avatars to the circles (Dynamic movement), while for the other half of the groups, the players could only see the final choice that the other player made, as in traditional game-theoretic tasks (Ballistic movement).

The efficiency of a dyad’s choices can be measured as the average of their payoffs across rounds. We found that the dynamic condition led to better overall earnings for both players than the ballistic condition (Figure 4A), with participants rarely colliding. In other words, coordination is promoted when the players have full information about one another’s moment-to-moment inclinations. Second, the fairness of a dyad’s choices can be measured by the ratio of the average earnings of the player who earned less relative to those of the player who earned more. A ratio of 1 would indicate a perfectly fair distribution of earnings. The dynamic condition led to fairer solutions than the ballistic condition, with players earning similar amounts of money (Figure 4B). One implication of these results under the CARMI framework is that giving the members in a group more information about what others are currently thinking about doing may help them to better Model others’ intentions and Adapt accordingly, allowing the group to achieve more well-coordinated and fair solutions. Reinforcement learning models that combine a feedback control loop to handle low-level

Repulsion with higher-level, long-term reward-maximization adaptation replicate these advantages for the dynamic over ballistic condition (Freire et al., 2020). This is something for politicians, social network sites, and amusement parks to consider when they are trying to design social spaces for their groups. Mutual visibility of group members is often an effective way to promote coordination.

What are the roles here? Broadly speaking, there were two ways of achieving stable, efficient, and fair outcomes. First, the players A and B might alternate over rounds who gets the large payoff (ABABABABAB). In this case, the roles correspond to temporal differentiation between ‘even trial’ vs. ‘odd trial’ winners (similar to bakers self-organizing into a morning shift and an afternoon shift). Second, they could exploit the coordinated equilibrium introduced by the random spatial assignment of payoffs. If A always goes to the ‘top’ target and B always goes to the ‘bottom’ target, their payoffs from round to round will vary but they will roughly even out over the course of the session. In this case, the roles correspond to ‘top’ and ‘bottom’ locations (similar to different basketball players consistently defending their respective targets). But how do we know these roles are stable? We captured the notion of stability in a graded, quantitative measure using the information-theoretic measure of surprisal, which Shannon (1948) defined as the negative logarithm of the probability of an event. In particular, we computed the average surprisal of a Markov Chain encoding the transition probabilities between events on successive rounds. Intuitively, under stable roles, knowing the outcome of a social interaction at one point in time should reduce one’s uncertainty about what will happen at other points in time. An unlikely event will have low probability and an observer would therefore be highly surprised to see it happen, given knowledge of other events (see also Vanderschraaf & Skyrms, 2003).

In terms of developing stable roles, there was a striking interaction between payoffs and visibility (see Figure 4C). In the low-stakes condition, choices in the ballistic condition were more stable than in the dynamic condition. Players in the dynamic condition simply relied on moment-to-moment visual information to figure out who should get the larger payoff on any given round. They did not feel a strong pressure to develop a consistent policy because they could use the continuously available information about player positions to help them coordinate in an ad hoc fashion. However, when the stakes were high, with one circle earning four times what the other circle earned, then the dynamic condition developed even more stable solutions than the ballistic condition. For these higher stakes situations, it was useful for the players to develop strong norms to help them

coordinate, and moment-to-moment information about player positions helped to create these norms. Interestingly, the ‘temporal alternation’-based roles were more prevalent in the dynamic conditions, whereas the ‘spatial direction’-based roles were more prevalent in the ballistic conditions.

One finer-grained metric of ‘conflict’ in a group is the length of time both players spend moving toward the high payoff option before one “peels off” and lets the other player have the high payoff prize. Using “peel off” time as an objective measure, groups generally have more conflict at the beginning of the experiment session than the end. The higher-stakes condition has more conflict early on than the lower-stakes condition, but by the end of the experiment, the ordering is flipped. Groups that have more conflict at the beginning of the experiment tend to have less conflict than other groups by the end of experiment, and are more likely to develop stable, fair strategies. One take-home message is that conflict in groups is not necessarily something to be avoided; especially in high-stakes situations, it may spur negotiation that leads to more stable, long-term solutions. It may be tempting to try to pave over disagreement in a group, but letting the group work through these conflicts is often key to giving them the motivation and insight that they need to develop creative, well-coordinated roles and norms. However, as a corollary, when factors such as existing power hierarchies or social group stereotypes determine how those conflicts are resolved, long-term inequalities may be baked into the resulting roles (e.g., O’Connor, 2019).

With respect to the CARMI framework, the Battle of the Exes is particularly useful for illuminating the relationship between timescales in Multi-level planning and Adaptation. Although participants were generally not thinking about this across-round strategy prior to their interactions, when they stumbled across it during early rounds, they recognized its value and worked to make it persist despite occasional disruptions. If participants had only been thinking in terms of within-round strategy, then this alternation strategy would not have been nearly as stable as it was.

Battle of the Exes also provides a good example of multiple, interacting temporal scales of strategies. Within-round strategies of generosity, bluffing, and chicken in the dynamic condition led to across-round strategies such as alternation and grim reaper (“If you don’t alternate with me, I’m going to always choose the high payoff option”) (Fudenberg & Tirole, 1991). Within-round conflict turned out to be instrumental in achieving across-round complementary coordination. The importance of dynamic, real-time interaction for establishing complementary roles (e.g., peel-off

times within each round) becomes apparent when we focus on the local processes by which roles emerge rather than their stable final form (e.g., stability across rounds). Local social interaction processes like imitation, preemption, argumentation, negotiation, proposing, straw polling, and voting all help groups reach effective role specializations over longer timescales. Dyads of actively interacting people find creative role specializations well beyond those created when individuals are restricted to producer and receiver roles (Theisen, Oberlander, & Garrod, 2010).

c. Find the Unicorn

In the Battle of the Exes game, there were two primary strategies for developing complementary roles that served as a major focal point attractor (Binmore & Samuelson, 2006), namely temporal or spatial differentiation. In addition, on any single round, player actions collapsed to only two outcomes – choosing the high or low payoff option. These simplifications made the patterns of dyadic interaction robustly analyzable. However, in our final paradigm for exploring role specialization (Andrade-Lotero & Goldstone, 2021), we considered a situation with a more open-ended space of strategies and with multiple, incompatible focal attractors. If players still adopted complementary roles when many role decompositions are possible, it would shed light on the way partners are mutually shaping one another’s strategies and the mechanisms players used to achieve such effective coordination.

We devised a task called Find the Unicorn, a two-player game in which players search for a unicorn in an 8x8 grid of tiles. On some trials, we hid a unicorn below a randomly chosen tile; on other trials, we removed the unicorn entirely. The two players simultaneously uncovered tiles, revealing whether the unicorn lies beneath them. Critically, we withhold information from the players. A player cannot see which tiles their teammate has uncovered, or whether a unicorn was found. If both players uncover the same tile, the overlap is signaled to both players by changing its color. At any point during a given round, either player can make a guess about whether the unicorn is present in or absent from the current grid. The other player sees the guess and may use it to inform their own guess. The round ends when both players announce that their guess is a final decision, at which point they receive feedback about their individual scores. The score depends on whether the player’s guess is correct, subtracting the number of tiles that were uncovered by both players, thus incentivizing fast and accurate completion at the group-level. This paradigm is inspired by

collaborative search tasks (Brennan, Chen, Dickinson, Neider, Zelinsky, 2008) where participants must efficiently divide the visual field, such as two bodyguards splitting the range of vision to scan for potential threats, or two children looking at a picture book (e.g., *Where's Waldo?*). It may also be interpreted in terms of a more abstract 'division of labor' by assuming that each tile represents some degree of effort.

Most dyads eventually settled on a strategy where players jointly uncovered the full grid without redundantly uncovering the same tiles, effectively partitioning the grid. In the top row of panel (A) in Figure 5, we show the 4 classes of coordinating strategies that were observed, and the bottom row shows unsuccessful dyads. To measure how well a dyad of players self-organized into complementary roles (i.e., sub-regions of the grid) we devised a Division of Labor Index (DL index), which operationalizes the idea that perfect coordination (a high DL Index value) is achieved when both players uncover the entire grid and do not overlap at all. The DL index significantly increased over successive rounds of the game, as members of a dyad learned to search non-overlapping, complementary regions of the grid. The DL index was positively related to the two players' consistencies, operationalizing consistency as the percentage of tiles that one player treats the same, either turning over or not, over two successive rounds. When there was substantial overlap in the tiles across players, performance on subsequent trials improved when the two players changed their selected tiles to differing degrees. That is, better coordination was observed when one player re-selected the same tiles (stubbornness), while the other uncovered a different set of tiles from the previous round (flexibility). Similar to the behavior of effective groups in the Group Binary Search game, we observed that considerable tile overlap gave rise to complementary degrees of reactivity, with one player being relatively stubborn and the other player being flexible. In panel (B) in Figure 5 we can understand why this difference in degrees of reactivity led to successful coordination. In this case, Player A stubbornly followed the "Choose Right Side" strategy, while Player B reacted to the overlapping tiles and decided to uncover only the tiles on the left side on the subsequent round.

There are two important determinants of reactivity. The first one is a graded version of a Win-Stay/Lose-Shift strategy (Nowak & Sigmund, 1993; Sloman, Goldstone, & Gonzalez, 2021) in which players prefer to keep using the previous round's strategy if it has been high scoring. Once a dyad has devised an effective tile divvying strategy, they decrease their reactivity to lock it in place.

Secondly, players were less reactive as the set of tiles they uncovered more closely resembled one of an inventory of possible coordinating strategies. For example, as the set of tiles uncovered by a player increasingly resemble those expected if they were adopting a “Choose Right Side” strategy, then the player is less reactive on the subsequent round. If their tile selections are consistent with their half of a fair strategy, then they tend to “lock in” that strategy, allowing their partner to eventually take up the complementary strategy.

We developed a series of nested agent-based models to explore both the fit of the models to the empirical results and how efficient they each are to achieve coordination. Our baseline model assumed that agents’ shared knowledge of the focal points (Schelling, 1960) indicated by the four pairs of strategies outlined in the top row in panel (A) in Figure 5. This model could account for some sets of tiles being selected far more often than predicted by chance, but could not account for wide early variation in strategies followed by gradual settling down into a single stable strategy. Thus, we extended the baseline focal points model to incorporate a graded win-stay-lose-shift strategy. While fitting human results better as well as performing better than the baseline model, this model missed certain patterns, such as predicting which player would tend to shift their strategy when two players overlapped considerably in their selected tiles. Consequently, a third model was devised, in which as an agent’s strategy becomes more similar to a focal point, it becomes increasingly likely to keep its strategy and its partner agent becomes more likely to shift its strategy to the complementary strategy of that focal point. This third model fit humans significantly better than the second model and also performed better in the sense of its dyads attaining higher scores.

Find the Unicorn brings together multiple insights from the CARMI framework. Successful players Adapted to the feedback received via scores and overlapping tiles. They were Repelled away from their partner’s strategy by inefficient overlaps, and had to construct a Multi-level plan. This plan potentially involved the longer timescale consideration that being closer to a coordinating strategy would help their partner recognize their Intentions and lead them to divide the grid in the expected way.

4. General Discussion

The three experimental paradigms reviewed above point to the importance of group members specializing to solve their problems. They also point to some of the factors that determine whether

the members will, in fact, be able to specialize. Experiment-long experience of members with each other promotes specialization, as does providing moment-to-moment information about members' strategies to all members. Human participants performed better at both individual and group levels when they adopted complementary roles, and the computational models that fit empirical results the best also incorporate mechanisms that lead to agents developing complementary roles.

4.1. Role Specialization in Biology and Artificial Intelligence

The benefits of role specialization within groups have been observed across a wide range of contexts, from eusocial insects to robot swarms. For insects, intrinsic variation in the ability of different castes to perform colony tasks can result in clear divisions of labor and resulting efficiencies (Page & Mitchell, 1991). Additionally, many species have mechanisms for individuals acquiring specializations based on environmental contingencies such as what food they are given to eat, temperature, predation risk, resource availability, or, as in our experiments, information about the tasks being performed by others (Duarte, Weissing, Pen, & Keller, 2011). Mammals such as bottlenose dolphins hunt in groups where individuals play highly delineated and contextually acquired roles such as drivers of prey fish and barriers to fishes' escape (Gazda, Connor, Edgar, & Cox, 2005).

Role specialization is especially critical for developing autonomous artificial agents who can coordinate effectively with human teams (Li, Kwon & Sadigh, 2021). Consequently, a number of classical multi-agent coordination algorithms involve explicit planning over role assignments (Grosz & Kraus, 1996; Kinny et al., 1992). Contemporary approaches to multi-agent reinforcement learning formalize roles as continuous embeddings that guide each agent's downstream policy (e.g. Wang et al., 2020, 2021). A number of insights emerge from studying these artificial systems. In one study by Ferrante et al. (2015), task partitioning emerged was found among evolved computational agents even when the agents were initially identical, and was favored when there were task switching costs. Task partitioning occurs even when strategy selection is based only on overall group performance with no prior information on possible subtask decompositions. Another factor that promotes role specialization among robots is frequency dependent profit margins in which agents assess the profitability of actions contingent on other agents' actions (Dadvar, Moazami, Myler, & Zargarzadeh, 2021). Interestingly, much of the research on swarm robotics has emphasized

4.2. Developmental trajectories for role specialization capacities

The emergence of differentiated roles has been identified as a key factor driving the evolution of group traits (Smaldino, 2014; O'Connor, 2019) with a strong evolutionary pressure for individual skills that promote the formation and maintenance of groups (Caporael, Dawes, Orbell, & Van de Kragt, 1989). Indeed, the underlying ability to differentiate roles and accommodate one's partner to pursue joint goals emerges early in development (see Tomasello & Hamann, 2012; Warneken et al., 2014; Meyer et al., 2016). Collaboration motivates children to take on harder tasks (Butler & Walton, 2013). By 4-5 years old they are able to reason about differences in knowledge and appropriately divide up roles according to skill levels (Magid et al., 2018), although full competence may emerge later in other (non-physical) task domains (Baer & Odic, 2022). Further exploration of developmental trajectories associated with each element of the CARMI framework may help to elucidate the specific cognitive mechanisms upon which effective collaboration is built. For example, early heuristics driving rigid 'repulsion' dynamics may give way to more flexible context-dependent inferences about others' mental states and plans as these capacities come online.

4.3. Consequences for real-world groups

The CARMI framework suggests several directions for promoting beneficial role specialization in real-world groups. First, facilitating communication among individuals in a group is often the most reliable way to decompose a group's task into subtasks. Over time, a group will establish the appropriate conventions and vocabulary to describe the relevant roles for the problems it needs to solve (Hawkins, Goodman, & Goldstone, 2019), and infrastructural technologies can dramatically facilitate communication. Second, an awareness of the need for role specialization may allow groups to "pre-differentiate" members. For example, a team can be explicitly assembled so that its members have diverse skills, backgrounds, and experiences (Page, 2007).

Third, division of labor is often most effective when members do not treat each other as complete black boxes, but rather have some understanding of each other's tasks, skills, and constraints (Faraj & Sproull, 2000). In each of our paradigms, modeling others' intentions was essential to role specialization. For Group Binary Search, individuals appear to learn that others are more/less reactive and then accommodate accordingly. In Battle of the Exes, the dynamic condition allowed some signaling of intentions, but even in the ballistic condition, dyads were able to coordinate to focal strategies across rounds. In Find the Unicorn, there were multiple focal strategies, but players

were often successful across games by mutually recognizing their closeness to a particular focal strategy and then making adjustments to pursue that strategy more efficiently. Even if individuals do not fully understand what experts on their team are doing, they may nonetheless come to rely on experts to achieve their own goals, and relying on experts itself becomes a norm that benefits the entire group (Andrade-Lotero, Velasco-García, Ortiz-Duque & Goldstone, 2022). Given these results, a recommendation is for groups to consider interventions that help members correct misconceptions and refine internal models of each other, such as task rotation, job shadowing, and cross-cutting activities bringing together multiple teams.

Conversely, understanding the processes underlying emergent role differentiation may shed light on the systematic forces that lead inequitable divisions of labor (and divisions of resources) to form and persist. For example, Cailin O'Connor (2019) shows in a simple computational model how roles easily sort along social categories (like gender) in ways that enforce long-run inequity and discrimination. For example, imagine a group dance where men have the role of stepping forward first, and women have the role of stepping back (p. 47). Locally, it becomes easy to avoid stepping on the toes of future partners of the opposite gender. But once role-based expectations are entrenched, they are highly robust to deviations (see also Bai et al., 2022). Similar mechanisms may drive rich-get-richer feedback loops in status hierarchies (Ridgeway, 2019), as when groups preferentially allocate roles (e.g. the 'leader' role) to individuals with type-consistent attributes. Thus, role specialization is best understood as a double-edged sword, and an important direction for future work is to better understand the forces driving inequalities to counteract harmful outcomes.

4. Conclusions

The CARMI framework identifies key factors in the bottom-up emergence of specialized roles within groups. The final row of Table 1 compares the three reviewed paradigms in terms of the CARMI components that were strongly implicated. All three paradigms crucially involved collective adaptation (A) of a group based on feedback supplied to it, and repulsion (R) of individuals from each other's strategies. In "Group Binary Search," these two processes suffice to yield agents with specialized roles that conjointly react appropriately to environmental signals. To A and R, "Battle of Exes" crucially adds multi-level (M) plans in which agents are not simply trying to maximize their payoffs on any given round, but are also trying to find stable, fair, and efficient strategies that work well in the long run. This gives rise to turn-taking solutions missed by agents who focus only on the

scale of a single game round. Finally, “Find the Unicorn” provides evidence for agents also developing models of each other’s intentions (I). Agents that infer goals and intentions from their partners’ actions help their group by accurately predicting how their partners will behave in novel situations, and successfully coordinating with them even in the presence of noise and environmental perturbations. Finally, direct communication (C) is notable in our three paradigms only because of its conspicuous absence. We stand by our decision to minimize communication in our experimental “parables” of collective coordination because communication would have likely led to ceiling levels of coordination. However, this is also an admission that direct communication of plans is often the single most potent tool of collective coordination (Vorobeychik, Joveski, & Yu, 2017).

In our experimental paradigms, we operationalized the notion of an individual’s role within a group as consistent behavior within the individual (self-similarity) and systematically different behaviors across individuals (other-dissimilarity). Roles, thus defined, emerged as a result of individuals’ tendencies to differentiate themselves from their partners, dynamically react to their partners’ behaviors, model their partners’ knowledge and intentions, respond to feedback, and develop multi-level strategies. For most real-world groups, these abilities are supplemented by explicit communication between members about how the group’s task should be decomposed into subtasks. Role specialization through explicit communication and enforcement of agreements may seem at odds with the self-organized emergence of roles emphasized by our work and others’. However, we would argue that a decentralized, self-organizing approach to the structured division of labor is still highly relevant even when communication and contracts are involved. One of the main activities that decentralized networks of agents do to promote their prosperity is build systems of rules, contracts, communication protocols, and organizations (Ostrom, 1990). Sophisticated communication systems and binding contracts are themselves the result of a complex, self-organized interplay between contracts, monitors, courts, constitutions, sanctions, channels, routing rules, and protocols. Our highly simplified role specialization paradigms provide a close-up view of some beginning steps that decentralized, disorganized groups take on the way to becoming more organized.

Acknowledgements

We are grateful for insightful comments from Marina Dubova, Seth Frey, Aileen Oeberst, Matt Setzler, Steven Sloman, Georg Theiner, Natalia Velez, and two anonymous reviewers.

5. References

- Allen, K., Jara-Ettinger, J., Gerstenberg, T., Kleiman-Weiner, M., & Tenenbaum, J. B. (2015). Go fishing! Responsibility judgments when cooperation breaks down. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, Austin, TX, 2015 (pp. 84-89). Cognitive Science Society.
- Andrade-Lotero, E., & Goldstone, R. L. (2021). Self-organized division of cognitive labor. *PLoS ONE*, 16(7): e0254532. <https://doi.org/10.1371/journal.pone.0254532>
- Andrade-Lotero, E., Velasco-García, J.A., Ortiz-Duque, J.M. & Goldstone, R. L. (2022). The Division of Linguistic Labor for Offloading Conceptual Understanding. *Philosophical Transactions of the Royal Society B*, 378, 20210360.
- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological monographs: General and applied*, 70(9), 1.
- Baer, C., & Odic, D. (2022). Mini managers: Children strategically divide cognitive labor among collaborators, but with a self-serving bias. *Child Development*, 93(2), 437-450.
- Bai, X., Fiske, S. T., & Griffiths, T. L. (2022). Globally Inaccurate Stereotypes Can Result from Locally Adaptive Exploration. *Psychological Science*, 09567976211045929.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 0064. <https://doi.org/10.1038/s41562-017-0064>
- Binmore, K. (2013). *Natural justice*. Oxford University Press.
- Binmore, K., & Samuelson, L. (2006). The evolution of focal points. *Games and Economic Behavior*, 55(1), 21-42.
- Bowling, M., & McCracken, P. (2005). Coordination and adaptation in impromptu teams. In *Proceedings of the 20th national conference on Artificial intelligence* (pp. 53-58).
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327-341.
- Brennan SE, Chen X, Dickinson CA, Neider MB, Zelinsky GJ. (2008). Coordinating Cognition: The Costs and Benefits of Shared Gaze During Collaborative Search. *Cognition*, 106(3): 1465–1477. doi:10.1016/j.cognition.2007.05.012.
- Butler, L. P., & Walton, G. M. (2013). The opportunity to collaborate increases preschoolers' motivation for challenging tasks. *Journal of Experimental Child Psychology*, 116(4), 953-961.

- Campbell, C., Izquierdo, E. & Goldstone, R. L. (2022). Partial copying and the role of diversity in social learning performance. *Collective Intelligence*, 1:1, 1-15.
- Caporael, L. R., Dawes, R. M., Orbell, J. M. & Van de Kragt, A. J. C. (1989) Selfishness examined: Cooperation in the absence of egoistic incentives. *Behavioral and Brain Sciences* 12:683–739.
- Chartrand, T. L., & Lakin, J. L. (2013). The Antecedents and Consequences of Human Behavioral Mimicry. *Annual Review of Psychology*, 64(1), 285–308. <https://doi.org/10.1146/annurev-psych-113011-143754>
- Csibra, G. and Gergely, G. (2007). ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60–78.
- Dadvar, M., Moazami, S., Myler, H. R., & Zargarzadeh, H. (2021). Exploration and Coordination of Complementary Multirobot Teams in a Hunter-and-Gatherer Scenario. *Complexity*, <https://doi.org/10.1155/2021/9087250>
- Dafoulas, G. A., & Macaulay, L. A. (2001). Facilitating group formation and role allocation in software engineering groups. In *Proceedings ACS/IEEE International Conference on Computer Systems and Applications* (pp. 352-359). IEEE.
- Dong, W., Lepri, B., Pianesi, F., & Pentland, A. (2012). Modeling functional roles dynamics in small group interactions. *IEEE transactions on multimedia*, 15(1), 83-95.
- Duarte, A., Weissing, F. J., Pen, I., & Keller, L. (2011). An Evolutionary Perspective on Self-Organized Division of Labor in Social Insects. *Annual Review of Ecology, Evolution, and Systematics*, 42(1), 91–110. <https://doi.org/10.1146/annurev-ecolsys-102710-145017>
- Durkheim, E. (1893). *The Division of Labor in Society*. New York: Free Press.
- Faraj, S. & Sproull, L. (2000) Coordinating Expertise in Software Development Teams. *Management Science*, 46(12):1554-1568.
- Ferrante, E., Turgut, A. E., Duéñez-Guzmán, E., Dorigo, M., & Wenseleers, T. (2015). Evolution of Self-Organized Task Specialization in Robot Swarms. *PLOS Computational Biology*, 11(8), e1004273. <https://doi.org/10.1371/journal.pcbi.1004273>
- Freire I. T., Moulin-Frier C., Sanchez-Fibla M., Arsiwalla X. D., Verschure P. F. M. J. (2020) Modeling the formation of social conventions from embodied real-time interactions. *PLoS ONE* 15(6): e0234434.
- Frey, S., & Goldstone, R. L. (2013). Cyclic game dynamics driven by iterated reasoning. *PLoS ONE*, 8(2): e56416. doi:10.1371/journal.pone.0056416
- Fudenberg, D., & Tirole, J. (1991). *Game Theory*. MIT Press.

- Galesic, M., Barkoczi, D., & Katsikopoulos, K. (2018). Smaller crowds outperform larger crowds and individuals in realistic task conditions. *Decision*, 5(1), 1–15. <https://doi.org/10.1037/dec0000059>.
- Gazda, S. K., Connor, R. C., Edgar, R. K., Cox, F. (2005). A division of labour with role specialization in group-hunting bottlenose dolphins (*Tursiops truncatus*) off Cedar Key, Florida. *Proc. R. Soc. B*, 272, 135–140.
- Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, 11, 508-514.
- Goldstone, R. L. & Theiner, G. (2017). The multiple, interacting levels of cognitive systems (MILCS) perspective on group cognition, *Philosophical Psychology*, 30:3, 334-368, DOI: 10.1080/09515089.2017.1295635
- Grosz, B. J., & Kraus, S. (1996). Collaborative plans for complex group action. *Artificial Intelligence*, 86(2), 269-357.
- Hawkins, R. D., & Goldstone, R. L. (2016). The formation of social conventions in real-time environments. *PLoS One*, 11(3): e0151670.
- Hawkins, R. D., Goodman, N. D., & Goldstone, R. L. (2019). The emergence of social norms and conventions. *Trends in Cognitive Science*, 23, 158-169.
- Hawkins, R. D., Franke, M., Frank, M. C., Goldberg, A. E., Smith, K., Griffiths, T. L., and Goodman, N. D. (2022). From partners to populations: A hierarchical bayesian account of coordination and convention. *Psychological Review*.
- Helmreich, R. L., & Schaefer, H. G. (1994). Team performance in the operating room. In M. S. Bogner (Ed.), *Human error in medicine* (pp. 225-254).
- Henrich, J. (2016). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.
- Ho, M. K., Cushman, F. A., Littman, M. L., and Austerweil, J. L. (2021). Communication in action: Planning and interpreting communicative demonstrations. *Journal of experimental psychology. General*, 124(1), 60-78.
- Hoda, R., Noble, J., and Marshall, S. (2012). Self-organizing roles on agile software development teams. *IEEE Transactions on Software Engineering*, 39(3), 422–444.
- Hove, M. J., & Risen, J. L. (2009). It's All in the Timing: Interpersonal Synchrony Increases Affiliation. *Social Cognition*, 27(6), 949-960.

- Janssen, M. A., Goldstone, R. L., Menczer, F., & Ostrom, E. (2008). Effect of rule choice in dynamic interactive spatial commons. *International Journal of the Commons*, 2, 288-312.
- Kinny, D., Sonenberg, E., Ljungberg, M., Tidhar, G., Rao, A., & Werner, E. (1992). Planned team activity. In *European Workshop on Modelling Autonomous Agents in a Multi-Agent World* (pp. 227-256). Springer, Berlin, Heidelberg.
- Lakin, J. L., Jefferis, V. E., Cheng, C. M., & Chartrand, T. L. (2003). The chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior*, 27, 145-162.
- Laland, K. N. (2017). *Darwin's Unfinished Symphony: How Culture Made the Human Mind*. Princeton University Press.
- Lee, W., Vaughan, N., & Kim, D. (2020). Task Allocation into a Foraging Task With a Series of Subtasks in Swarm Robotic System. *IEEE Access*, 8, 107549–107561. <https://doi.org/10.1109/ACCESS.2020.2999538>.
- Li, M., Kwon, M., & Sadigh, D. (2021). Influencing leading and following in human–robot teams. *Autonomous Robots*, 45(7), 959-978.
- Liemhetcharat, S., & Veloso, M. (2014). Weighted synergy graphs for effective team formation with heterogeneous ad hoc agents. *Artificial Intelligence*, 208, 41–65.
- Magid, R. W., DePascale, M., & Schulz, L. E. (2018). Four-and 5-year-olds infer differences in relative ability and appropriately allocate roles to achieve cooperative, competitive, and prosocial goals. *Open Mind*, 2(2), 72-85.
- McKee, K. R., Hughes, E., Zhu, T. O., Chadwick, M. J., Koster, R., Castaneda, A. G., Beattie, C., Graepel, T., Botvinick, M. & Leibo, J. Z. (2021). Deep reinforcement learning models the emergent dynamics of human cooperation. arXiv preprint arXiv:2103.04982.
- Meyer, M., van der Wel, R. P., & Hunnius, S. (2016). Planning my actions to accommodate yours: joint action development during early childhood. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1693), 20150371.
- Mirsky, R., Macke, W., Wang, A., Yedidsion, H., & Stone, P. (2021). A penny for your thoughts: the value of communication in ad hoc teamwork. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence* (pp. 254-260).
- Misyak, J. B., Melkonyan, T., Zeitoun, H., & Chater, N. (2014). Unwritten rules: virtual bargaining underpins social interaction, culture, and society. *Trends in cognitive sciences*, 18(10), 512-519.

- Nowak, M., Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* 364, 56–58 (1993). <https://doi.org/10.1038/364056a0>
- O'Connor, C. (2019). *The origins of unfairness: Social categories and cultural evolution*. Oxford University Press, USA.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Ostrom, E. (2006). The value-added of laboratory experiments for the study of institutions and common-pool resources. *Journal of Economic Behavior & Organization*, 61(2), 149-163.
- Ostrom, E.; Walker, J.; Gardner, R. (1994). *Rules, games, and common-pool resources*. Ann Arbor: University of Michigan Press.
- Page, R. E. & Mitchell, S. D. (1991) Self organization and adaptation in insect societies. In: Philosophy of Science Association 1990, Vol. 2, ed. A. Fine, M. Forbes & L. Wessels, pp. 289–98. Philosophy of Science Association.
- Page, S. E. (2007) *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton University Press.
- Pugh, D. S., Hickson, D. J., Hinings, C. R., and Turner, C. (1968). Dimensions of organization structure. *Administrative science quarterly*, 13(1), 65–105.
- Rabb, N., Fernbach, P. M., & Sloman, S. A. (2019). Individual Representation in a Community of Knowledge. *Trends in Cognitive Sciences*, 23(10), 891–902. <https://doi.org/10.1016/j.tics.2019.07.011>
- Raveendran, M., Puranam, P., & Warglien, M. (2016). Object salience in the division of labor: Experimental evidence. *Management Science*, 62(7), 2110-2128.
- Richardson, M. J., Harrison, S. J., Kallen, R. W., Walton, A., Eiler, B. A., Saltzman, E., & Schmidt, R. C. (2015). Self-organized complementary joint action: Behavioral dynamics of an interpersonal collision-avoidance task. *Journal of Experimental Psychology: Human Perception and Performance*, 41(3), 665.
- Ridgeway, C. L. (2019). *Status: Why is it everywhere? Why does it matter?*. Russell Sage Foundation.
- Roberts, M. E., & Goldstone, R. L. (2011). Adaptive Group Coordination and Role Differentiation. *PLoS ONE*, 6, 1-8.
- Schelling, T. C. (1960). *The strategy of conflict*. Cambridge: Harvard University Press
- Scott-Phillips, T. C., Kirby, S., & Ritchie, G. R. (2009). Signalling signalhood and the emergence of communication. *Cognition*, 113(2), 226-233.

- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in cognitive sciences*, 10(2), 70–76.
- Setzler, M., & Goldstone, R. L. (2020). Coordination and Consonance Between Interacting, Improvising Musicians. *Open Mind: Discoveries in Cognitive Science*, 4, 88–101. https://doi.org/10.1162/opmi_a_00036
- Setzler, M., & Goldstone, R. L. (2022). Tonal Emergence: An Agent-Based Model of Tonal Coordination. *Cognition*, 221, 1-19. 104968.
- Shannon C.E. (1948). *A mathematical theory of communication*. Bell system technical journal, 27. doi: 10.1002/j.1538-7305.1948.tb01338
- Sloman, S., & Fernbach, P. (2017). *The Knowledge Illusion: Why we never think alone*. Riverhead.
- Sloman, S. J., Goldstone, R. L., & Gonzalez, C. (2021). A Social Interpolation model of group problem-solving. *Cognitive Science*, 45, 1-30. E13066.
- Sloman, S. A., & Rabb, N. (2016). Your understanding is my understanding: Evidence for a community of knowledge. *Psychological Science*, 27(11), 1451–1460.
- Smaldino, P. E. (2014). The cultural evolution of emergent group-level traits. *Behavioral and Brain Sciences*, 37(3), 243–254. <https://doi.org/10.1017/S0140525X13001544>
- Smaldino, P. E. (2019). A Modeling Approach that Integrates Individual Behavior, Social Networks, and Cross-Cultural Variation. *Trends in Cognitive Sciences*, 23(10), 818–820. <https://doi.org/10.1016/j.tics.2019.07.005>
- Stone, P., & Veloso, M. (1999). Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2), 241–273. [https://doi.org/10.1016/S0004-3702\(99\)00025-9](https://doi.org/10.1016/S0004-3702(99)00025-9)
- Strogatz, S. H. (2003). *Sync: How order emerges from chaos in the universe, nature, and daily life*. Hyperion-Acquired Assets: New York.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. Cambridge, MA: MIT press.
- Theisen, C.A., Oberlander, J., & Garrod, S. (2010) Systematicity and arbitrariness in novel communication systems. *Interaction Studies*, 11, 14–32
- Tomasello, M., & Hamann, K. (2012). The 37th Sir Frederick Bartlett lecture: Collaboration in young children. *Quarterly Journal of Experimental Psychology*, 65(1), 1-12.
- Valentine, M. A., & Edmondson, A. C. (2015). Team scaffolds: How mesolevel structures enable role-based coordination in temporary groups. *Organization Science*, 26(2), 405-422.

- Vanderschraaf P., Skyrms B. (2003). Learning to take turns. *Erkenntnis*, 59(3), 311–347. doi: 10.1023/A:1026046625024
- Vorobeychik, Y., Joveski, Z., & Yu, S. (2017). Does communication help people coordinate? *PLoS ONE* 12(2): e0170780. doi:10.1371/journal.
- Wang, T., Dong, H., Lesser, V. R., and Zhang, C. (2020). ROMA: Multi-agent reinforcement learning with emergent roles. In *Proceedings of the International Conference on Machine Learning* (pp. 9876-9886).
- Wang, T., Gupta, T., Peng, B., Mahajan, A., Whiteson, S., & Zhang, C. (2021). RODE: Learning roles to decompose multi-agent tasks. In *Proceedings of the International Conference on Learning Representations*.
- Warneken, F., Steinwender, J., Hamann, K., & Tomasello, M. (2014). Young children's planning in a collaborative problem-solving task. *Cognitive Development*, 31, 48-58.
- Wegner D.M. (1987). Transactive memory: A contemporary analysis of the group mind. In: *Theories of group behavior*, ed. Brian Mullen & George R. Goethals, pp. 185-208.
- Wu, S. A., Wang, R. E., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., & Kleiman-Weiner, M. (2021). Too Many Cooks: Bayesian Inference for Coordinating Multi-Agent Collaboration. *Topics in Cognitive Science*, 13(2), 414-432.




Figure 1. An overview of the CARMI Framework for understanding how groups dynamically self-organize themselves so that members adopt complementary specialized roles.

Feature / Paradigm	GBS	Exes	Unicorn
Number of players	Many	2	2
Number of player outcomes	Medium (50) integers	Small (2) top/ bottom	Large (2^{64}) powerset of grid
Feedback coarseness	Manipulated (Exact vs. directional)	Exact	Partial (intersecting tiles only)
Real-timeness	No (end of round)	Manipulated	Yes
Goal	Collective	Individual	Individual
Stressed CARMi components	Adaptation Repulsion	Adaptation Repulsion Multi-level Plans	Adaptation Repulsion Multi-level Plans Intention-modeling

Table 1. An overview of dimensions along which the three experimental paradigms differ.




Figure 2. Results from Roberts and Goldstone’s (2011) Group Binary Search task. Groups improve in their ability to coordinate over successive plays of the game. Groups do better when given exact numeric feedback (e.g. “The group’s sum is too high by 23”) than directional feedback (e.g. “The group’s sum is too high”).




Figure 3. Participants' reactions to numeric feedback in the Group Binary Search game. The horizontal axis shows the feedback given to participants by the computer (e.g. -10 means that the group's sum was lower than the target number by 10). The vertical axis indicates how the group responded on the next round by increasing (positive numbers) or decreasing (negative numbers) their sums relative to the previous round. The dotted line shows the line expected if the group was perfectly calibrated to the feedback. The solid line shows the best fitting line to the empirical results.




Figure 4. Results from Hawkins and Goldstone (2016). The efficiency and fairness of dyads' solutions are higher when they can see each other's moment-to-moment positions. For stability, there was a strong interaction between payoff differential (high vs low) and the moment-to-moment visibility of players.

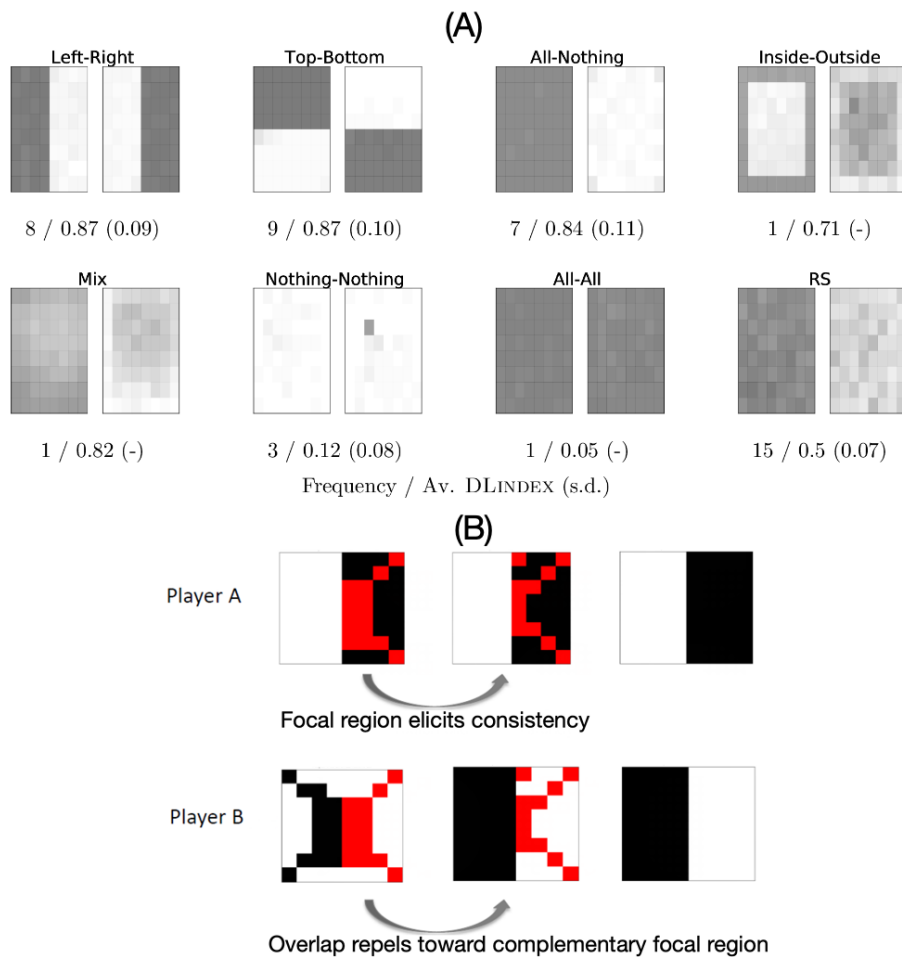


Figure 5. (A) Example decompositions of the 8 X 8 grid by dyads. The focal attractor strategies in the top row are efficient, complementary strategies in which the two players jointly uncovered all of the tiles without any overlap in their coverages. RS stands for a Random (or uncategorizable) Strategy. (B) An actual game play in which the two mechanisms of stubbornness and repulsion-to-the-complementary-strategy are present in the emergence of different, complementary roles.

