The unrealized value of centralization for coordination

Frederic Schneider, School of Management, Yale University

Eva Ranehill, Department of Economics, University of Gothenburg

Roberto Weber, Department of Economics, University of Zurich

January 18, 2018*

Preliminary draft. Please do not cite or circulate.

Abstract

We study the extent to which groups and individuals implement centralization to turn a seemingly complex coordination problem into a much simpler one. In a laboratory experiment, participants repeatedly form bilateral network ties and play a coordination game. Profitable interaction requires agreeing on a "convention," and participants benefit from developing (?) a homogenous convention to be used for all interactions. Homogenous conventions are most easily achieved under a centralized strategy in which a central participant "teaches" all others a common convention. We find little evidence that participants spontaneously identify and implement a centralized strategy, both in experiments using the above task with a student population, and in survey experiments using a stylized version of the task with a sample of experienced managers and game theorists. We conclude that individuals have substantial difficulty to recognize the benefits of centralization for aiding coordination and efficiency.

Key words: coordination, conventions, centralization experiment, network formation

JEL-Codes: C92, C72

* We thank participants at several seminars and conferences for helpful comments and suggestions. We also thank Nadja Jehli for valuable research assistance. We are also grateful for financial support from the Swiss National Science Foundation (Project 100018-140571).

The unrealized value of centralization for coordination

"The coordination of efforts essential to a system of cooperation requires, as we have seen, an organization system of communication. Such a system of communication implies centers or points of interconnection and can only operate as these centers are occupied by persons who are called executives."

- Chester Barnard, *The Function of the Executive* (1938)

1. Introduction

It has long been recognized that centralization and hierarchy can play a critical role in facilitating coordination. For example, in order to ensure the standardization of national languages, states have often established centralized coordinating entities, such as the Académie Française or the Real Academia Española. A primary function of these bodies is to coordinate the use of the French and Spanish languages across diverse populations, geographical regions and national boundaries. Similarly, the United States Constitution specifies a centralized coordinating function for the federal government, including giving Congress the power to "fix the standard of weights and measures."

Centralization is also an important mechanism for *organizational* coordination. As early as Barnard (1938), organizational scholars recognized that coordination of interdependent units in firms is often most easily achieved through hierarchical systems in which upper-level managers direct and align action. Chandler (1977) argued that the "visible hand" of management and hierarchy arose from the increased need for coordination produced by large firms with varied geographical locations, products and activities. Malone (1987) further noted that communication costs necessary for coordination are often minimized when information transmission occurs through a hierarchy. More recently, economic models of organization acknowledge a critical role for hierarchy and centralization in coordinating activity across units, while also recognizing the costs of centralization (Crémer, Garicano, and Prat 2007; Alonso, Dessein, and Matouschek 2008; Hart and Holmstrom 2010).

Implicit in these strands of research is the idea that people will utilize centralization when it is beneficial. In this study, rather than asking whether centralization can be profitable for organizations—or the conditions under which it is so— we ask to what extent people recognize that seemingly complex coordination problems can be greatly simplified through the adoption of hierarchy and centralization. Simply put, we study to what extent individuals recognize and implement centralization to tackle coordination problems.

To answer this question cleanly, we abstract from features of natural and theoretical environments (e.g., uncertainty, communication costs, misaligned incentives) that can make centralization inefficient. More precisely, we construct a situation in which coordination is profitable and then study whether people actually adopt a centralized strategy. We develop a group task in which individuals decide first with whom to engage in bilateral interactions and second which of many possible "conventions" to employ in their chosen interactions. The problem of selecting among the possible conventions is a pure-matching coordination game (Schelling 1960; Mehta, Starmer, and Sugden 1994), in which coordination on any common convention is efficient but where, *a priori*, no unique common convention exists. Interacting with another individual is efficient when subjects share a common convention, but not otherwise. We design the game so that interacting with others, and attempting to coordinate on a common convention is, despite an initial small cost, substantially more profitable than no interaction, in expectation. Additionally, the more rapidly a pair of interacting individuals coordinates on a convention the higher their payoffs. We also provide an added incentive to develop a homogeneous convention that works across all interactions with other group members. Efficiency is maximized in our experimental task when *(i)* individuals interact with all others in their group, *(ii)* quickly learn a convention for each interaction, and *(iii)* have access to a common convention that can be used by all group members. Given the possibilities of multiple independent interactions and many possible conventions in each interaction, solving this coordination problem can appear challenging.

However, adopting a rather simple centralized strategy greatly facilitates the coordination problem and allows full efficiency to obtain in a small number of periods. Under such a strategy, one group member "teaches" all others in the group a common convention so that everyone can engage in homogeneous interactions. Hence, centralization—i.e., selecting one group member to serve the same function as the Académie Française or the Real Academia Española—makes the coordination problem in our setting fairly trivial.

Our main research question is whether people confronted with a coordination problem of the above type adopt this centralized strategy. To address this question, we study the ability of groups and individuals to identify and implement the centralized strategy through a series of laboratory experiments and survey studies.

We begin with a laboratory experiment in which we confront groups with a challenging version of the task. In Study 1, implementing centralization requires both identifying such a strategy, but also solving—through tacit agreement—a higher-order coordination problem of who should serve the central function. This requires a manner in which to break symmetry among strategically identical group members.

Participants in our first study only have access to a subtle instrument to break the symmetry of group members, in the form of ordered ID numbers, such that one participant is labeled as “1,” another as “2,” and so on, with the number also reflecting the group member’s position on the screen on which subjects make their decisions. In such a case, group member “1” is a natural focal point for the central role.[1] Hence, subjects in our first study have a mechanism to break symmetry between the group members and implement the centralized approach that makes coordination simple. However, the subtlety of ordered integers as a tacit coordination device makes implementing the centralized strategy difficult. In accordance with our expectations, we find that groups fail to implement the centralized strategy in this environment. Hence, the ability to identify the benefits of centralization and self-organize is not so strong that this outcome spontaneously arises in challenging settings.

Our second and third studies attempt to greatly facilitate groups’ implementation of the centralized strategy, by introducing mechanisms present in organizations that potentially greatly reduce the problem of identifying who should serve the central, teaching role. Our goal is to provide empirical evidence on whether these mechanisms do, in fact, serve this function and facilitate the adoption of efficient centralized strategies.

In Study 2, we impose structure on the group interactions. Specifically, we impose two rules governing the interactions between individuals. The first rule—*Positional Structure*—requires that Group members with labels 2-6 interact with the group member labeled with the ID number “1” in all periods, while other interactions remain voluntary. Hence, this restriction solves the problem of which participant should serve the central coordinating role. This is analogous to having a person in an organization with whom everyone must regularly meet or who regularly communicates with everyone.

The second rule—*Action Structure*—requires the person with ID 1 to use a single, homogeneous, convention in all of his or her interactions with the other five group members. This means that everyone interacting with group member 1 learns the same convention, which is then helpful for interacting for other group members.

[1] Mehta, Starmer and Sugden (1994) demonstrate that “1” is the modal choice in a pure matching coordination game involving picking “any positive number.” This concords with Schelling’s (1960) intuition that “1” is unique as the “‘first’ or ‘smallest’ number.”

The combination of both rules (*Combined Structure*) greatly simplifies the coordination problem by requiring everyone to interact with group member 1, and to do so using the common convention employed by group member 1. At that point, the coordination problem between all other group members is trivial—they only need to start interacting with one another using this common convention. Since the combination of these two rules seemed, *a priori*, to make the remaining coordination problem almost trivial, we thought this study would demonstrate that these kinds of rules are valuable means for solving coordination problems.[2]

Despite this great simplification, we find that imposing structure has limited benefits for groups' ability to solve the coordination problem. Aside from what the rules force participants to do, they realize few further benefits. Specifically, participants 2-6 fail to realizets the benefits produced by the forced interactions with group member 1 in their interactions, despite thatthe forced structured interactions guarantee that they learn a common convention. Hence, despite our prior belief that the structure we imposed would make transparent and easily attainable the benefits of centralization, we find very little evidence that participants recognize its value.

Study 3 addresses the problem that, so far, coordination is tacit, which may prevent groups from realizing the benefits of centralized coordination. We therefore implement a *Communication* condition where subjects can coordinate explicitly on centralization. Specifically, we let group members spend five minutes at the beginning of the experiment discussing how to approach the task. This collective reasoning process makes it more likely that the group discovers the centralized solution and provides them with the opportunity to remove strategic uncertainty by explicitly appointing one group member as the centralized teacher, and to specifying how each player should act under a centralized strategy. Indeed, we find some groups clearly identify such a strategy and agree to implement it. These groups do well in terms of rapidly developing a common convention and obtaining high earnings. However, this only applies to about one-quarter of the groups. The remaining groups do no better than the groups in Study 1. Hence, despite giving groups ample opportunity to self-organize, we find that a large majority of groups fails to do so.

In combination, Studies 2 and 3 demonstrate that most groups fail to realize the benefits of centralization even when doing so is substantially simpler than in Study 1. We interpret this as evidence that the benefits of centralized organization for solving coordination problems are

---

[2] We also implemented only Position Structure and onlyAction Structure as separate conditions as a way to identify the relative importance of these two kinds of interventions.

not readily apparent to many people. This interpretation of our results is consistent with the observations that even subjects in Study 2 who have been taught a homogenous convention through their mandatory centralized interactions fail to recognize that this easily allows them to interact with all other group members. Moreover, we also find that those groups that *are* able to agree on the centralized strategy with Communication have little trouble implementing it. Thus, it seems to be that the limits arise from failing to understand that centralization makes the complex coordination problem a simple one.

Our results should also be interpreted in light of the fact that our study involves an environment designed to eliminate many of the reasons why groups and individuals may avoid centralizing in more natural environments. For instance, the task in our experiment is one in which decentralization does not confer groups with better access to, and use of, dispersed information, which may limit the benefits of centralization in real organizational contexts (Milgrom and Roberts 1992; Alonso, Dessein, and Matouschek 2008). Thus, in our study, unlike in many non-laboratory contexts, there is little reason not to centralize.

Nevertheless, Studies 1 through 3 do not rule out the possibility that groups fail to implement the centralized strategy due to some other reason than simply being unaware of its benefits. Therefore, our fourth and fifth studies attempt to isolate the role of awareness of the benefits of centralization from other preference-based or strategic reasons not to implement such a strategy.

Study 4 provides groups, at the beginning of the experiment, with a description of the centralized strategy. That is, this version of the experiment was identical to Study 1, except that, at the outset, we explained the strategy whereby group member 1 teaches everyone a common convention and the remaining players then use this homogenous convention to coordinate with one another. This recommendation also briefly described the payoff benefits of pursuing this strategy. However, as in Study 1, we left it entirely up to subjects to decide how to interact with one another and what choices to make in the coordination game. Nevertheless, *every* group in Study 4 implemented the centralized strategy. A simple explanation of the benefits of centralization for coordination is thus sufficient to get people to adopt it in the context of our experiment.

Study 5 employs a different approach, moving our focus from a strategic setting—where coordinating through the centralized strategy involves sufficient confidence that others will also pursue this strategy—into a non-strategic decision. We modify the laboratory experiment into a scenario-based individual task in which participants have to provide a recommendation to a hypothetical work team confronting a stylized version of our laboratory experiment. Specifically, we describe a task in which individuals have to decide with which other team members to collaborate and which software to use in their collaboration. As in our laboratory task, the hypothetical work team do well by coordinating on a centralized strategy in which a single team member teaches all other group members which software to use, which then rapidly facilitates the adoption of a homogeneous software choice among all team members. We allow subjects to provide written strategies for how a team should confront this task and we provide substantial monetary incentives for providing an effective recommendation. Hence, this study provides a test of whether individual subjects understand the value of centralization and can describe a centralized strategy in their recommendations.

Our first wave of data collection uses the same student population as our laboratory experiments and finds that only 29 percent of subjects recommend the centralized strategy. Hence, even in an individual task without strategic uncertainty, where subjects do not need to rely on others to realize the benefits of centralization, most people fail to recognize centralized strategies as a means for greatly simplifying a complex coordination problem.

The survey-based method also allows us to test whether a population with greater managerial experience more easily identifies the centralized strategy. One natural concern about our use of a student population in the above studies is that this population is not responsible for a great deal of organizational decision making. Hence, if more experienced managers regularly recognize the benefits of centralization for coordination, perhaps it is less relevant that less experienced students do not. For this purpose, we propose the same individual task to a sample of Executive MBA students, with substantial work experience. In this sample, the frequency of recommendations to play the centralized strategy is 32 percent, which is statistically indistinguishable from the proportion in the student sample. Thus, the inability to identify centralization as a solution to coordination problems is not unique to student samples.

Taken together, our studies suggest that, for many people, centralization may not readily come to mind as a way to tackle complex coordination problems, even when it greatly facilitates such problems and yields considerably higher earnings. Nevertheless, we do find that roughly a quarter of people and groups identify and adopt centralized strategies, meaning that it is not simply too cognitively demanding to do so. Instead, we believe that our findings are consistent with a form of "organizational naiveté" that fails to recognize the value of simple ways of organizing work tasks or communication practices to help facilitate coordination (Heath and Staudenmayer 2000; Kriss, Blume, and Weber 2016).

The remainder of this paper is structured as follows. In the next section, we briefly review several strands of related literature. Section 3 describes the coordination game we study and identifies theoretical properties of different strategies, including centralization. Sections 4 through 9 present, respectively, Studies 1 through 5. Section 10 concludes and discusses some implications of our results.

2. Related literature

Our work shares features with many research streams at the intersection of experimental and organizational economics. First, as we note in the Introduction, our work follows numerous theoretical and empirical studies in economics and organization that note potential benefits of centralization for coordination (Malone 1987; Alonso, Dessein, and Matouschek 2008; Hart and Holmstrom 2010).

Our work is most closely related to research highlighting the idea that people fail to employ mechanisms that aid coordination. Heath and Staudenmayer (2000) introduce the term "coordination neglect" to refer to the tendency to "fail to understand that coordination is important" and to "fail to take steps to minimize the difficulty of coordination" (p. 155). Heath and Staudenmayer present evidence from several business cases and laboratory studies demonstrating an apparent neglect of simple coordinating processes and identify cognitive mechanisms that may underlie this phenomenon. More recent laboratory studies in economics further demonstrate that simple mechanisms that can aid coordination—such as pre-play communication—may be underutilized (Kriss, Blume and Weber, 2016; Fehr 2016).

Also related is an extensive experimental literature that uses simple games to model the kinds of coordination problems that arise in organizations. This research often identifies ways in which groups of individuals fail to obtain efficient equilibria and then tests the effectiveness of alternative organizational interventions for facilitating efficient coordination (Weber 2006; Brandts and Cooper 2006; Brandts, Cooper, and Weber 2014). Our studies of mechanisms such as organizational structure (Study 2) and communication (Study 3) as means for aiding coordination relate to this work.

Somewhat related, though not explicitly employing coordination games, is research that studies how groups of individuals learn in contexts where learning must occur in a coordinated manner. For instance, Leavitt (1951) studied five-person groups attempting to identify a symbol jointly observed by all five group members by passing messages back and forth. Leavitt studied different communication networks, identifying their relative benefits and drawbacks in terms of efficiency of learning (errors and speed) and group-member satisfaction. More centralized structures generally yielded better performance, though also lower satisfaction. In other related work, Weber and Camerer (2003) and Selten and Warglien (2007) study the development of simple conventions and codes in laboratory groups, using procedures that bear some similarity to the task in our experiments. Also relevant is a study by Blume, Duffy, and Franco (2009) that studies a joint search process, whereby group members have to identify a combination of individual actions that yields a successful outcome. Their study restricts such learning to decentralized (symmetric) strategies, finding evidence that subjects often adopt optimal decentralized search, which is also consistent with our findings that behavior is often consistent with a symmetric, decentralized equilibrium.

Less closely related, but relevant for our work are other laboratory experiments on games played on endogenous networks. Like in this work, our experiment creates a situation in which people voluntarily form links, and these links then determine play of a subsequent strategic game (Kirchkamp and Nagel 2007; Corbae and Duffy 2008; Choi and Lee 2014). However, unlike in these studies, our focus is not on the particular outcomes that arise in any specific instance of the coordination game, but rather in whether a specific general type of super-game strategy arises across a finitely repeated game.

3. The Game

The main part of our laboratory experiment involves a two-stage game that is repeated for 30 periods among a fixed group of 6 players. The game has complete information, and the interests of the players are aligned.

In each period of the game, players first decide whether or not to interact with each other; an interaction occurs only if both players agree to it. If no interaction takes place, the two players receive a fixed outside payment of 30 Experimental Currency Units (ECU). If interaction takes place, the players enter the second stage and play a two-person coordination game with the objective of identifying a joint mapping from "symbols" to "objects." Specifically, subjects interacting in the second stage are shown three colored circles (red, blue and yellow), or objects, and can assign one unique abstract symbol from seven possibilities ("$", "!", "=", "%", "@", "?" and "#") to each object. If an interaction between two players takes place, the objective is to select the same symbol for each object as the other player. If two players assign the same symbol to an object, they have a match and receive a payoff of 20. Thus, players' payoffs in the second stage can range from 0 (if no matches occur) to 60 (if they match all three symbol-object choices). If two players interact, they learn at the end of the period whether they matched symbols for a particular object, but not what symbol the other player chose in case there is no match.

Thus, in this coordination game players have to figure out whether it is worthwhile to interact with the other five players and, if interacting, identify a common mapping of symbols to objects for each interaction. More precisely, assume that two identical players, 1 and 2, have chosen to interact with each other. They see a common set, O, of 3 objects and a common set, S, of 7 > 3 symbols. Let's call player i's assignment of symbols to objects in an interaction with player j, f_i^j.[3] If both players assign the same symbol to a particular object $o \in O$ (that is, $f_1^2(o) = f_2^1(o)$), they have a match, $m_{1,2}(o) = 1$. Call the total number of matches between two players, $M_{1,2} = \sum_{o \in O} m_{1,2}(o) \in \{0,1,2,3\}$. After choosing their functions, f_1^2 and f_2^1, the players observe for which objects they matched symbols and for which they did not, and their corresponding payoffs are realized. Players receive a payoff for each successful match in their interactions, $\pi_1 = \pi_2 = 20 M_{1,2}$. Thus, the players' goals are aligned. They both prefer achieving as many matches as quickly as possible and, if able to match, prefer interacting to not interacting. If both players match symbols for all 3 objects, $f_1^2(o) = f_2^1(o)\ \forall o \in O$, we call this equilibrium mapping a "convention."

Note that, in principle, a player can interact efficiently with each of the other group members while employing a different convention in each interaction. Our implementation of the game adds one further element, which mirrors the feature that having access to a unique common convention among members of a society or organization is often desirable. Specifically, we provide an incentive to identify a single, homogeneous symbol -object mapping that all group members can use to interact with each other. To do this, we append 5 additional periods (31-35) of the repeated game, where players are rewarded for having access to a common convention. In this second phase, we introduce two modifications to the game: *(i)* players are forced to interact with all five other group members, and *(ii)* they interact using a single mapping for all five interactions, $f_i^j = f_i\ \forall j \neq i$. That is, players need to coordinate with everybody using a single convention. Since players are aware of Phase 2 at the outset of the game, we expect that players recognize that they can benefit not only from developing conventions that allow efficient interaction with each other group member, but also from developing a homogeneous convention that works across all interactions.

Our primary interest is in the super-game strategies that players employ in this game, and the impact of these strategies on earnings across all 35 periods. There are many strategies

[3] That is, she chooses an injective and surjective mapping $f_i^j: O \to S$ between the two sets. The function is not bijective because there are more symbols than objects, leaving some symbols unassigned.

and equilibria in both the stage game and the super game. For instance, any network structure, including both the empty network (in which nobody interacts) and the complete network (in which all interactions take place), is supported by Nash equilibria, as is any combination of bilateral network links. When two players interact continuously, a natural strategy in the coordination game is a symmetric learning strategy in which players choose symbol-object pairings at random until matching; once they match for all three symbols, they continue using that convention in their bilateral interaction. Under this strategy, the probability that two players will initially match symbol choices on any single object is $\frac{1}{7}$ and the probability that they will match symbol choices on all 3 objects in the first attempt is $\frac{1}{7} \times \frac{1}{6} \times \frac{1}{5} = \frac{1}{210}$. In expectation, this strategy yields 0.43 matches in the first period. Since players will not undo existing matches, reaching 3 matches (i.e., a convention) is an absorbing state, which is approached asymptotically over time. In expectation, two continuously interacting players achieve the payoff from the outside option, corresponding to 1.5 matches, by period 5, and the ex ante probability of achieving all three matches by period 30 is 99.8%. In what follows, we refer to this random learning strategy as the *symmetric* learning strategy.

To provide broad insight into the benefits of centralization, we focus on three benchmark super-game equilibria: no interaction, interaction with symmetric strategies, and interaction with a centralized strategy. Figure 1 depicts the expected per-period payoffs from these strategies. While there are many possible alternative equilibrium strategies in the super-game, these three capture the relative benefits of decentralized and centralized approaches.

Figure 1. Expected per-person payoffs over time, under three benchmark strategies.

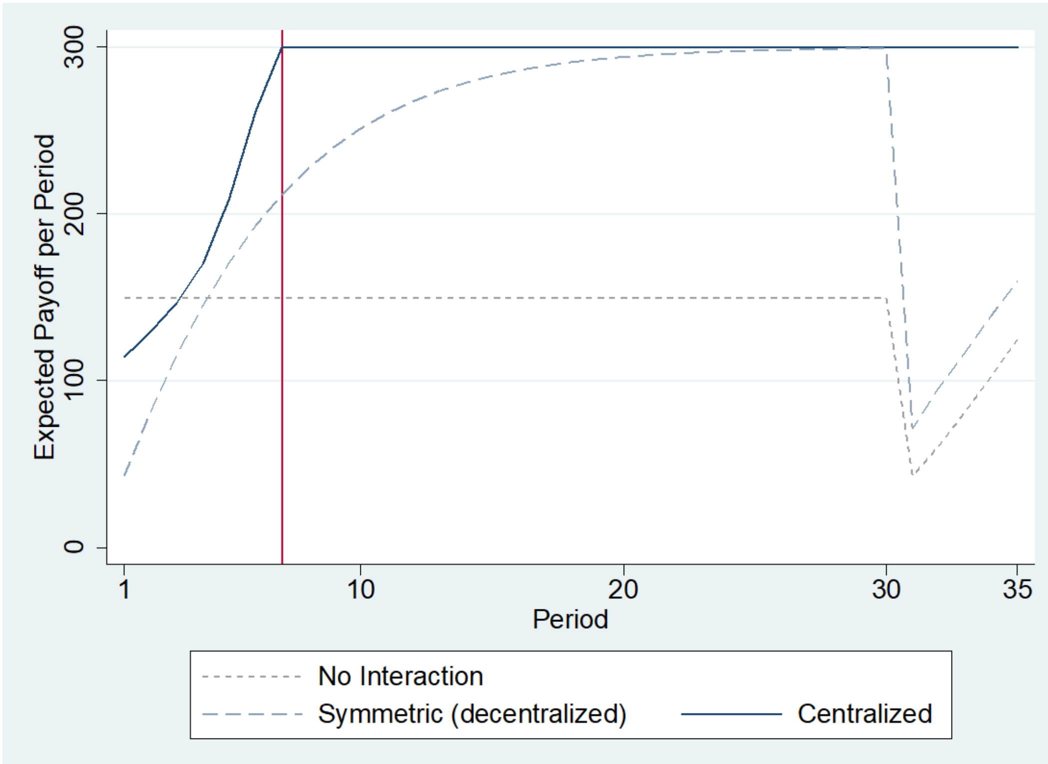

No interaction. An equilibrium exists in which players forgo interacting entirely during Phase 1 (Periods 1-30). This guarantees the fixed outside payoff of $30 \times 5 = 150$ in each period. The payoffs for this strategy are shown by the line labeled, "No Interaction," in Figure 1, which is flat at 150 for Phase 1.

In Phase 2 (periods 31-35), players are required to interact but cannot use a history of interactions from Phase 1 to aid coordination. Hence, in period 31 they start from scratch, with an initial expected match rate for each object of $\frac{1}{7}$. Subsequently, in periods 31-35 they might employ the symmetric learning strategy that we describe above—Figure 1 presents the expected Phase 2 earnings under such symmetric learning.[4]

Interaction with symmetric strategies. As an alternative, consider the equilibrium in which all players in a group choose to interact in all periods and adopt the symmetric learning strategy across all periods in Phase 1. Players start by choosing random symbol-to-object mappings in all their interactions, and keep doing this until they have at least one match. Once two players have matched symbols for a given object and interaction, they keep assigning this

[4] It is of course possible that players might employ alternative strategies in Periods 31-35. For instance, they might employ a centralized learning strategy in which Player 1 chooses the same convention repeatedly and Players 2-4 "learn" from Player 1. However, we expect players with access to such a strategy would likely have employed it in Phase 1, so we consider the symmetric strategy more likely in Phase 2 following no interactions in Phase 1.

symbol to this object for the rest of the game, and continue to randomly assign the remaining symbols to the remaining objects.[5] Convention development in this case would follow the Markov process we described above, indicated by the "Symmetric" line in Figure 1.

Upon reaching Phase 2, groups learning with symmetric, decentralized strategies will revert to a lower level of coordination. This is because bilateral conventions in Phase 1 develop independently, meaning there is no unique convention group members can employ in Phase 2. Hence, they need to start at a much lower level of coordination, and only improve gradually over time. Nevertheless, as indicated in Figure 1, players can achieve slightly better-than-random coordination by exploiting historical frequencies of mappings they observed in their bilateral interactions in periods 1-30, leading to a matching probability of 27% in period 31 (see Appendix for details).

Note that even a simple, symmetric, search strategy yields an expected payoff that is much higher than the outside option. The expected per-period average payoff across all 35 periods under this strategy is XXXX, which is substantially higher than the certain payoff of XXX from all 35 periods of the No Interaction equilibrium described above. While the centralized learning process under this strategy involves some risk, the probability of obtaining lower payoffs than under no interaction is less than 0.001 percent, meaning that risk aversion alone is unlikely to account for a preference for no interaction.

Interaction with a centralized strategy. As a final benchmark, the centralized learning strategy achieves faster coordination, reduces risk relative to the decentralized benchmark above, and guarantees a single, homogenous convention to arise in all interactions within 7 periods. A key aspect of this strategy is that it relies on players using asymmetric strategies. For that purpose, we introduce (ordered) labels for players, such as the numbers 1 to 6 that we use in our experiment. As noted earlier, we assume that a natural focal player to serve the central role is Player 1.

Under a centralized strategy, Player 1 interacts with all other players in all periods of the game and serves the function of teaching a convention to all other players. To do so, she fixes one symbol-to-object mapping in the first period, and then uses this "central" mapping, $\bar{\square}_c$, for all interactions, in all periods. The remaining "peripheral" Players 2-6 interact with Player 1 in all periods, to learn this convention. Specifically, they select symbol-object combinations at random in the first period. If they successfully match with Player 1 for any

[5] The solution of the game where both players adhere to this strategy is a subgame-perfect Nash Equilibrium (SPNE). Using the restriction to symmetric equilibria together with the principle of insufficient reason also selects this solution as the unique SPNE of this game.

object, they retain the successful symbol in all subsequent matches. Otherwise, they try one of the previously unused symbols, repeating this process until they have matched on all objects. Note that this asymmetric search process *guarantees* agreement on a complete convention between Player 1 and all peripheral players in no more than 7 periods, after all the symbols have been exhausted for each object.

The peripheral players, 2-6, decline to interact with one another in period 1 and in all subsequent periods, until they have identified the complete convention employed by Player 1, At that point, they agree to all interactions with other peripheral players and use the same convention for these interactions. Hence, after a maximum of 7 periods, all players are efficiently interacting with all other players with a complete, homogeneous convention. This is reflected in the line labeled "Centralized" in Figure 1.[6]

This learning strategy also greatly simplifies Phase 2. Since all players have access to the unique convention determined by Player 1, they continue to employ this in all their interactions in periods 31-35.

As Figure 1 makes clear, there are considerable gains to a group from identifying and implementing the centralized learning strategy, meaning that groups with an awareness of the benefits of centralization should implement it. The benefits of centralization are further highlighted in Figure 2, which shows the average overall expected earnings over the course of the experiment under the three benchmark strategies. While the Symmetric strategy implies 64 percent higher earnings in comparison to No Interactions, the Centralized strategy yields an additional 21 percent higher expected earnings and considerably lower variance, in comparison to the Symmetric strategy.[7]

Figure 2. Expected total payoffs under three benchmark strategies.

[6] Note that there are additional variants that all fit within a broader centralized approach and constitute Nash Equilibria. For instance, equilibria exist in which any other player assumes the central role. More substantively different, other equilibria prescribe alternative strategies for Players 2-6 while they wait to identify Player 1's convention. Other centralized strategies transform *all* interactions into asymmetric ones by designating one player in each interaction as the central player, the other as the peripheral player. These schemes can be more efficient but are also more complex in terms of coordination. These equilibria yield similar payoffs—all higher in expectation and with lower variance than symmetric learning. We describe these alternative strategies in more detail in the Appendix.

[7] Because peripheral players wait until the coordination game yields more than the outside option, they can expect a slightly higher payoff (XXX) than the central player (XXX).

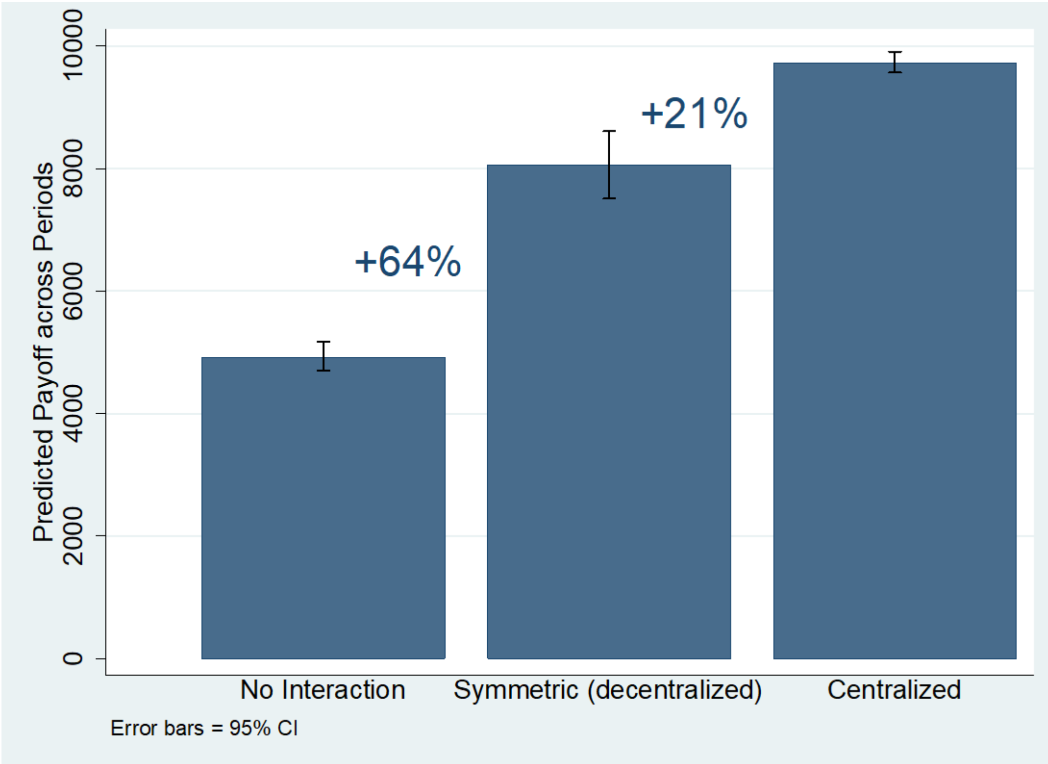

4. The Studies

Across all of our studies, the focus of our research is to what extent participants implement the centralized strategy. The most challenging case, Study 1, closely follows the above description of the game, using minimally labeled players and no other mechanisms to facilitate coordination on a centralized strategy and. The remaining studies then investigate features that may make implementing the centralized strategy easier. Therefore, Study 1 also serves as a useful benchmark against which to compare the effectiveness of the subsequent treatments.

4.1. Study 1: A Challenging Baseline

A. Experimental Design

At the onset of our experiment, participants were randomly allocated to fixed groups of 6 participants each, and interacted only with the five other members of their group throughout the study. Throughout all periods of the game, group members could identify each other by randomly assigned ID numbers 1 through 6, which were always presented in their natural order on the screen.

Study 1 implemented 2 treatments. While the matching game remained the same in both treatments the voluntary nature of interactions varied with treatment as described below.

Interaction Choices in Phase 1 (Periods 1-30). Study 1 varied by treatment whether subjects had the opportunity to select with which other group members to interact in each period. Specifically, in a *Voluntary Interaction* treatment, participants were free to choose whether to engage in the bilateral coordination game with each of the five other players in their group. That is, each participant decided how many, and which, interactions to have. In each period, group members made the decision to interact simultaneously, prior to the coordination stage. In this interaction choice stage, participants had to click "yes" or "no" next to an ID number for each of the other group members. Link formation was reciprocal; that is, both participants had to click "yes" for their interaction to materialize. If at least one of the two participants declined interaction, the interaction did not take place, and each of the two participants received the fixed outside option of 30 ECU. No explicit costs were associated with declining an interaction, or attempting to interact and being declined.

Figure 3 presents the interface through which subjects made choices in a period. At the beginning of the period, they only observed the boxes on the left of the screen, in which they could indicate their choice of whether to interact with each of the other five group members.

Figure 3. Choice interface in a period (Voluntary Interaction).

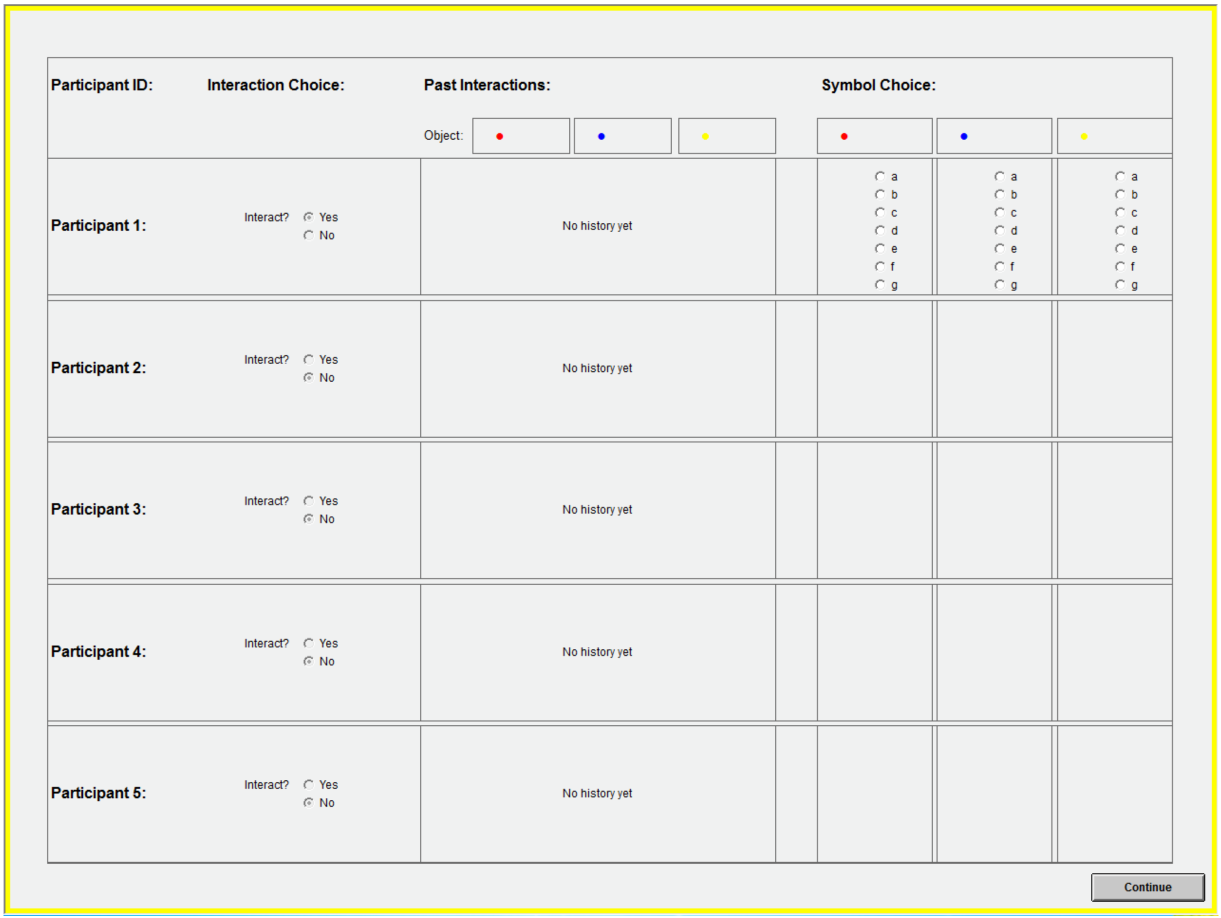

We also conducted an alternative condition, *Forced Interaction*, in which interactions were compulsory. This meant that the first stage selection of interactions was omitted, and instead each participant had to play the second-stage coordination game with each of the five other group members. This benchmark condition with an exogenously induced, complete network, allows us to establish the empirical learning rates among players who are forced to interact.

The Matching Game in Phase 1 (Periods 1-30). After the interaction choice stage, the coordination stage ensued, in which those participants who mutually agreed to interact in the Voluntary Interaction condition, and all participants in the Forced Interaction condition, played the coordination game.

The rightmost three columns in Figure 3 demonstrate the choice interface through which subjects selected symbol-object combinations in any interactions that occurred in a period.[9] In this example, which corresponds to the screen observed by the subject with ID 6,

[9] We were careful not to reveal the actual symbols and objects used in the experiment before the coordination game started. In the instructions, we used the dummy symbols ``a,'' ``b,'' ``c'' etc., and dummy objects (``L,'' ``M,'' and ``R.''). We also chose the objects to differ in color and the symbols to differ in shape so that no

there is only one interaction, with the subject with ID 1. All interacting participants chose a symbol for each of the objects, simultaneously and in private, for each period and interaction. A symbol could be assigned to at most one object within an interaction. Across interactions, symbols could be chosen independently.

The Matching Game in Phase 2 (Periods 31-35). The second phase of the experiment provided an incentive to develop a homogeneous convention easily accessible to the entire group. In this phase, participants interacted with all group members, and were required to choose a single symbol-object correspondence that applied to all interactions with the 5 other group members. Payoffs were, as during Phase 1, 20 ECU per successful match. Phase 2 was the same across all treatment conditions, and was known to participants from the beginning.

Participant Information. Participants received full information about the entire experiment at the outset of the study.[10] This included a clear explanation of the requirement in Phase 2 to interact with all other group members using a single symbol-object correspondence.

At the end of each period, participants received feedback for each interaction about the objects on which they matched successfully with the other participant. Importantly, they did not receive information about the specific symbol the other participant chose, only whether or not it matched their own symbol choice. This information was provided in a table, in which participants could scroll up or down and see their previous choice of symbols per object, as well as the outcome of these choices (match or no match), for all previous periods. The table also indicated periods in which participants did not interact with each other.

Elicitation of Beliefs. Before the first period, we asked participants to state their beliefs about the learning trajectory for two individuals who interacted continuously for 30 periods. Specifically, participants were asked to consider two participants who interacted in each period of the study, and to state, on average, how many successful matches these participants would have at three different points in time; in the 10th period, in the 20th period, and in the 30th period. Participants made their guesses without monetary incentives.

Procedures. The experiment took place at the laboratory for experimental economics at the University of Zürich.[11] In total, 162 students from the University of Zurich and the Swiss Federal Institute of Technology in Zurich took part in the study. We employed a between-

---

geometric similarities could make a particular symbol-object combination more salient than another. We also randomized the order of presentation for each participant so that subjects could not use order as a coordination device. Furthermore, we pre-tested the matching game to verify that no symbol-object pairing occurred significantly more frequently than expected by chance.

[10] Instructions for the Voluntary Interaction condition are provided in the Appendix.

[11] The experiment was conducted in English. Recruitment was conducted using h-root (Bock, et al., 2012). The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

subjects design (no subject participated in both conditions), with 27 independent matching groups in total: 120 participants (20 groups) in the Voluntary Interaction condition, 42 participants (7 groups) in the Forced Interaction condition. Sessions lasted about 2 hours.

At the end of the experiment we converted ECU into Swiss francs at the rate of 150 ECU = 1 CHF. Participants in the Forced Interaction condition earned on average CHF 61.15; those in the Voluntary condition earned CHF 55.23. We excluded students in psychology or economics, as well as participants with previous experience with experimental coordination games. Each session was preceded by a series of comprehension questions; the experiment did not start until all participants had answered the questions correctly.

B. Results

Our principal objective is to investigate whether subjects in the Voluntary condition attempt to use centralized strategies when playing the coordination game. Although difficult, coordination on a centralized strategy is at least theoretically possible, based on the asymmetry introduced by the ID numbers 1 to 6.

Figure 4 presents average earnings over the course of the experiment for the Voluntary and Forced Interaction conditions. The figure also includes the benchmark equilibria for comparison. As the Figure makes clear, average earnings fall well below those attainable under centralization. In the Forced Interaction condition, earnings are quite close to those predicted under the decentralized symmetric strategy. In the Voluntary Interaction condition, average earnings are slightly higher than predicted by the decentralized strategy at first. After period 5, however, average payoffs in Forced Interaction are higher as interacting pairs start to develop conventions. By the end of Phase 1, after period 20, subjects who are forced to interact obtain average payoffs very close to the maximum of 300 per period.

Overall, subjects earned less money per period —on average, 194 ECU — in Voluntary, versus 219 ECU in Forced ($p<0.01$, Mann-Whitney-U (MWU) test, N=27). This difference does not seem to be driven by different learning trajectories but instead by the number of forgone interactions in Voluntary. In fact, the average number of interactions per subject in the Voluntary condition is close to 3 throughout the experiment, and does not differ much by ID number (see Appendix Figure A.1).[12] The, average earnings *per interaction* and

[12] Part of this is driven by a drop in interactions over the first four periods, where subjects discontinue interactions with low numbers of matches. More precisely, subjects dissolved a total of 120 ongoing interactions in periods 2-4; of these, 99 had zero matches in the period before the dissolution, 20 had 1 match and only one interaction had 2 matches. Moreover, the low number of initial interactions is driven partly by overly pessimistic beliefs about convergence on a convention (see Appendix Table A.1).

period is 44.2 ECU in Voluntary versus 43.8 in Forced (p=0.96, MWU, N=27). Hence, when subjects chose to interact, they did equally well in both conditions; but, they interacted only about half of the time when doing so was voluntary.[13]

Figure 4. Mean earnings in Study 1 relative to benchmark equilibria.

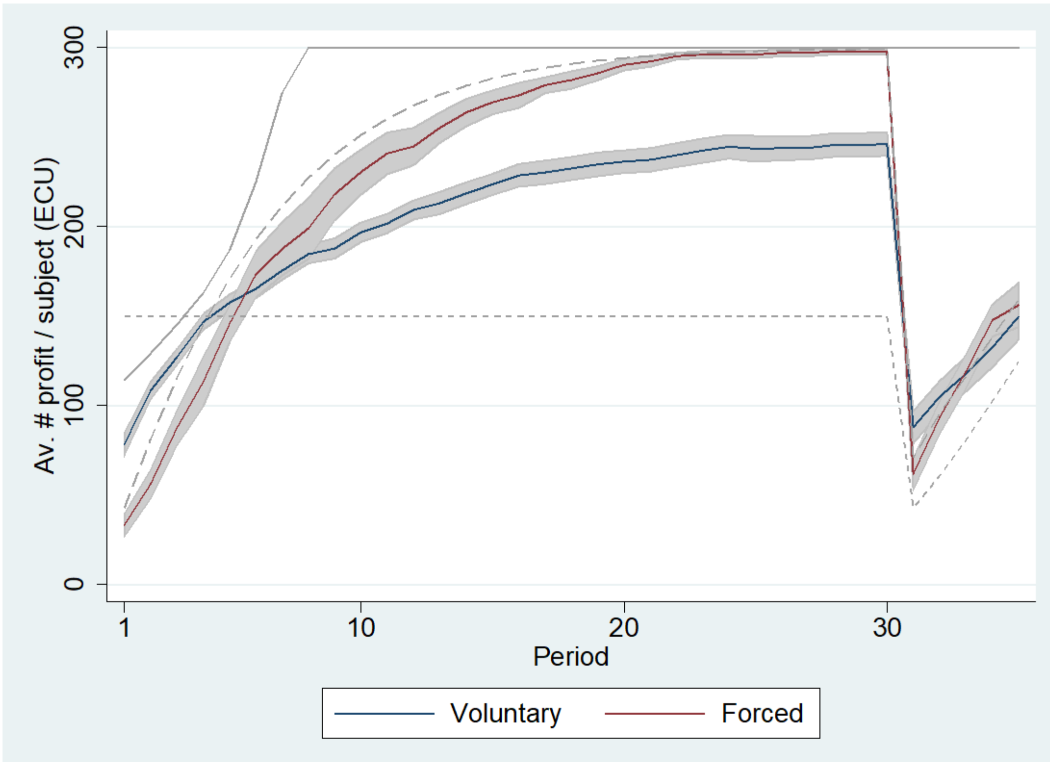

The average earnings provide little evidence of widespread implementation of the centralized strategy.[14] However, we can go further by studying whether we observe any evidence of the asymmetry in roles required by centralization. Table 1 shows, by ID number, whether subjects in either of the two conditions exhibited any systematic tendency to play in a manner consistent with either a central or peripheral role under the centralized strategy. Recall that centralization requires central players to retain their object-symbol correspondences across periods. On the other hand, peripheral players should always try a different correspondence following a failure, as they search for the convention employed by the central player. As Table

[13] Under Voluntary interaction, pairs of subjects who interacted throughout the experiment obtained payoffs very close to pairs under Forced Interaction.
[14] Of course, it is also possible that groups are heterogeneous and some implement the centralized strategy while others do not. In Appendix Figure A.2, we provide a graph showing earnings across period for each group in Study 1. Examination of the learning trajectories reveals little evidence that any groups effectively adopted a centralized strategy.

1 reveals, there is little evidence of specialization by ID type. Players all tend to retain their strategy choices more often than expected under a symmetric learning strategy and there seems to be no specialization by role.

Table 1. Percentage of subjects choosing the same symbol following a mismatch

ID	Voluntary	Forced
1	36	49
2	40	40
3	48	53
4	40	35
5	31	40
6	28	38
Overall	37	43

Returning to Figure 1, we observe another consequence of the absence of a centralized strategy: when subjects reach Phase 2, their earnings drop sharply due to the lack of a homogeneous convention. The absence of convention homogeneity can be further explored by looking at the symbol-object pairings employed by a subject in Phase 1, across interactions. As can be expected due to the apparent low degree of centralization in strategies, the homogeneity of conventions within groups is also low. For instance, if we randomly pick two "unrelated" correspondences from a group—i.e., two correspondences that come from two different subjects, were not used in the interaction with each other, and were not directed at the same subject—that subjects used in their interactions in period 30, we get on average 0.808 matches out of 3 possible matches in Voluntary and 0.560 in Forced. This is far lower than predicted under the centralized strategy, where the development of a homogeneous convention predicts the maximum of 3.0 matches.

Taken together, our first study finds that, though attainable, the benefits of centralization do not arise spontaneously. However, as we note earlier, Study 1 presents a highly challenging setting for centralization to arise endogenously, since the design provides subjects with a very minimal way of breaking the symmetry between the players and, thus, requires tacitly solving a higher-order coordination problem. Our next two studies introduce mechanisms intended to facilitate centralized coordination.

5. Study 2: Imposing Structure

Our next two studies aim to test whether mechanisms often present in organizations can aid groups in realizing the benefits of centralization. In Study 3, we investigate the effect of requiring groups to engage in communication prior to confronting the task, and whether such communication allows them to implement centralization.

In Study 2, we consider imposing varying degrees of structure on the interactions between group participants, taking them part of the way toward the adoption of the centralized strategy. Specifically, we designate one group member (Player 1) as explicitly occupying an asymmetric, central, role, thus further breaking the symmetry between players in Study 1 and making it potentially easier to coordinate on a centralized strategy with Player 1 in the central function. We also impose elements of the centralized strategy on interactions involving Player, thus restricting players to play at least some elements of the centralized strategy. However, we leave all other interactions voluntary and free of any structure.

A. Design

We impose two kinds of structure onto the design employed in the Voluntary Interaction condition of our first experiment. We implement three conditions: positional structure, action structure and combined structure.

First, we introduce *Positional Structure* by making the interaction with participant 1 mandatory. That is, in this condition all group members are required to interact with the group member with ID 1 in every period. Thus, in the interaction choice stage (see Figure 3), all interaction radio buttons on Player 1's screen, and all others' buttons for interactions *with* Player 1 were exogenously set to "yes" while the other buttons remained freely adjustable. Otherwise, the period proceeds identically to the Voluntary Interaction condition: all other participants decide whether to interact with one another and the second stage of each period proceeds as before, with all interacting players selecting symbol-object combinations while attempting to match with one another. Note that positional structure entirely solves the higher-order coordination problem regarding which group member should play a central role in teaching all others a convention—all players are required to interact with Player 1 and this is common knowledge.

Second, we introduce *Action Structure* by requiring Player 1 to use a single symbol-object mapping in all her interactions in a period. Player 1's screen in the coordination game in the second stage required allowed to set only a single symbol-object correspondence that would then apply to any interactions she had in that period. This explicitly establishes that there is a player who is required to employ the same convention in all her interactions, but leaves open whether others interact with this potentially central player.

In addition to separate *Positional* and *Action Structure* conditions, we also implemented a *Combined Structure* condition where we imposed both positional and action structure. That is, players were all required to interact with Player 1 in all periods and Player 1 was required to employ the same symbol-object mappings in all interactions in a period. Thus, essentially, in the Combined Structure condition all group members interacted with a central player and had access to a common convention once they had coordinated with Player 1. This treatment implements key parts of the centralized strategy we identified above: as long as Player 1 fixes his or her correspondence across periods, then it becomes very easy for all other players to learn this homogeneous convention and utilize it in all other interactions.

In all three conditions, Phase 2 (Periods 31-35) proceeded identically to before. All six participants in a group were required to interact with one another, and each participant had to employ a single symbol-object mapping in all five interactions.

While our first study involved a situation in which coordinating on and adopting the centralized strategy is quite challenging, this second study greatly simplifies doing so. In the Combined Structure condition, every group had an exogenously imposed central player (thus solving the higher-order coordination problem) and this central player had to use the same symbol-object combination in all five interactions. Hence, once a peripheral player (Players 2-6) coordinated with the central player, she knew that she had access to a convention that would work in interactions with all other players. We expected these interventions to result in faster development of conventions, more interactions and greater earnings in Phase 1 of the experiment and more efficient coordination throughout Phase 2, relative to the Voluntary Interaction condition of our first study.

B. Results

Figure 5 shows the realized average earnings in all three conditions of Study 2, across both Phases of the experiment. For comparison, the figure also includes the average earnings in the Voluntary Interaction condition of Study 1. A couple of observations are clear from the graph.

Figure 5. Average payoffs in Structure (Study 2) and Voluntary Interaction (Study 1) Conditions

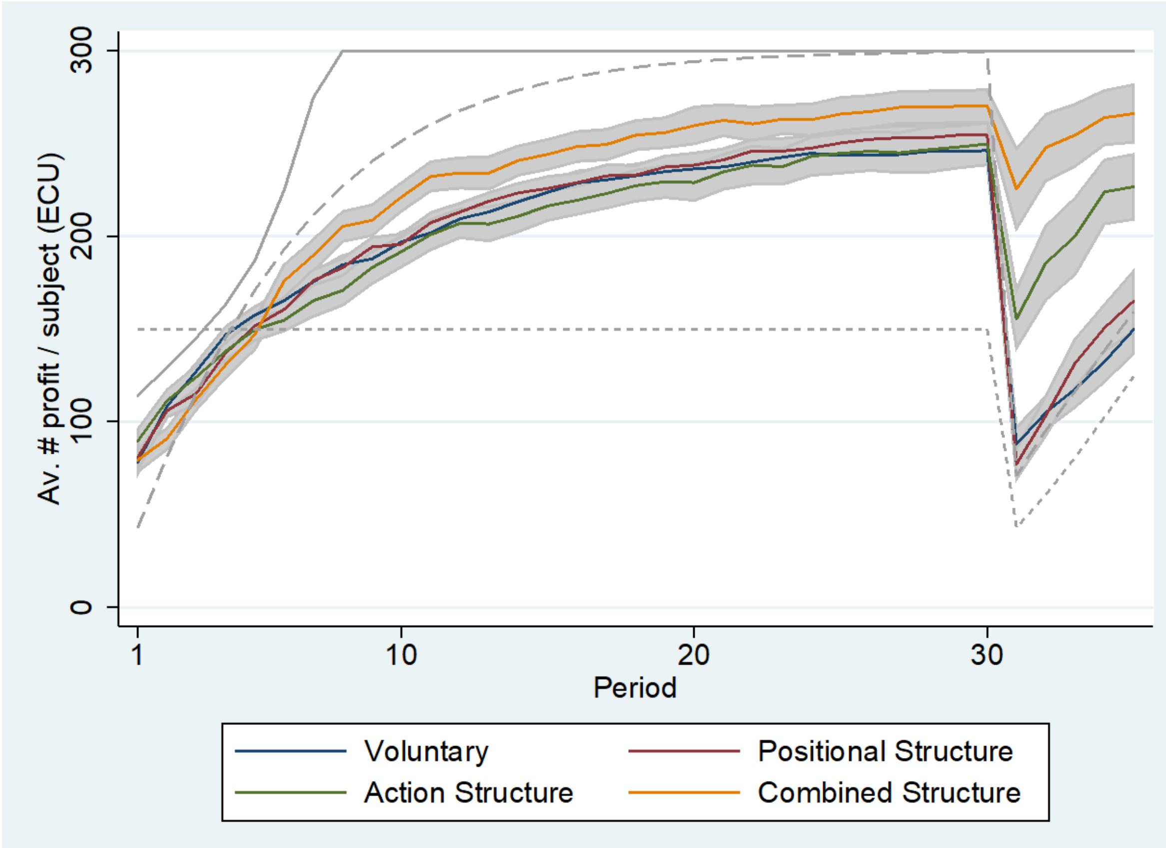

First, in Phase 1, imposing structure on the group interactions generates slight benefits only in the Combined Structure condition; though, even here, average earnings nevertheless fall far short of what they are under the centralized strategy. In fact, average earnings remain close to those under Voluntary Interactions and below those under the decentralized strategy in all three structure conditions. Thus, despite forcing subjects to adopt a great deal of the centralized strategy, particularly in the Combined Structure condition, groups remain far from the potential earnings they might obtain from taking the remaining steps and fully adopting the centralized strategy.

Turning to Phase 2, we see that both treatments that involve Action Structure yield substantial benefits relative to Voluntary Interaction. Recall that in Phase 2, all the conditions are identical—subjects are required to have all 5 interactions and use a homogeneous correspondence for all interactions. Under Action Structure, all players know that Player 1 is required to use a homogeneous correspondence in all interactions. It seems that in Phase 2, individuals at least partly recognize the value in having a common central convention that they can all access: as revealed in the right side of Figure 4, groups do substantially better in Phase

2 under Action Structure and Combined Structure. However, if the central player is not required to use a homogeneous correspondence but is instead simply required to interact with all other players (in the Positional Structure condition), there is no discernible benefit compared to Voluntary Interaction in Study 1.

To better understand the sources of increased earnings from conventions in the conditions with structure, Figure 6 shows average earnings over the experiment separately for interactions involving the central participant (i.e., all interactions involving Player 1) and all "peripheral" interactions not involving the central participant (i.e., interactions involving only Players 2-6). The figure includes earnings from unrealized interactions in which at least one player chose not to interact.

Looking first at the central participants, we see some of the benefits of requiring elements of the centralized strategy. For instance, by requiring central participants to interact with all other players, they do roughly as well as under the symmetric strategy in the Positional Structure condition. Moreover, in the Combined Structure condition, the central interactions yield profits in between those predicted under the symmetric and centralized strategies. Indeed, after a couple of periods of initial exploration, central participants in the Combined Structure condition exhibit an earnings trajectory that closely follows the trend expected under a centralized strategy. This is the result of these central participants starting to retain their symbol choices, while peripheral participants, in their mandatory interactions with Player 1, search by cycling through the symbols. Thus, with a brief delay, participants succeed fairly well in implementing the centralized strategy in interactions involving Player 1, which is reflected in Player 1's higher earnings.

Figure 6. Average payoffs in Structure conditions by player role (Study 2)

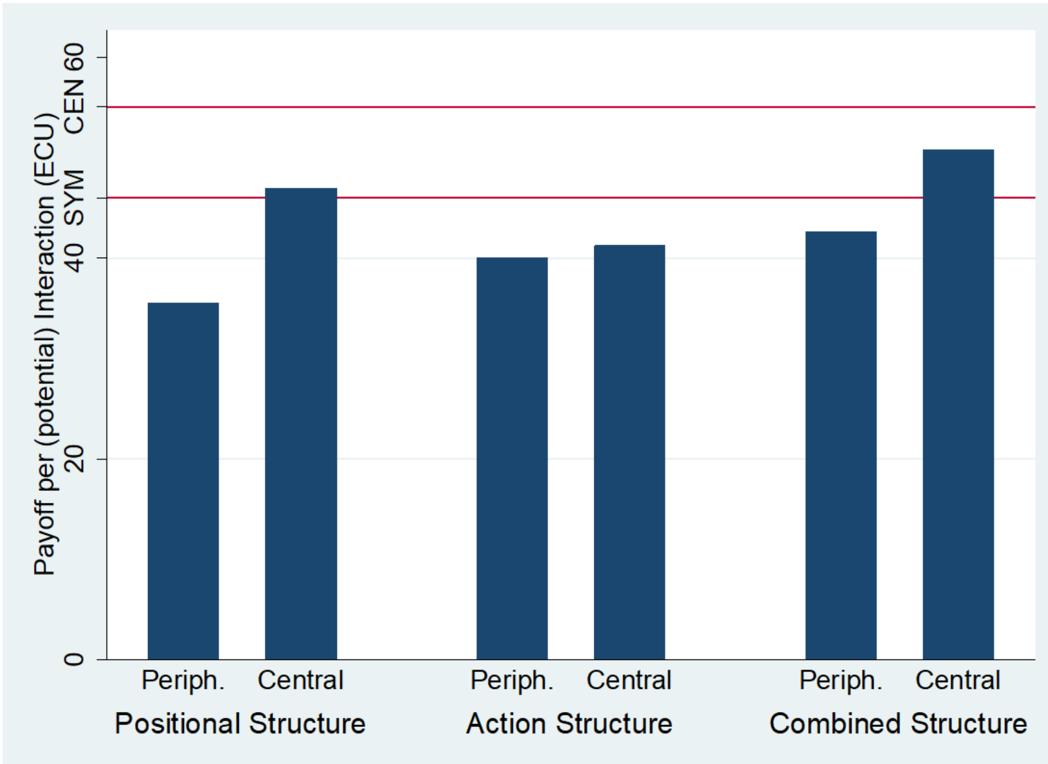

Interestingly, however, we see little spillover from the degree of centralization imposed by structure in interactions with the central Player 1 to the potential interactions between the peripheral Players 2-6. In every case, potential peripheral interactions yield earnings that are below those predicted under the symmetric learning strategy, and far worse than under the centralized strategy. This is particularly surprising in the Combined Structure condition, where the fact that everyone is interacting with Player 1 and that Player 1 uses the same convention in all interactions—which means that after a few rounds everyone has access to the same convention employed by Player 1—does not seem to translate into more efficient peripheral interactions. Instead of leveraging common knowledge of their convention with the central participant to coordinate with other peripheral subjects, Players 2-6 achieve only an average of 2.35 matches per peripheral interaction in periods 1-30, even significantly less than the 2.49 predicted under the symmetric decentralized strategy ($p = 0.04$, Wilcoxon Rank-sum, N=12).

Interestingly, interactions between peripheral subjects in the Structure conditions are no more likely than in the Voluntary Interaction condition of Study 1. In all cases, participants have roughly 2.5 interactions per period. This is despite the fact that, after a peripheral player acquires the convention employed by Player 1 in the Combined Structure condition, she knows with certainty that she can interact efficiently with any other peripheral participant who has also acquired the convention.

Moreover, even when they do interact, peripheral participants do not always employ the conventions that they developed with Player 1. In the Combined Structure condition, if one understands the coordinating benefits of centralization, then peripheral interactions are trivial once players acquire the convention developed by Player 1. However, many interactions between peripheral players exhibit a slow learning trajectory, in which players attempt to discover novel symbol-object correspondences, rather than relying on the convention they have both already obtained from Player 1.

6. Study 3: Communication

Our third study attempts to facilitate coordination on the centralized strategy using a different approach. In this case, we impose no structure or restrictions on the players, but instead allow them to resolve coordination on a centralized strategy through a pre-play communication phase. Specifically, after receiving instructions, but before playing the game, the six participants in a group have access to an electronic chat in which they can discuss possible strategies for confronting the coordination problem in our experimental task. If the main impediment to the implementation of centralized strategies in Study 1 is the higher-order coordination problem of whom to select as the central player, then allowing group members to discuss and plan on how to confront the task should greatly mitigate this problem. Moreover, if some subjects are aware of the benefits of centralization and can formulate such a strategy, then pre-play discussion might allow them to convince others of this benefit.

A. Design

The study proceeded identically to the Voluntary Interaction condition of Study 1 with one major exception. After reading the instructions, subjects were given the opportunity to communicate via an electronic chat for 5 minutes. They were not instructed on what they should discuss, but at this point in the experiment they had received extensive instructions regarding both phases of the coordination game and had completed the comprehension questions.

Importantly, the instructions also made clear that the objects ("L, M, and R") and symbols ("a, b c, d, e, f, and g") were not the same as those they would encounter in the actual study. Therefore, spending time during the chat on creating a symbol-object mapping was not beneficial.[15]

We conducted a total of 19 groups with 114 subjects.[16] Sessions again lasted approximately 2 hours, and average earnings were CHF 63.30.

B. Results

With Communication, participants can use the opportunity to discuss the upcoming game to reduce strategic uncertainty by coordinating on an asymmetric strategy profile that exploits the benefits of centralized learning. To identify possible heterogeneity in how groups use communication, we had the content of the 5-minute discussion period classified by a research assistant unaware of the research questions and without access to the actual behavior implemented by groups. This classification included, as its focus, a category for the centralized strategy that we discussed above. The research assistant identified 4 of 19 groups (21 percent) as reaching an agreement to play the centralized strategy.

Figure 7 shows the average earnings per period in the Communication condition, separately for groups that reached an agreement to implement the centralized strategy (labeled, "Comm. (Centralization)") and for the groups that are not classified as having reached such an agreement. Among the relatively small number of groups with such an agreement, the earnings trajectory is fairly close to that expected under the centralized strategy. This is true both in Phase 1 (Periods 1-30) and in Phase 2 (Periods 31-35). Therefore, allowing groups the opportunity to communicate prior to playing the game clearly allows some groups to identify and implement the centralized strategy, reflected both in their communication usage and in the earnings such groups obtain.

However, for the remaining majority of groups (79 percent), the addition of communication has only very slight benefits for realized earnings. As Figure 7 shows, the

[15] In addition, we elicited participants' expectation of matches at different points in time during the experiment both before and after communication took place, in contrast to a single elicitation in Study 1. Thus, before the experiment started, participants in the Communication condition first answered a set of control questions, then reported their beliefs about matching, communicated for 5 minutes, and reported their beliefs again, before they played the game, the two last steps being specific to the Communication condition.

[16] In two sessions we added a subtle hint with respect to the optimal organizational structure. Specifically, at the top of the screen containing the chat window, we added a simple organizational chart. The results from these two sessions are similar to the other sessions in the communication treatment, and we therefore pool the data in the subsequent analysis. Our sample also includes one additional group that is excluded from the analysis. In the first session, a bug in the software meant that the message window also appeared after the first period of the game. The experimenter instructed subject not to use the communication mechanism at that stage; however, a subject in one group did so. We therefore exclude this group from the analysis, although including this group does not substantively change the results.

earnings trajectory lies only slightly above those present in the Voluntary condition and generally below the expected earnings even under the decentralized symmetric learning strategy.

Figure 7. Average payoffs in Communication (Study 3) and Voluntary Interaction (Study 1) Conditions

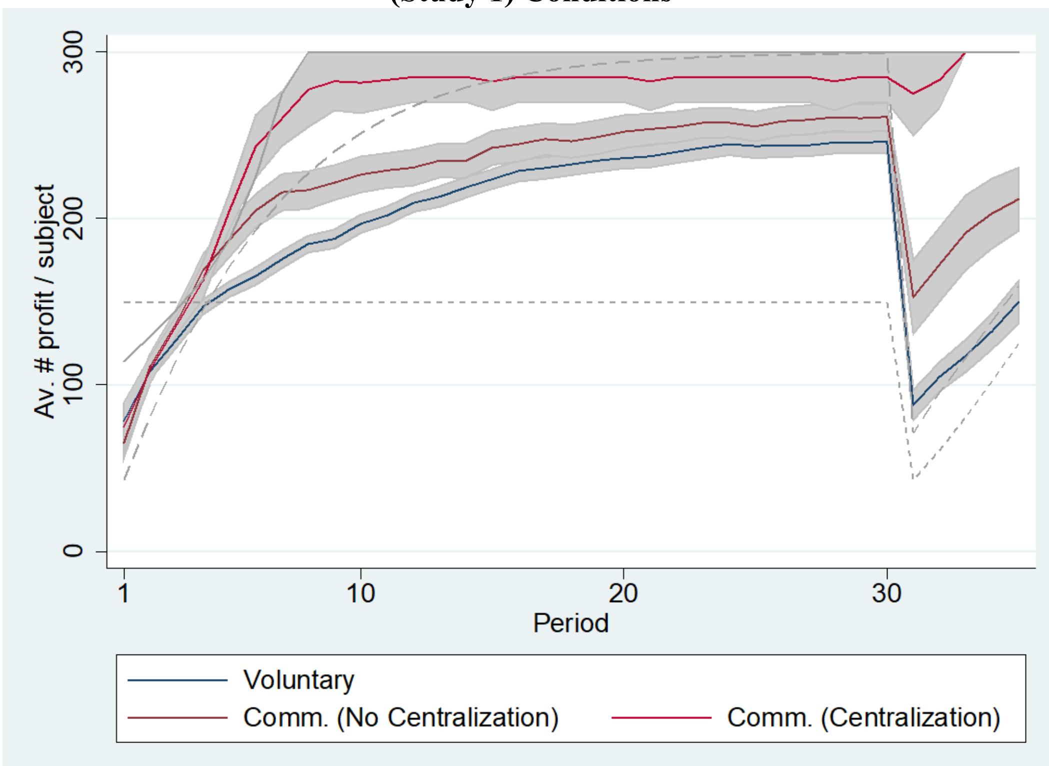

To summarize, Studies 2 and 3 introduce mechanisms that seem, *a priori*, likely to greatly facilitate coordination on the centralized strategy. However, on aggregate, there is only modest evidence that groups provided with the interventions in these two studies achieve the earnings possible under the centralized strategy.

7. Study 4: Explaining the Benefits of Centralization

As the results of the Communication condition suggest, the few groups that identify the benefits of the centralized strategy go on to implement this strategy. Moreover, the failure to obtain the benefits of the centralized strategy in the Combined Structure condition—where doing so is trivial once one realizes that everyone in the group has learned a convention from the central player—suggests that the problem in implementing such a strategy is not necessarily one of higher-order coordination challenges, but rather that many people simply don't realize the benefits of centralization for confronting this type of problem.

However, we have not ruled out that there may be other reasons why groups may fail to implement the centralized strategy, even if members are aware of such a strategy and its benefits. For instance, perhaps subjects are reluctant to play the centralized role, because it involves a slightly lower payoff than that obtained by subjects in the periphery. Or, perhaps everyone wants to play the role of the central player because of a desire to exert control or authority over the group.

To further explore whether the impediments to implementing centralization go beyond awareness of the benefits of such a strategy, we conduct an additional variant of the experiment in which subjects are described a centralized strategy and told of its benefits. However, it remains entirely up to them whether they are willing to implement it.

A. Design

Our Explanation condition is identical to the Voluntary Interaction condition of Study 1, with one modification. After receiving the instructions describing the game and completing the comprehension quiz, but before commencing play, subjects receive a brief "Explanation" before commencing Period 1. The explanation, contained on a single page, specifically stated:

> *In prior sessions of this experiment, we discovered that most groups have a hard time quickly adopting a common set of symbol-object combinations that work for all interactions in Stage 1 and Stage 2. This is despite the fact that there is a fairly simple way to achieve this outcome quite quickly, guaranteeing success in no more than 7 periods.*
>
> *Our goal in this session is to see if groups adopt this strategy when we explain it to them. Therefore, in the following, we describe this strategy to help your group solve this problem and earn the most money in the experiment. You are free to decide whether or not to follow this recommendation. However, if your group follows it, you will all earn more money than if you choose to do something else.*
>
> *The strategy involves the person with ID number 1 "teaching" a common set of symbol-object combinations to all other group members, by interacting with everyone and using the same symbol-object combinations in all interactions and in all periods. Once everyone learns the same symbol-object combinations from this the person with ID 1, then everyone can interact.*

The explanation went on to describe the centralized strategy we presented earlier in the paper in detail, with instructions for how Player 1 can confront the task in the central role and for how the remaining players should act in the peripheral functions.[17]

[17] The complete text of the explanation is provided in the Appendix.

We collected data from two sessions, involving 42 participants in 7 groups. Except for the explanation, all instructions, procedures and student population were otherwise identical to the Voluntary Interaction condition of Study 1.

B. Results

As shown in Figure 8, groups provided with an explanation of the benefits of centralization achieve the earnings feasible under such a strategy almost universally. The earnings trajectory tracks the earnings predicted under the centralized strategy almost perfectly. This reflects the fact that groups nearly perfectly implement the centralized strategy.

Figure 8. Average payoffs in Explanation (Study 4) and Voluntary Interaction (Study 1) Conditions

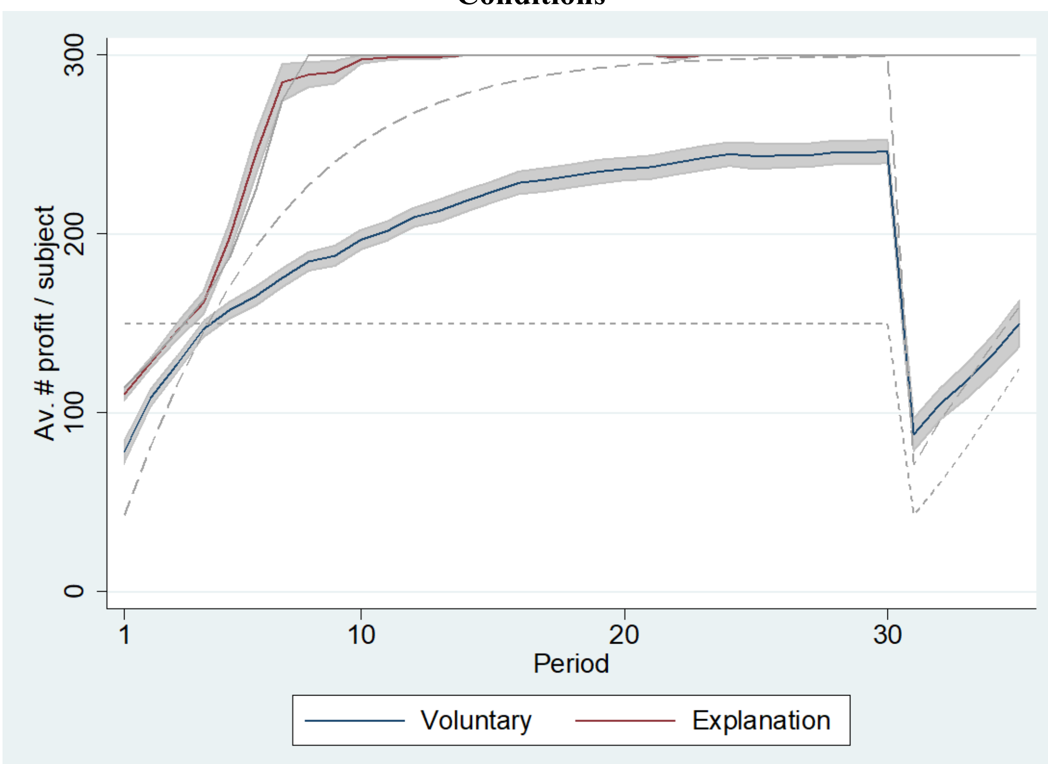

Thus, it seems unlikely that a significant reason for the failure to implement the centralized strategy in our prior studies is based on preferences against doing so. Instead Study 4, where the procedures and payoffs are unchanged and the only change is an explanation of the benefits of the centralized strategy, suggests that once people are made aware of the benefits of centralization, they implement such a strategy without fail.

8. Aggregate results across laboratory studies

Before proceeding to our survey experiment, which provides further evidence that the impediment to implementing the centralized strategy is mainly that subjects simply don't realize it, we pause to take stock of our aggregate findings. Figure 9 presents the average earnings across all of the conditions in Studies 1 through 4, along with the predicted mean earnings under the three benchmark strategies.

Figure 9. Average Earnings across Laboratory Studies 1 through 4

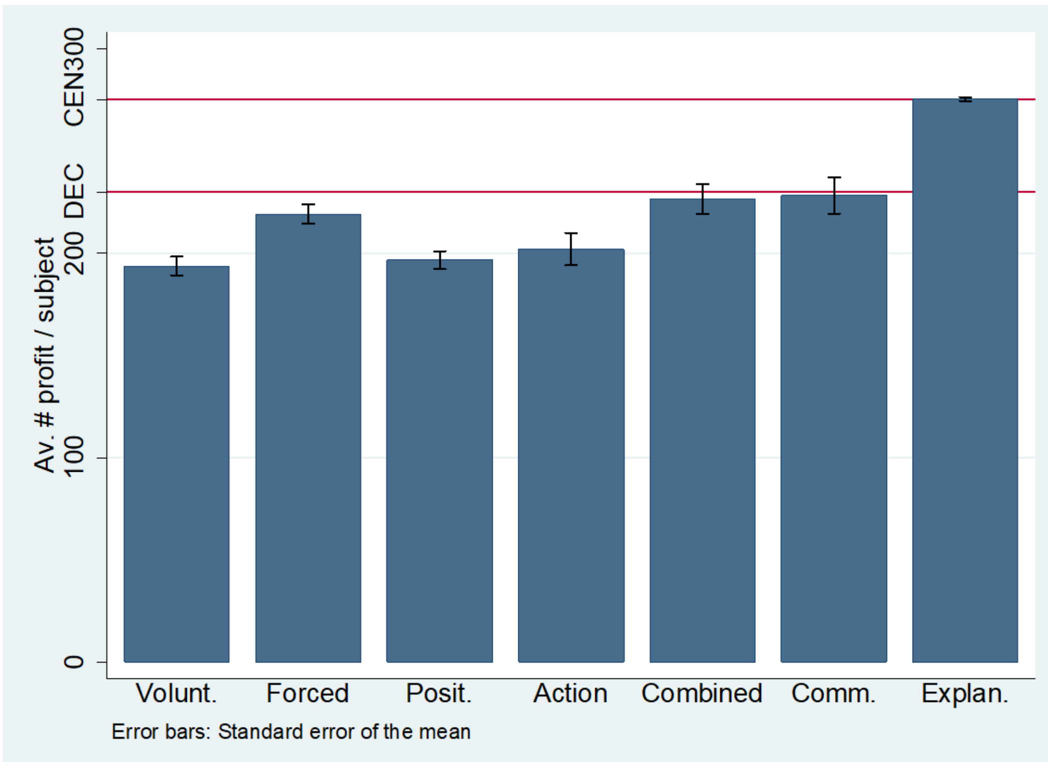

As the Figure demonstrates, most of our treatment manipulations yield at least slight improvements above the earnings realized in the Voluntary Interaction condition of Study 1. However, in all conditions other than the Explanation condition of Study 4, the mean earnings are well below those attainable under centralization and they almost never exceed even the earnings attainable under the decentralized symmetric strategy.

These observations are further confirmed in Table 2, which provides regression results of average group earnings obtained in the experiment on the different treatment conditions. The regression uses each group as an observation and uses the Voluntary Interaction condition of Study 1 as the omitted category. As the results show, and consistently with Figure 9, there are modest benefits of the first three treatment conditions, though these are not statistically significant. The benefits of Combined Structure and Communication are larger and achieve statistical significance. However, the earnings in all conditions are far below, and statistically distinct, from those in the Explanation condition. Moreover, the earnings in the Explanation condition are close to and statistically indistinguishable from the earnings predicted under centralization, which is not true of any other condition. Thus, our observation that the Explanation treatment is the only one that yields group performance close to that under centralization is confirmed in this analysis.

Table 2. Regressions of mean group earnings by condition

	Coeff.	**Std. Error**
Forced	25.38	(11.47)**
Positional Structure	2.88	(9.54)
Action Structure	8.27	(9.54)
Combined Structure	32.82	(9.54)***
Communication	37.38	(8.26)***
Explanation	81.76	(11.47)***
Constant	193.83	(5.84)***
R2	0.456	
Observations	90	

9. Study 5: Survey Experiments with Students and Experienced Professionals

The above results suggest that our subjects are unaware of the benefits of centralization for simplifying the coordination problem. A natural question is whether these findings are the result of our use of a population, university students that is inexperienced in actually organizing and coordinating teamwork. Perhaps more seasoned professionals with years of practice in managing groups may be aware of the benefits of centralization for efficiently coordinating groups.

To explore this possibility, we therefore conducted an additional, survey-based experiment in which we compare the frequency with which subjects identify the centralized approach. We examine two populations: the student population in our earlier laboratory studies and participants in the Executive MBA (EMBA) program from the University of Zürich. The latter population is considerably older (average age = 39.2), has substantial work experience (13.8 years) and draws from a variety of sectors and areas of education (e.g., 25% from banking and finance, 18% from pharmaceuticals and healthcare, 17% from engineering, 6% from telecommunications). Hence, the comparison can be helpful for identifying whether greater professional and managerial experience provides a greater ability to identify and implement centralized strategies.

As we had to fit the experiment into the limited time frame of the EMBA class, we created a hypothetical "case study" scenario that represents a stylized version of our laboratory game. In particular, it involves a group of 6 people who repeatedly decide to interact and then face a matching game similar to the one used above. The subjects' task in this experiment is simply to provide a written recommendation to the team on how they should confront the task. We collected subjects' recommendations, evaluated them, and provided a substantial monetary prize for providing the detailed recommendation that yields the greatest earnings for the team in expectation. This allows us to then code these recommendations for whether they recommend a centralized approach. In contrast with our laboratory experiments, by removing the reliance of subject's earnings on others being able to implement the centralized strategy—in this study, it is sufficient for a subject to simply recommend it—we can evaluate whether subjects indicate awareness of the benefits of centralization absent any strategic uncertainty.

A. Experimental Design

We constructed a scenario describing a team project setting in which, as in our laboratory experiments, a group of six individuals has the potential to engage in bilateral interactions over a series of days. Aside from their interactions, team members have no way of communicating with each other. Each person needs to make decisions each day for a total of 20 days. The first decision each day is about engaging in bilateral "collaborations" with the other team members. Team members give up a sure $100 in daily earnings if they interact with another team member rather than working independently. As in our game, team members make these decisions independently, and a collaboration only takes place when both parties indicate that they want to collaborate.

The second decision is about the piece of software they use for the collaborations that are realized. As in the laboratory, these decisions are made simultaneously, and are independent across collaborations. There are five different software products, each of which is equally adequate for the task, provided both parties use the same software. Collaboration between two team members is profitable, yielding $200 in earnings, only when the two people use the same software. If two team members attempt to collaborate but choose different software, then they earn $0 for their interaction. Finally, to introduce an incentive to develop a homogeneous convention, subjects are told that, at the end of the 20 days, a manager will inspect the software used on day 20 and give each team member a $1000 bonus if all six members are collaborating with everybody else, and all are using the same piece of software.

We presented this scenario, as part of a voluntary exercise, at the end of an economics class in EMBA program. The class took place directly before lunch and there were 28 students present at the time of the presentation. After receiving a written description of the scenario and hearing it read aloud, participants were instructed to provide a free-form recommendation to the team's problem by 8pm the same day. They could submit their responses using an online form. We emphasized that their solution should be as detailed as possible, describing a complete action plan for all team members for all 20 days. We incentivized the task by announcing a CHF 500 prize to the person who provided the solution that yielded the highest expected team payoff if implemented by the team members. Participants were told to submit responses individually, and to not discuss the problem among themselves.[18]

Centralization provides a simple way to solve this coordination problem, in a manner very similar to our laboratory experiment. Specifically, the problem is greatly reduced under a centralized approach in which (i) one team member assumes a central role and always uses the same piece of software in all collaborations, (ii) everybody collaborates with this central player and cycles systematically through the software options until they match; and (iii) once matched with the central player, they then attempt to use this software options in their interactions with the other, peripheral team members. In this case, a group is guaranteed to have complete coordination with a homogeneous software convention after 6 days.

In order to provide a more direct comparison with our student sample, we also conducted a version of this scenario-based experiment in the laboratory, using subjects drawn from the same university population as in our earlier studies. In this session, participants received a show-up payment of CHF 10 and competed for a prize of CHF 100 for providing the best solution.[19] Our student subjects had two handicaps compared with the executive participants. First, in the laboratory, we were able to enforce individual deliberation and prevent subjects from discussing the problem. Second, subjects in the laboratory had

---

[18] We had no way to enforce this but we assume that discussing the problem and developing solutions jointly would bias our results upwards, in the sense that we should see more efficient solutions than with truly individual solutions.

[19] The potential prize for the students is smaller than that for our EMBA sample. However, a direct comparison, given differences in income and opportunity costs for the two populations is not straightforward, even with equivalent monetary amounts. Our primary objective was to use incentives that served the function of motivating participants, and to have any potential difference lie in the direction of having stronger incentives for EMBAs.

considerably less time to provide a solution. We did not give subjects an explicit time limit for finishing their solution but suggested they should come to an end after approximately 12 minutes of writing, so that their total deliberation time was about 15 minutes. This means that, if anything, we bias our results in favor of the EMBA subjects being more likely to identify better solutions.

B. Results

Figure 10 summarizes the results. In the session involving the standard student sample, 5 out of 17 students (29 percent), provided a centralized strategy as their answer. Hence, in this individual exercise, the proportion identifying a centralized strategy is slightly higher than the frequency with which groups did so in our Communication condition in Study 3 (21 percent).

Figure 10. Frequencies of recommendations of centralized strategies across populations

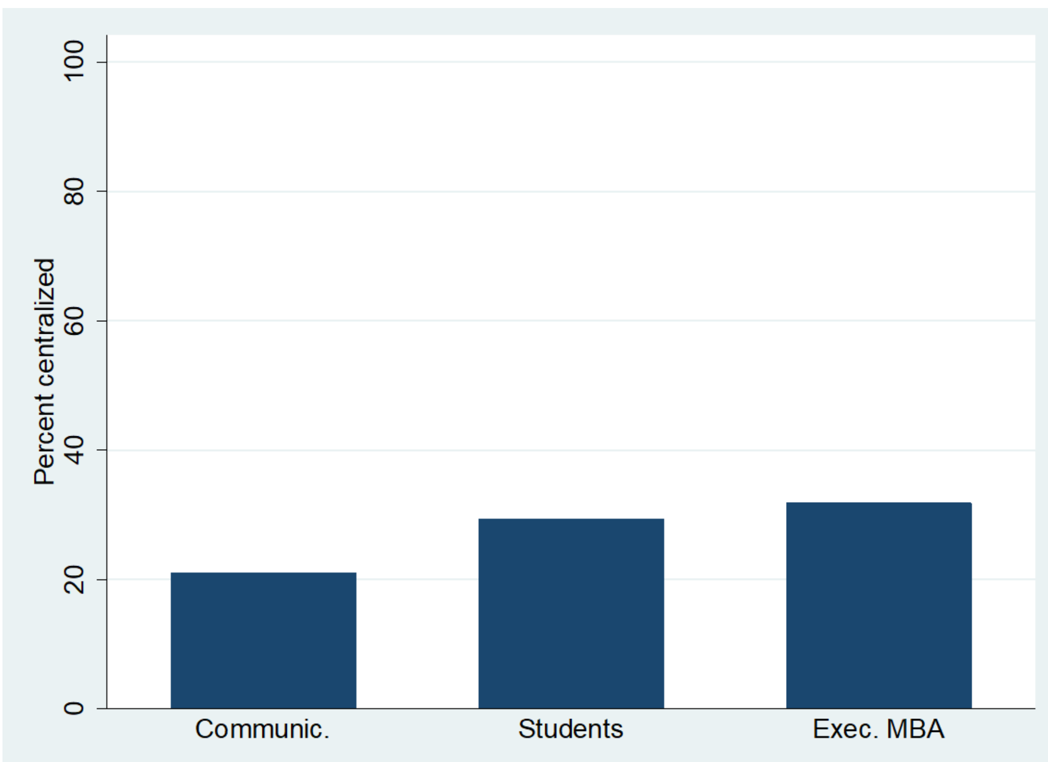

A more important comparison is with the EMBA students. We obtained 22 submissions from the EMBA students, meaning that 6 students present for the announcement did not make a submission. We treat the selected sample of 22 as the relevant group, since we lack information on the motives or thought processes behind the decisions of the remaining 6 participants.[20] We received 7 recommendations consisting of a centralized strategy, which yields a percentage of 32, quite close to the 29 percent in the standard student sample.

Thus, the results of this study yield two important findings. First, when we change our task to an individual decision problem that requires subjects to merely identify a centralized strategy, rather than implement it in a coordinated manner, we again find that few subjects spontaneously recognize the benefits of centralization. Second, a comparison between the students who comprise the population in our first four studies and Executive MBA students with more work and managerial experience indicates that the latter are not substantively more likely to recommend centralization.

Discussion and Conclusions

We conduct four laboratory studies and one survey study to investigate the extent to which groups and individuals spontaneously adopt centralization. We use a novel repeated two-stage game, in which players decide whether to interact bilaterally with others and interactions involve developing conventions to efficiently coordinate their interactions. We also add an incentive to develop a homogeneous convention. Despite the seeming complexity of the underlying repeated game, groups can easily solve the coordination problem by exploiting any asymmetry between the players to implement a centralized strategy for learning a common convention across periods.

Across all of our studies, a coherent picture emerges: despite several experimental treatments in which it should be easy for any group or individual subjects who recognize the benefits of centralization to increase their earnings by adopting such a strategy, we see only modest adoption of centralization. Even when we impose on groups a significant part of the centralized strategy (Structure conditions in Study 2) or give them the ability to organize through a lengthy pre-play chat period (Communication in Study 3), we see only modest improvements in average earnings in these conditions. We also see, in Study 5, that only about 30 percent of individuals—university students and Executive MBA program participants alike—recognize the value of centralization. However, when we give players a simple, non-binding recommendation to play the centralized strategy, in Study 4, we see near universal adoption of the strategy. Viewed jointly, the above results all suggest that the main problem is simply that subjects are unaware of the benefits of centralization.

---

[20] However, if we include them as cases that did not identify the centralized solution, this lowers the success rate for the EMBA sample to 25 percent, which remains close to that with the standard student sample.

This observation is consistent with other findings that suggest subjects fail to realize that simple mechanisms can greatly facilitate group coordination (Heath and Staudenmayer, 1998; Kriss, Blume and Weber, 2016). In these studies, as in ours, groups face seemingly complex coordination problems, but there exist relatively simple ways of using communication or organizing team activity in order to greatly resolve the problem. Nevertheless, both unawareness of the benefits of these mechanisms mean that they are underutilized and the coordination failure in groups is substantial. We similarly observe groups leaving a substantial amount of potential earnings unrealized—roughly, between 20 and 30 percent—because they fail to recognize and implement centralized strategies.

Our findings are important for understanding one potentially significant source of coordination failure in organizations. When centralization and hierarchy provide mechanisms to greatly reduce coordination problems, this efficiency gain may nevertheless not realize because of individuals' unawareness of such benefits. For instance, the head of a regional chain of stores may not schedule enough meetings or utilize technology in a manner that allows all local employees to be on the same page in marketing or sales. Similarly, where a strong culture is important for an organization, a disparate set of local cultures may arise due to a lack of proper use of centralization to establish a superordinate culture and instill its values and beliefs in everyone throughout the organization.

References

Appendix

Figure A.1. Number of Interactions in Voluntary Interaction (Study 1)

Table A.1. OLS Regression of number of chosen interactions in period 1 (Study 1)

	Beliefs	Risk	Cognitive	All
Belief: Period 10	-0.077			-0.038
	(0.242)			(0.244)
Belief: Period 20	0.199			0.038
	(0.424)			(0.414)
Belief: Period 30	0.963***			1.015***
	(0.366)			(0.358)
Risk: General		0.081		0.074
		(0.077)		(0.078)
Risk: Driving		-0.118*		-0.059
		(0.066)		(0.072)
Risk: Financial		0.023		0.006
		(0.079)		(0.065)
Risk: Sports/Leisure		-0.011		0.006
		(0.080)		(0.077)
Risk: Career		0.015		0.008
		(0.099)		(0.096)
Risk: Health		0.094		0.068
		(0.059)		(0.058)
Risk: Hypoth. Lottery		-0.070		-0.022
		(0.128)		(0.118)
Cogn.: CRT Score			0.246*	0.172
			(0.145)	(0.138)
Cogn.: Quant. Major			0.206	0.300
			(0.446)	(0.335)
Age	0.004	0.039	0.022	0.021
	(0.053)	(0.057)	(0.052)	(0.054)
Female	-0.506**	-0.592*	-0.451	-0.288
	(0.255)	(0.341)	(0.315)	(0.298)
Constant	1.765	2.985**	3.165**	0.302
	(1.414)	(1.477)	(1.376)	(1.523)
R^2	0.289	0.098	0.070	0.339
Adj. R^2	0.258	0.024	0.037	0.250
N	120	120	120	120

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Prior to Period 1, we elicited subjects' expectations about matching success over time and the exit survey included questions on risk attitudes (Dohmen et al. 2011) and the Cognitive Reflection Task (Frederick 2005). With these data, we ran an OLS regression of the number of first-period interaction choices on the expected number of matches at different points in time, the seven risk attitude items, the CRT score and the sociodemographic controls age, gender, and quantitative major. This regression suggests that the key reason for choosing not to interact was the belief that continued interaction would result in few matches: raising a subject's ex-ante expectation by one match in an average interaction in period 30 increases the number of interactions chosen in period 1 by 1 ($p<0.01$). Risk attitudes and cognitive capabilities, on the other hand, seem to be irrelevant.

Appendix – Additional Information provided in the Explanation Instructions

Note that, across the 35 periods of this task, your group does best when

- every group member interacts with all five other group members and successfully matches on every symbol-object combination (i.e., you each have 5 interactions and get 3 x 20 = 60 ECU from each of them in a period, or 300 ECU in total), and
- all 6 group members use a single common set of symbol-object combinations (this allows the group to perform better in Stage 2, obtaining 300 ECU in every period).

You, and the whole group, do best when you achieve these outcomes as quickly as possible.

In prior sessions of this experiment, we discovered that most groups have a hard time quickly adopting a common set of symbol-object combinations that work for all interactions in Stage 1 and Stage 2. *This is despite the fact that there is a fairly simple way to achieve this outcome quite quickly, guaranteeing success in no more than 7 periods.*

Our goal in this session is to see if groups adopt this strategy when we explain it to them. Therefore, in the following, we describe this strategy to help your group solve this problem and earn the most money in the experiment. You are free to decide whether or not to follow this recommendation. However, if your group follows it, you will all earn more money than if you choose to do something else.

The strategy involves the person with ID number 1 "teaching" a common set of symbol-object combinations to all other group members, by interacting with everyone and using the same symbol-object combinations in all interactions and in all periods. Once everyone learns the same symbol-object combinations from this the person with ID 1, then everyone can interact.

Specifically, the precise strategy is

- Everyone **interacts with the participant with ID 1 in all periods.** That is, everyone selects "yes" in all interactions with ID 1 in all periods and ID 1 selects "yes" in all interactions with all group members.
- The participant with ID 1 chooses a single set of symbol-object combinations and **uses this same combination in all 35 periods and in all interactions.**
- When interacting with the participant with ID 1, **all other group members try to find this combination** by rotating systematically through each of the 7 symbols, trying each one once until they succeed.
- Once a participant with ID 2, 3, 4, 5 or 6 matches symbols for **at least two objects** with participant 1 (that is, earnings from interacting are 40 ECU, which is higher than the fixed payment of 30 ECU from not interacting), then (s)he **attempts to have all interactions will all other participants, and uses the same symbol-object combinations (s)he is using with participant 1** in all interactions.

If everybody in the group follows this plan, this will guarantee that every participant can learn the symbol-object combinations that participant 1 is using in at most 7 periods (once they have tried all the symbols). At that point, since participant 1 has taught all group members the same combinations, they can all interact with one another and all earn the maximum of 60 possible ECU from each interaction. In Stage 2, everyone can again use the same symbol-object combinations and earn the maximum amount of money possible.

Appendix – Instructions for MBA Survey Study

Introduction

Thank you for participating in this decision exercise.

In the following, I will read a hypothetical scenario about a work team facing a challenging task. You should provide a written recommendation to the team about how to best approach their task. Specifically, I ask you to write down a strategy for the team, which allows them to learn to collaborate as quickly and effectively as possible in the task.

We will then evaluate all of your recommendations, in order to identify the *best* recommendation—that is, the one that will yield, in expectation, the highest total profits for the team. The person in this room who proposes the best recommendation will receive a **payment of CHF 500**. Below I first present a summary of the problem, and then explain each step in detail. *Remember that you are competing with others for the best response. Therefore, it is important that you think about this task individually and not discuss your recommendation with someone else before you have sent it to us.*

Scenario Summary

The work team consists of 6 people. Each day, over a space of 20 days, each member of the group needs to make two decisions.

- First, each member has to decide whether to work independently or try to collaborate bilaterally with each of the other team members. This means that a team member can collaborate bilaterally with between 0 and 5 other team members on each day.
- Second, for each collaboration in a day, each group member decides what software to use. Collaboration between two team members is profitable only when the two people use the same software; otherwise, both team members would have been better off working independently. The problem is that each pair of group members attempting a collaboration cannot plan ahead which software to use, meaning that each member must try to guess which software the other one is using.

Your job is to provide a written recommendation for the work team, containing a strategy about how to best approach this problem.

Detailed overview

Step 1: Choosing with whom to interact

The 6 people in the work team have the opportunity to collaborate in pairs each day for 20 days. All group members are identified with a unique ID-number “1,” “2,” “3,” “4,” “5” and “6,” and interact only via a computer interface.

The first decision for each team member, at the beginning of each day, is to **decide with which of the other 5 individuals he or she wishes to collaborate that day**. An individual can try to collaborate with all other 5 individuals, with any number fewer than 5, or with no one. So, for example, **individual "1"** begins the work day by seeing a computer screen that has five boxes, one for each group member, like this:

☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6

Individual 1 can click each box to indicate with which other team members he or she wants to collaborate that day. All team members thus simultaneously indicate with which group members they would like to collaborate that day.

A collaboration between two team members occurs only if **both** members indicate that they would like to collaborate. This means that each group member will have between 0 and 5 separate collaborations each day, based on how many boxes that individual clicks and how many other people click their boxes that correspond to that individual. Each day, all team members make a new set of decisions about whether to collaborate with others. Each group member only knows about his own collaboration choices, not about interactions between other group members.

Each person **starts the day with $500**. Collaborating with someone **costs $100**. So, for each collaboration that an individual has, his earnings decline by $100. However, each collaboration also creates a **possibility to make even more money**, as described below.

Step 2: Matching software

The second decision for a team member is to choose, for each separate collaboration he or she has that day, a software to use in that collaboration. **A collaboration is successful only if both people use the same software**. Specifically, there are 5 different kinds of software that team members can use. The names of the software are unknown in advance, so they are instead indicated by the symbols:

* % & / ?/

If, for example, individual 1 is interacting with individual 2, then their collaboration is successful if they both select the software denoted with a "?" and is unsuccessful if one of them selects "*" and the other selects "&."

Specifically, for every interaction that takes place in a day, all participants see the set of possible software to choose from, and choose which software to use for each specific collaboration. However, the software types are presented in a random order to all group members, so each individual sees the software in a different order. Therefore, group members cannot simply choose the first or last in order to pick the same. Note that, since there are 5 types of software, then two people who are collaborating and who select software *at random* will have a 20 percent chance of matching.

For example, if the group member with ID 1 is interacting with the group member with ID number 2, he or she sees the following on the computer screen:

Interaction with 2: ☐ * ☐ % ☐ & ☐ / ☐ ?

Group member 1 must now select one software type. Group member 2 will have the same set of software choices, but will see them in a *different order*. Both members decide **simultaneously**, without knowing what the other chooses.

At the end of the day, the computer will tell the two group members if they matched in their software choice. If they did not match, it will not tell them what choice the other person made. So, each member knows only his own software choices, and whether a choice matches the other's choice or not in a particular interaction. Without any additional information, the only way for two people to learn how to match software is through *trial and error*. Once a collaboration is successful, those two team members can continue to successfully collaborate in all subsequent days by selecting the same software again.

Note that it is possible for an individual to use two different software types in two different interactions. For example, if individual 1 is collaborating with individuals 2 and 3, he may use "*" in his collaboration with individual 2 and "&" in his collaboration with individual 3. Or he may use the same software in both collaborations. Each of these collaborations may be successful or unsuccessful, independently of whether the other one is successful or unsuccessful. So, an individual can have anywhere between 0 and 5 collaborations, each of these can be either successful or unsuccessful, and each of them can involve using the same or different software than in the other collaborations.

If two individuals **successfully** collaborate by picking the same software, they each receive **$200** for that collaboration.

Remember that, in order to collaborate, both individuals give up $100. This means they are **better off (by $100)** if they successfully collaborate than if they had not collaborated. But if the collaboration is **unsuccessful**, then they each receive nothing, meaning that they are **worse off (by $100)** than if they had not attempted to collaborate.

Individuals on the team have no way of communicating prior to beginning the task for 20 days. They also cannot communicate about what software they have used or are likely to use.

The manager

At the end of the 20 days, a manager will inspect what software the different team members are using. The manager will give **each group member a bonus of $1000** if, on the last day, they are **all collaborating** and are **all using the same software**. The manager will not interfere before this point, but has hired you to prepare a brief memo to distribute to the group regarding how they should confront this problem.

Your Task:

Your task is to write a few sentences to the team of 6 workers regarding how they should attempt to perform this task.

You need to provide a **detailed plan** for how team members should approach the decision of (1) with whom to collaborate and (2) how they should approach the decision of which software to use in any collaborations that they have.

Please be precise. Your recommendation can include rules about how to try to search for software matches with collaborators and about with which other group members to interact. We ask you to provide your answer at http://goo.gl/forms/E73bwXFI44 before 20.00 today.

Remember that ***you do not yet know*** *the names of the software that team members will use, so you cannot tell them to use a particular one. Also, the software will show up at different places on the screen for different people and may not have names that can be ordered alphabetically or numerically, so you cannot tell them to pick one based on the order.*

Evaluating your suggestions

After we collect all of your responses, *we will identify the expected payoffs to the team*, in combination, from following your suggestion. Thus, your suggested plan has to be sufficiently precise for us to identify what every team member will do at any point. From all of the suggested plans you submit, we will select the plan that generates the highest earnings for the team, in expectation, if they follow that plan. The person who suggested this plan will receive CHF 500. This will be paid to you before the conclusion of the course. If two or more people have suggested plans that are equally effective, or they give the same plan, then one of them will be randomly selected to receive the prize.