

Collective Problem-Solving of Groups Across Tasks of Varying Complexity

Authors

Abdullah Almaatouq,^{1*} Ming Yin,^{2*} Duncan J. Watts^{3,4*}

Affiliations

¹Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139.

²Department of Computer Science, Purdue University, 305 N. University Street, West Lafayette, IN 47907.

³Microsoft Research, 641 Avenue of the Americas, 7th Floor, New York, NY 10011.

⁴University of Pennsylvania, Philadelphia, PA 19104

*Correspondence to: amaatouq@mit.edu, mingyin@purdue.edu, djwatts@seas.upenn.edu

Abstract

As organizations gravitate to group-based structures, the problem of improving performance through judicious selection of group members has preoccupied scientists and managers alike. However, it remains poorly understood under what conditions groups outperform comparable individuals, which individual attributes best predict group performance, or how task complexity mediates these relationships. Here we describe a novel two-phase experiment in which individuals were evaluated on a series of tasks of varying complexity; then randomly assigned to solve similar tasks either in groups of different compositions or as individuals. We describe two main sets of findings. First, while groups are more efficient than individuals and comparable “nominal group” when the task is complex, this relationship is reversed when the task is simple. Second, we find that average skill level dominates all other factors combined, including social perceptiveness, skill diversity, and diversity of cognitive style. Our findings illustrate the utility of a “solution-oriented” approach to identifying principles of collective performance.

Introduction

Problem-solving in groups is ubiquitous throughout the economy and society. Business firms have long been highly reliant on teams for functions as diverse as engineering, design, and marketing, but other domains including science, are also increasingly group-based (1). Naturally, questions about how to most effectively construct and manage groups have also preoccupied researchers across a variety of fields, including psychology, economics, management science, and more recently complexity science (1–6).

In spite of this attention, research on the collective performance of groups of problem-solvers has often reached inconsistent or conflicting conclusions. While some studies find that groups dramatically outperform individuals (1, 2), others find that “process losses” cause groups to underperform their best members (7, 8). Complicating matters further, other studies find that groups outperform individuals under some conditions, but not others (4, 9, 10). Thus, it seems that real groups can outperform nominal groups but that (a) they do so only rarely in the literature, and (b) the conditions under which they do are not well understood.

Another area of disagreement pertains to the effects of different group compositions on collective performance. For example, lab studies conducted between the 1980s and mid-2000s (11–15) found that average ability was the most consistent predictor of collective performance. More recent studies, however, have argued the opposite: that average ability is less relevant to collective performance than factors such as social perceptiveness (aka “emotional intelligence”) (3, 16–18), skill diversity (19, 20), and cognitive style diversity (21–24). Finally, because task complexity (25) itself is rarely varied systematically within a single study, little is known about its mediating effect on the relative performance of groups versus individuals, as well as the relationship between group composition and performance.

Reading this literature, a hypothetical manager would have difficulty deciding on whether or not to construct a team for some task, and if so, which of potentially many individual-level attributes to measure, how to optimally combine individuals with those attributes, and how that combination might depend on the complexity of the task at hand. Motivated by this practical managerial problem, we conducted a novel “two-phase” experiment to answer three main questions (pre-registered at [AsPredicted.org #13123](https://aspredicted.org/#13123)): (1) Do groups perform better than individuals in the context of a specific class of complex tasks, defined below? (2) Which of several competing

group composition attributes dominate collective performance? (3) Are these results mediated by task complexity? In phase one we measured several relevant attributes for individual workers (i.e., skill, social perceptiveness, and cognitive style); then in phase two, we used this information to construct groups with desired combinations of individual attributes (i.e., group-level skill, skill diversity, group-level social perceptiveness, and cognitive style diversity).

Our approach differs from previous work in several respects. First, by varying four widely studied attributes of groups simultaneously, we investigated the relative importance of these attributes individually and collectively in predicting collective performance. Second, we directly measured individual skill and problem-solving style (i.e., cognitive style) on the task in question before assignment to groups. Third, we used a block randomization scheme that intentionally oversampled infrequent combinations of individuals (e.g. “all high skill and high social perceptiveness”), thereby greatly increasing our statistical power. Fourth, by systematically varying the complexity of the task over a wide range (from “low” to “high” complexity) without changing the nature of the task, we determined how, or if, the relative performance of groups versus individuals or the importance of different attributes changes with task complexity (e.g., does social perceptiveness, or skill diversity, matter more for the most complex tasks than for simple tasks?). Fifth, by evaluating all group composition effects in terms of out-of-sample prediction of the outcome of interest (i.e., collective performance), we directly address the motivating problem of optimal group construction. Finally, we pre-registered our research questions and analysis plan (see SM Section 1 for exceptions), thereby increasing the replicability of our findings (26).

Room Assignment Task

The main task in question was a “*room assignment*” problem in which participants—first as individuals and then possibly in groups—were required to assign N “students” to M “rooms” where each student had a specified utility for each room. Their objective was to maximize total student utility while also respecting Q constraints (e.g., “Students A and B may not share a room or an adjacent room;” see Figure 1 for an illustration of the task; see Figures S1-S2; see SM Section 2.1 for more details). When the task was done in groups, participants were allowed to move different “students” simultaneously and, therefore, can perform parallel processing;

however, they were blocked from moving the same student at the same time (i.e., to avoid generating both human confusion and software errors).

We chose the room assignment task for four reasons. First, it is a specific instance of a more general class of complex problems known as Constraint Satisfaction and Optimization Problems (CSOPs), which are widely studied in artificial intelligence and operations research (27); thus, our findings will inform collective solution of CSOPs in general. The connection to operations research is useful because, unlike other “toy” problems in theoretical computer science, CSOP problems map in a relatively intuitive way to a range of practical resource allocation problems. For instance, CSOPs have been used to model many problems that are of practical interest, from staffing software projects where there are several developer-to-activity combinations to evaluate (28) to forming learning groups based on some criteria related to the collaboration goals (29) to the railway timetabling (30). Second, CSOPs are an abstraction of many resource allocation and optimization problems; thus, they capture important features of real-world group problem-solving exercises without requiring participants to have specialized skills. Third, as with other complex problems (31), the payoff function for CSOPs can be described as a “rugged landscape” characterized by many locally optimal but globally suboptimal solutions. Correspondingly, CSOPs are amenable to potentially many solution strategies and styles, where no single strategy is universally superior (32). Fourth, the complexity of CSOPs can be systematically varied by adjusting a few key parameters; in our case, by changing the number of students (N), the number of rooms (M), and the number of constraints (Q). The connection of the task complexity to algorithmic complexity is useful because it admits a formal definition of “complexity” (i.e., the run-time of an algorithm) which in turn allows for ranking problems of the same class in terms of increasing complexity. Specifically, we find a useful mapping between “component complexity” and “coordinative complexity”, as defined in (25), and parameters of CSOP. By contrast, assigning complexity scores to the parameters of arbitrary tasks is much more difficult to do in a systematic way, in part because we lack a formal language for task complexity. By drawing tasks from a field that already possesses a formal language for the complexity we made one aspect of our design simpler.

Figure 1 illustrates the variability in complexity with two instances of the room assignment problem: a “simple” instance, which involves assigning six students to four rooms subject to only two constraints (“B and E must be neighbors” and “C and F can’t live in the same room or be

neighbors”); and a “super complex” instance, which involves assigning 18 students to eight rooms subject to 18 constraints.

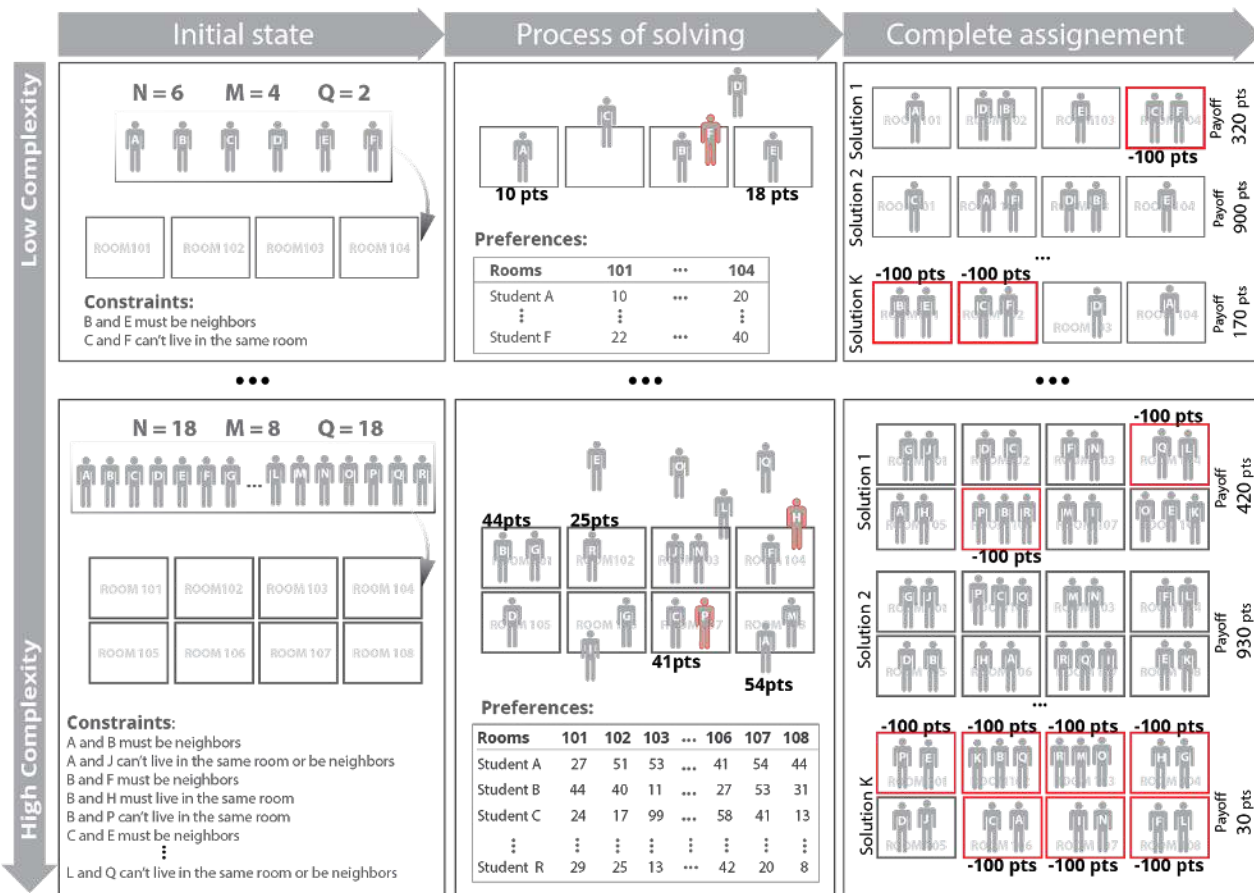


Figure 1. Illustration of the “room assignment” task. The task required to assign N “students” to M “rooms” so as to maximize the total utility of the students, each of who has a specified utility for each room, while also respecting Q constraints. The complexity of the task is characterized by the different number of students to be assigned (N), the number of dorm rooms available (M), and the number of constraints (Q). The first row shows a low complexity case in which six students are to be assigned to four rooms subject to two constraints. The second row shows a high complexity case in which 18 students are to be assigned to 8 rooms subject to 18 constraints. See also Figures S1-S2 for screenshots of the experiment interface.

Experimental Design

In phase one, 1200 participants recruited from Amazon’s Mechanical Turk completed five room assignment tasks (three “very low” and two “moderate” complexity tasks), as well as a standard “Reading the Mind in the Eyes” (RME) test (33), which is commonly used as a measure of social perceptiveness (see SM Sections 2.2 and 2.3). In the RME test, participants were shown 36 pairs

of eyes, where for each pair of eyes they had to choose one of four words describing the corresponding emotion (Figure S3). This test was used by a number of recent studies relating social perceptiveness to group performance (3, 16–18) and it has been shown that it is equally predictive of collective performance for both face-to-face groups (interacting freely in a room) and online virtual groups—interacting via text-based chat and could not see each other’s eyes or facial expressions at all (17, 18). These findings support that the RME test is capturing a deeper, domain-independent aspect of social reasoning, not merely the ability to recognize facial expressions of mental states. After the completion of phase one, we evaluated all participants on skill level, social perceptiveness, and cognitive style (see SM Section 2.4.2 for phase one details).

Skill. Our primary definition of skill was the sum of scores on the two moderately complex room assignment tasks. The score a participant earned in a room assignment task considers both the total utility of students decided by rooms they were assigned and penalties caused by violations of any constraints (see SM Section 2.1 and SM Section 2.4.3 for details). Individuals who scored above/below the median skill score were classified as high/low skill, respectively.

Social perceptiveness. We defined social perceptiveness simply as the number of RME questions correctly answered. As with skill, individuals above/below the median social perceptiveness score were classified as high/low social perceptiveness, respectively.

Cognitive style. Finally, based on the participants' answers to a post-experiment survey question, we defined an individual’s cognitive style as belonging to one of two categories: “optimize first,” who indicated a preference for allocating all students to rooms for which they had the highest utility before attempting to resolve conflicts; and “satisfy first,” who indicated a preference to first allocate all students with conflicts before moving students to higher-value rooms. Our definition of cognitive style proceeds from three criteria (23): it must be persistent for a given individual (i.e., consistent across tasks); it must be heterogeneous across the sample (ideally, roughly equal numbers would have each style); it must not be highly correlated with skill (see SM Section 3.4 for more details and alternative definitions). Our specific measure of cognitive style (“optimizer” vs “resolver”) had an average test-retest reliability of 0.72 (see SM page 20), which falls within the recommended range of 0.7-0.9. Also, the other measures of cognitive style that we reported in the supplementary materials also pass the test-retest reliability (0.74 for constraint violation tolerance; and 0.71 for preference for efficiency vs. perfection).

In phase two, we recruited the same 1200 participants and allowed 828 of them (as per our pre-registration; see SM Section 1 for sample sizes) to perform a second sequence of five room assignment tasks (task sequence is randomized), also of varying complexity (very low, low, moderate, high, very high; all tasks timed out at 10 mins in phase two, regardless of complexity). Based on each participant's phase one labels for skill and social perceptiveness, we first assigned each individual into one of the six blocks: HH (all individuals in this block are classified as high skill and high social perceptiveness, $N = 100$); MH (contains a mixture of high/low skill individuals with high social perceptiveness, $N = 213$); LH (all individuals in this block are classified as low skill and high social perceptiveness, $N = 90$); HL (all individuals in this block classified as high skill and low social perceptiveness, $N = 97$); ML (contains a mixture of high/low skill individuals with low social perceptiveness, $N = 221$); and LL (all individuals in this block are classified as low skill and low social perceptiveness, $N = 107$). Next, within each block, individuals were randomized to "group" ($N=591$ participants, forming 197 groups) or "individual" ($N=237$) conditions; and finally, groups were constructed also at random (see Figure 2 for overall experimental design; Section 2.4.4 for details on experiment phase two design). The main purpose of the block randomization scheme was to oversample statistically less frequent combinations (e.g., all group members had high skills or high social perceptiveness), thereby increasing the statistical power of our experiments (a secondary benefit was that it allowed us to match the distributions of participant types in phases one and two; see SM Section 2.4.5).

Three metrics were used to capture an individual or a group's performance in a room assignment task T of phase two: First, we considered the normalized score the individual or group obtained on this task, defined as $\frac{\text{score on task } T}{\text{max score for task } T}$; second, we measured duration (or time to completion), defined as the time elapsed from the start of the task until a solution is submitted (where we note that all tasks timed out at 10 mins in phase two regardless of complexity); and finally, we also computed efficiency as $\frac{\text{Normalized Score}}{\text{Duration}}$. Problem-solving time, and relatedly, problem-solving efficiency, are natural performance metrics for which one may wish to optimize for under some circumstances. For example, how to quickly come up with a reasonably good plan for resource allocation in disaster response is such a problem that requires taking the problem-solving time into consideration when measuring performance.

We defined a number of independent variables that capture various possible influencing factors of collective performance: 1) *group-level skill*: the average value of three group members’ skills; 2) *group-level social perceptiveness level*: the average value of three group members’ social perceptiveness level; 3) *skill diversity*: the variance of the three group members’ skills; and 4) *cognitive style diversity*: “homogeneous” groups comprised all individuals of the same cognitive style, as defined above, whereas diverse groups comprised different types (see SM Section 3.1 for more details on our independent and dependent variables).

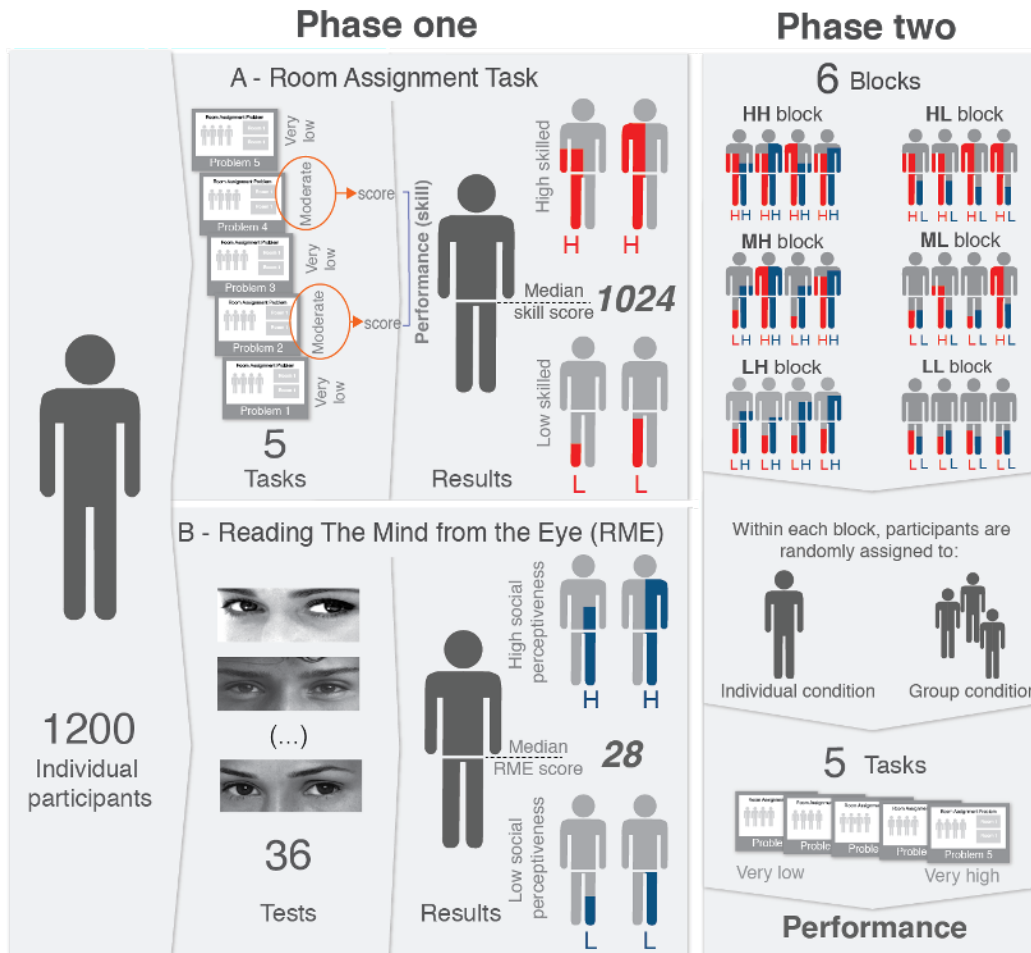


Figure 2. Schematic illustration of the study design. In phase one, participants completed a sequence of the “room assignment” task and a standard “Reading the Mind in the Eyes” (RME) test. In phase two, the same participants were assigned to blocks and randomized into “individual” or “group” conditions before performing the second sequence of five room assignment tasks

Results

Performance as a function of task complexity

Figure 3 shows how performance varied as a function of task complexity. Across all individuals and groups (see also Figures S6-S8), higher task complexity resulted in lower normalized scores (3A), longer duration (3B), and hence lower efficiency (3C). Although the direction of these results is unsurprising, the large and roughly linear dependency of three performance measures on complexity validates our design, in which overall complexity is manipulated by varying one or more task parameters (N,M,Q).

Moreover, the ability to vary human-experienced complexity by such substantial margins (on average, individuals and groups spent roughly three times as much work time on the most complex task as the least complex task, but obtained normalized scores that were roughly ten percentage points lower, which is 50% of the effective range; see caption Fig. 3) allows us to test for interaction effects with task complexity where theories of collective performance have been largely silent: that is, to what extent does the relative performance of groups versus individuals as well as optimal group compositions depend on the characteristics of the task being performed? Alternatively, one can view varying complexity as a robustness check on findings obtained for any single task (34).




Figure 3. Varying the room assignment task complexity. Increasing the task complexity reduces the normalized score (A), increases the time (B) required to complete the task, and reduces efficiency (C). Data is combined across both individual and group conditions across all 6 blocks. Error bars indicate the 95% confidence intervals. Groups and individuals scored at least 80% of the max score in over 85% of tasks, hence the effective range for the normalized score (i.e., the y-axis of A) is between 80% and 100%. The minimum time required for a solution to be submitted is one minute and the maximum is 10 min, hence effective range for the duration (i.e., the y-axis of B) is between 1 and 10 min. The difference in experienced difficulty between very low and very high complexity is very large: the average normalized score dropped by about 50% of the effective range of scores (from roughly 95% to 85% on an effective scale of 80-100), and the average time taken increased by 200% (from 2 min to 6 min).

Groups versus nominal groups and individuals

Figure 4 compares overall standardized collective performance (mean-centered within each task complexity level) with two natural benchmarks: (a) a “comparable individual,” defined as a randomly drawn individual from the same block; and (b) a “nominal group,” constructed by drawing three individuals randomly and without replacement from the same block, and then choosing the individual with the highest skill (i.e., CSOP score) from phase one. Nominal groups, therefore, simulate a situation in which a manager assigns the work to the best worker, as judged by past performance (i.e., phase one scores). The dependent variable used then is again each individual/group’s phase two performance, as measured by the normalized score, duration, and efficiency.

For all levels of task complexity, Figure 4A shows that groups score higher than comparable individuals overall ($P = 0.030$; see SM Table S4a) but lower than nominal groups ($P = 0.002$; see SM Table S4a), consistent with longstanding findings that nominal groups outperform real groups under various circumstances (8); although, the effects are small. Interestingly, however, Figure 4B shows that groups complete the most complex tasks—but not simpler ones—faster than either comparable individuals or nominal groups, suggesting that for tasks with many components (students and rooms) and many constraints the benefits of distributing work to a group outweigh the process losses (e.g., motivation loss, coordination cost) associated with groups (7). Finally, Figure 4C shows that for complex tasks the gains in speed exceed the deficits in score, resulting in a striking interaction between task complexity and configuration with respect to efficiency: for simple tasks groups are considerably less efficient than either individuals or nominal groups, yet they are considerably more efficient than either for the most complex tasks. This result is reminiscent of group decision making among social insects where a recent study has found that ant colonies outperform individual ants when the discrimination task is difficult but not when it is easy (35) (see also Figure S9-S11 for results broken down by block and Tables S4-S6 for statistical tests).

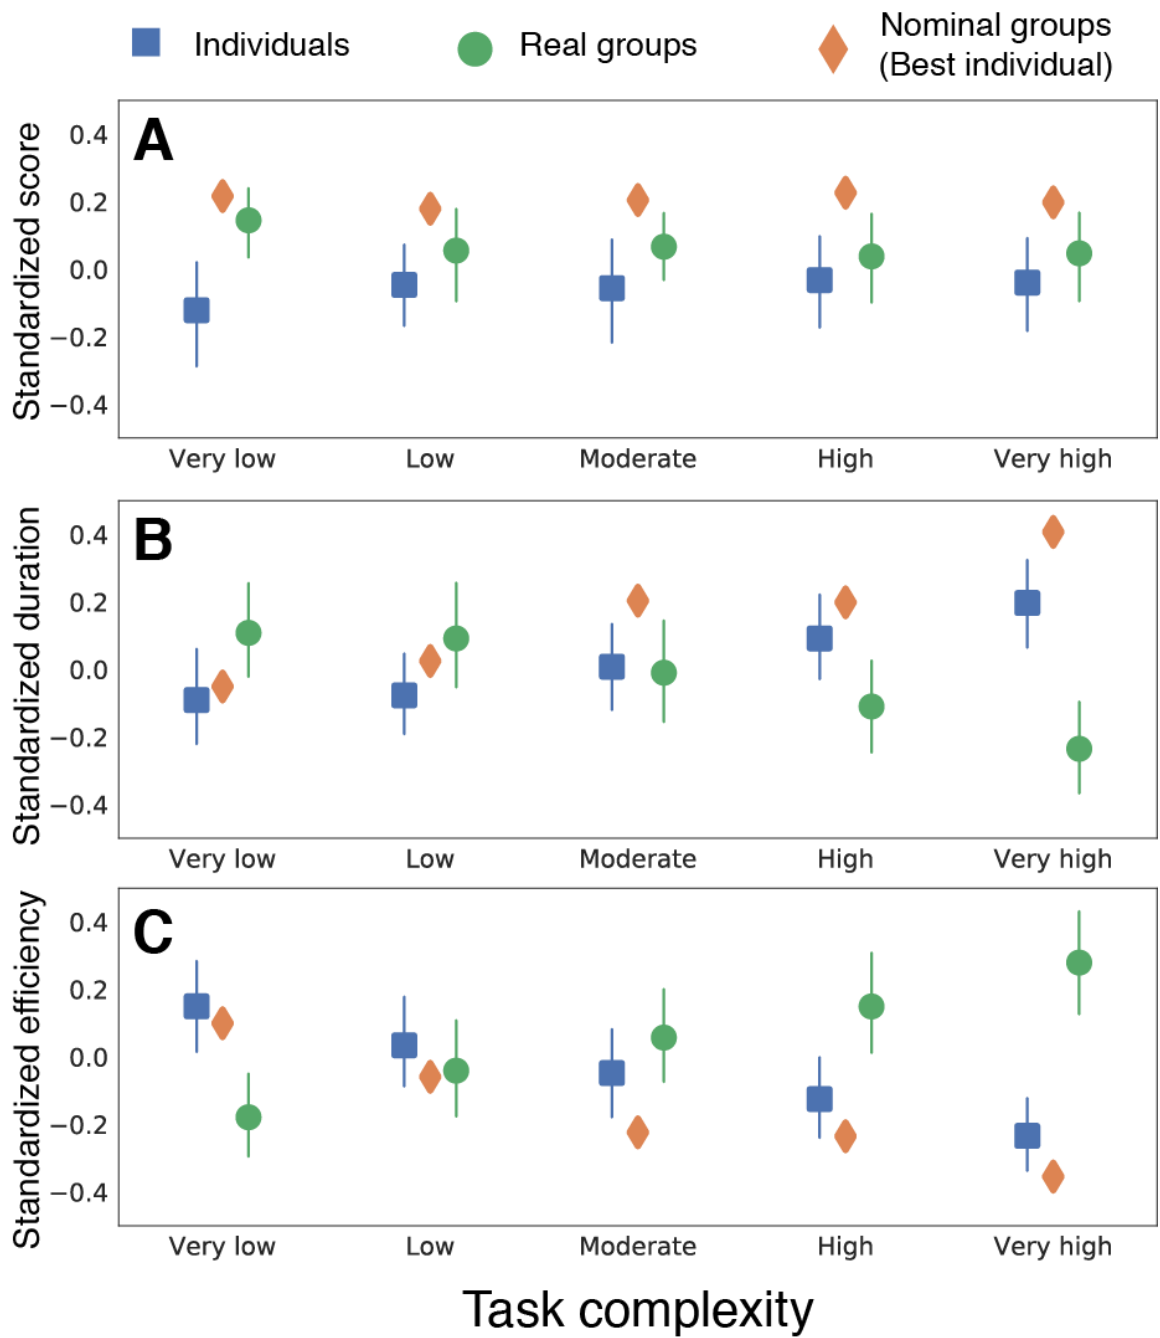


Figure 4. Comparing performance across individuals, real groups, and nominal groups. Data is combined across all 6 blocks and standardized (i.e., mean-centered) within each task complexity level (differences are relative within complexity, and should not be compared across complexity levels). Error bars indicate the 95% confidence intervals.

Group composition and collective performance

Focusing now on groups only, Figure 5 shows the absolute and relative effects of all pre-registered independent variables on collective performance, which is quantified as normalized score (Figure The 5A), duration of completion (Figure 5B), and efficiency (Figure 5C), respectively (and all three metrics are standardized within each task complexity level as per our pre-registration; see SM Section 3.4). Across all complexity levels, Figure 5A shows that average skill level had the largest effect on groups' scores, and was both positive and highly significant (Table S7). In addition, the effect of skill is consistently and significantly larger than that of social perceptiveness (Wald chi-square test; $\chi^2 = 6.35$, $P = 0.012$), which was also positive and significant (Table S8). In contrast, skill diversity (i.e., variance in group members' ability) has consistently and significantly negative effects on the score (Tables S9-S10, Figure S12) while cognitive style diversity has no consistent and significant effect (Tables S11-S15, Figure S13). Compared with group score, the effects of skill, social perceptiveness, and diversity on duration (Figure 5B) and efficiency (Figure 5C) are small and not significant at the $P < 0.05$ level (Tables S7-S15).

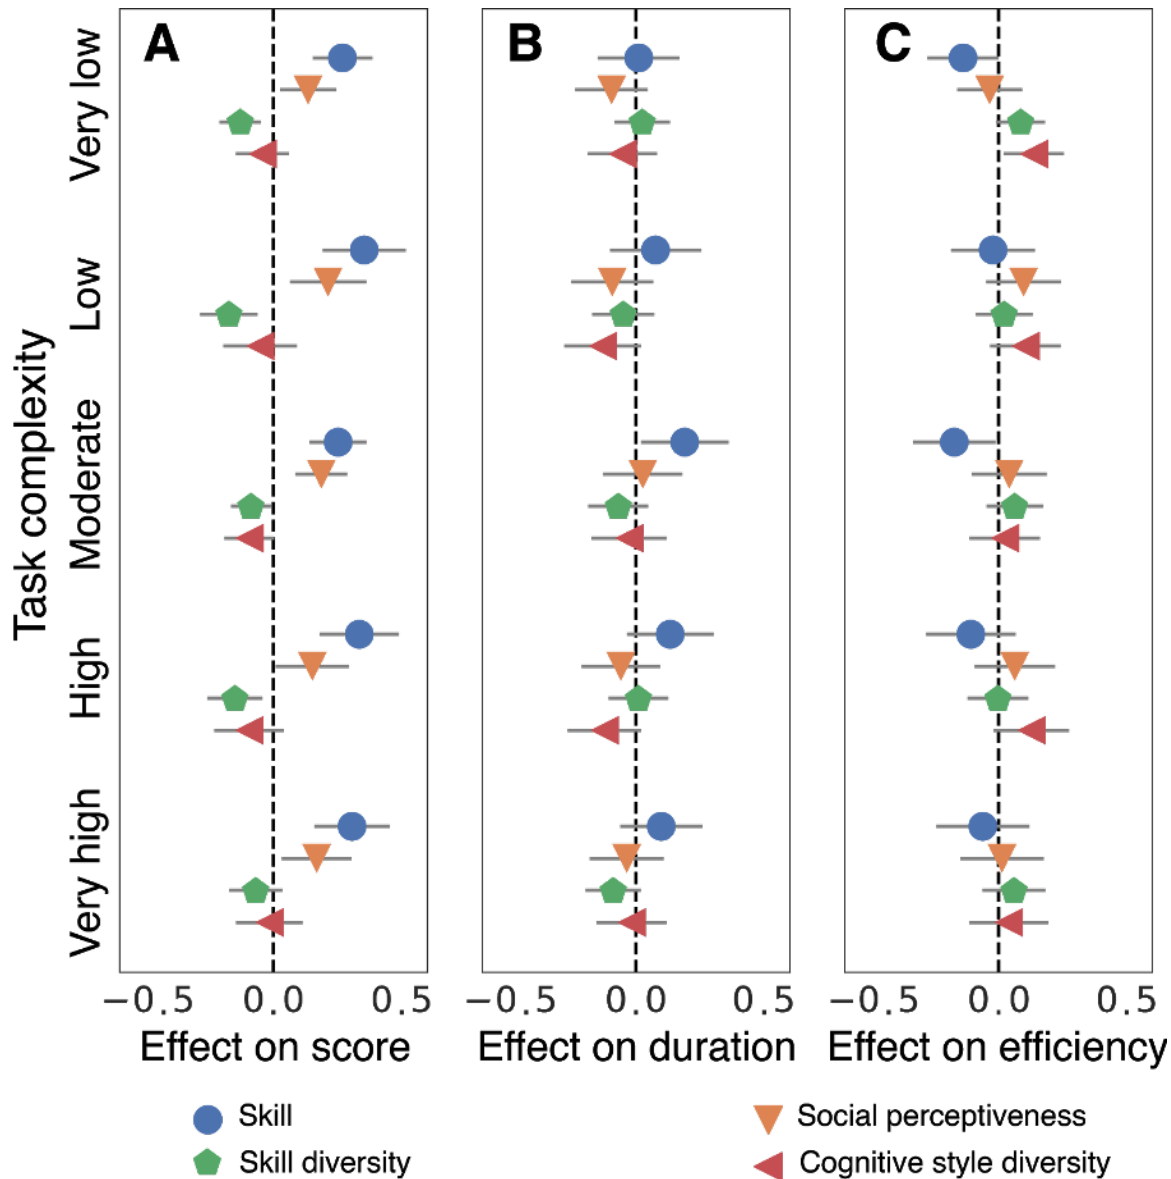


Figure 5. Group composition and collective performance. Standardized regression coefficients for skill, social perceptiveness, skill diversity, and cognitive style diversity as a function of task complexity when predicting (A) normalized score, (B) duration, and (C) efficiency. Error bars indicate the 95% confidence intervals. All regression coefficients are standardized within each task complexity level as per our pre-registration (see SM Section 3.4). See Tables S7-S15 for statistical tests (mixed effect model), and Figures S12-S13 for additional analyses on the effects of skill/cognitive style diversity.

Out-of-sample prediction of collective performance.

Standardized regression coefficients are helpful for comparing effect sizes; however, in practice, we may also care about predictive accuracy (36, 37). To illustrate, consider a hypothetical manager who wishes to compose a group for some task, and who has prior information about the skill, cognitive style, and social perceptiveness of prospective group members. In essence, the manager's task is to predict the collective performance of different combinations of individual traits. More specifically, the manager cares about two related questions. First, what is the predictive accuracy of his or her "model" (i.e., how much of the observed variance can be accounted for by all independent variables in combination)? Second, what fraction of overall predictive performance is accounted for by each independent variable? The answer to the first question quantifies the extent to which collective performance depends on the observed individual traits (versus unobserved traits, factors external to the individuals, and random noise, etc.), and hence to what extent it can be "engineered" at all. The answer to the second question indicates which of the observed variables to prioritize when selecting group members. The latter is particularly important when there is a cost associated with the measurement of the relevant variables.

Addressing the first question concerned with the degree to which collective performance depends on the observed individual traits, Figure 6A shows the out-of-sample R^2 for a simple linear regression model where the dependent variable is the cumulative normalized score (i.e., summed over all tasks), and all observed independent variables are included first independently (i.e., separate, univariate regressions; blue diamond symbol) and then cumulatively (orange square symbols) in order of increasing independent out-of-sample explanatory power (i.e., the R^2 of the corresponding univariate regression). Overall, the R^2 was approximately 0.24, meaning that the model "explained" about 24% of the out-of-sample observed variance in collective performance (more complex machine learning models scored similarly; see Figure S14). This figure is intermediate between recent attempts to predict individual life-course outcomes ($0.03 \leq R^2 \leq 0.23$) (38) and attempts to predict the size of Twitter cascades ($R^2 \approx 0.4$) (39).

Addressing the second question related to which of the observed variables to prioritize when selecting group members, Figure 6B shows cumulative out-of-sample R^2 for the same model but starting with the most explanatory variable (i.e., skill level) and adding variables in order of

decreasing explanatory power. Although social perceptiveness and skill diversity do visibly increase out-of-sample R^2 , these improvements are even smaller than one would surmise from the corresponding regression coefficients in Figure 5A: skill alone corresponds to an R^2 of 0.19 (or 80%) of all out-of-sample explained variance, a figure that compares favorably with earlier meta-analytical studies of general cognitive ability (13), which could explain only 8% of the within-sample variance in collective performance (out-of-sample performance would almost certainly be lower). In other words, our hypothetical manager could predict her group's performance almost as well knowing only skill as she could with all other variables combined (as expected, predicting duration and efficiency were considerably harder, see Figs S15-S16 for details).

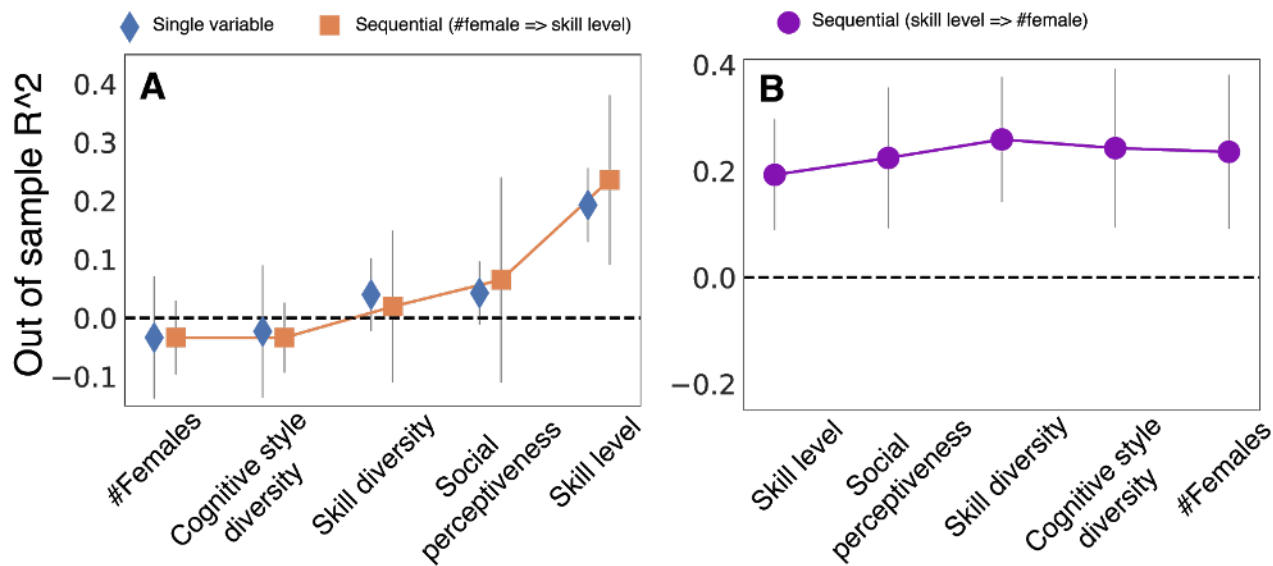


Figure 6. Out-of-sample prediction performance. Using linear regression (70% training, and 30% testing; randomized and repeated 5 times) to predict groups' cumulative normalized score with groups' skill level, skill diversity, social perceptiveness, cognitive style diversity, and the number of female group members. (A) Compares predictive performance for covariates regressed independently (i.e., in separate models; blue diamonds), and in a single model where covariates are added in order of increasing independent predictive performance (orange squares). (B) Predictive performance for a single regression model where covariates are added in order of decreasing independent predictive performance (purple circles). Error bars indicate 95% confidence intervals.

Discussion

These results provide mixed support for previous studies and also highlight some important building blocks (e.g. the two-stage design, block randomization, relative and out-of-sample predictive performance) from which a program that resolves these inconsistencies could be constructed. First, our results replicates prior findings regarding the effectiveness of groups vs. individuals: whereas we find that groups outperform comparable individuals selected at random, consistent with (1), we also find that groups score worse than the best individual selected from a nominal group of the same size, consistent with (7, 8). Although one might have expected that the relative performance, in terms of normalized score, of groups and individuals would be different across such large differences in task complexity levels; however, we find no evidence for such interaction. Interestingly, even as real groups (i.e., teams) underperform nominal groups in terms of score, for the most complex tasks—but not for simpler tasks—they attain higher efficiency by completing their work faster.

Second, our finding that the effects of average individual skill and social perceptiveness are positive and highly significant is consistent both with the aforementioned meta-analytical studies that favored ability (13–15) and also with the more recent experiments that emphasized social perceptiveness (3, 16–18). However, our ability to compare effect sizes and predictive performance across multiple effects resolve the apparent inconsistency between the two sets of results: skill dominates social perceptiveness by an order of magnitude. That is, the on-task skill of group members—accounting for roughly 80% of explained variance—far outweighs other factors, such as skill diversity, cognitive style diversity, and social perceptiveness, that have been emphasized in recent years. Third, our findings that (a) skill diversity is negatively associated with collective performance and (b) cognitive style diversity measures are neither positively nor negatively associated with performance appear to contradict widely cited claims regarding the performance benefits of diversity (19, 20). Naturally, the lack of diversity effect might not generalize to all types of tasks, groups can be diverse with respect to attributes other than skill and cognitive style (e.g., demographics, specialized skills, worldview, etc.), and diversity can affect outcomes other than performance on task (e.g., satisfaction, legitimacy, social equity, etc.); thus our results should not be construed as finding no effect of diversity in general. Nevertheless, they

add to other recent results (40) that positive performance effects of diversity are surprisingly difficult to detect in carefully controlled empirical studies.

We note that the results of laboratory experiments, including ours, rarely translate directly into the real world. Obtaining results of immediate practical relevance would require running a far more extensive and complicated series of experiments than the one we have presented, in which we would vary the available time, group size, task type, group interaction parameters, and many other potentially moderating variables. Nonetheless, our experiment is more realistic than previous work in one important sense: that when some hypothetical manager is faced with a situation where she must select individuals, about which she has some prior information, to combine into a group, it is effectively out-of-sample predictive performance that she is seeking to maximize. In other words, if the problem we have studied did arise in a real-world context then the quantity we are measuring would be exactly the quantity that a hypothetical manager would care about. In this sense, our work exemplifies a “solution-oriented” approach (41): by forcing theoretical conjectures to confront the sort of practical questions that a manager trying to assemble a team might ask, our objective has been to advance basic understanding of collective problem-solving. In concrete terms, our approach accomplishes this goal through a unique combination of (a) a class of tasks that exhibits well-defined and easily manipulated complexity; (b) a two-phase design in which we measure individual on-task skill and cognitive style as well as social perceptiveness prior to group assignment; (c) a large sample size and block randomization to increase power; (d) a focus on out-of-sample predictive accuracy; and (e) a comparison of relative effect sizes and predictive performance across multiple attributes simultaneously.

An obvious limitation of our study, however, is that although our results were extremely robust with respect to task complexity, we only studied one type of task. However, the value of our novel approach is that it offers a template for a research program that focuses on out-of-sample prediction performance and leverages the power of “virtual lab” platforms to execute complex experiments with multiple treatments. In future work, we hope to apply the same approach to qualitatively different tasks (e.g., creative problem-solving vs process improvement) as well as varying other parameters of interest (e.g., group size, incentives, communication patterns, type of group interactions, division of labor, leadership, etc.). A research program that systematically varied task types along with other contextual factors would advance the basic science of collective problem-solving while also addressing practical applications. Although such a program

is well beyond the scope of this paper (or any single paper), our paper introduces some important building blocks (e.g. the two-stage design, block randomization, relative and out-of-sample predictive performance), which in combination with emerging “open science” practices such as pre-registration, data availability, open code, and “many-labs” style collaborations, offer a promising route forward.

Materials and Methods

Human participant approval

The study was reviewed by the Microsoft Research Ethics Advisory Board and approved by the Microsoft Research Institutional Review Board (MSR IRB; Approval#: 0000019). All participants provided explicit consent to participate in this study and MSR IRB approved the consent procedure.

Data and code availability

All of the data, analysis code and the pre-registration plan are publicly available at the Open Science Framework (OSF #c6sqb) repository. Our main hypotheses, experimental design, and analyses were pre-registered before the collection of the data (AsPredicted #13123).

The experiment was developed using Empirica (<https://empirica.ly/>), an open-source “virtual lab” framework and a platform for running multiplayer interactive experiments and games in the browser (42). The source code for the “Room Assignment” task can be found at <https://github.com/amaatouq/room-assignment>, and the source code for “Reading the Mind in the Eye” (RME) test can be found at https://github.com/amaatouq/RME_test.

Standardized Coefficients

We note that in our analysis, we use the original (continuous) scores for individuals and groups that we obtained from phase one (where a group’s score is the average of the group members’ scores), not the block labels that were used for the group assignment by differentiating high-skilled/low-skilled (or high social perceptiveness/low social perceptiveness) individuals and groups. To estimate the effect size of each influencing factor (i.e., skill, social perceptiveness, skill diversity, cognitive style diversity) on collective performance for tasks of varying complexity levels, we fitted linear regressions on the collective performance data obtained on each of the 5 room assignment task instances that we included in phase two. Specifically, given a

task instance (e.g., the high complexity task in phase two), a particular influencing factor (e.g., social perceptiveness), and a performance metric (e.g., normalized score), we regressed each group's standardized performance in that task on the group's standardized value of the chosen influencing factor to estimate how this factor affects collective performance. Standardization of performance metric values and influencing factor values is conducted within the chosen task. For example, given a task instance t at a particular complexity level, we pooled the normalized scores obtained by all groups on t , and a group's standardized score on t can be computed as (the group's normalized score on task t minus the mean of all groups' normalized scores on task t) divided by the standard deviation of all groups' normalized scores on task t . Such standardization enables meaningful comparisons of effect sizes across tasks of different complexity levels.

Out-of-sample prediction procedure

The out-of-sample R^2 we reported in Figure 6 was obtained through the following procedure: we randomly split the entire set of groups that participated in our phase two experiment into a 70% training set on which the predictive model was learned, and we then tested the model on the remaining 30% testing set. This procedure was repeated 5 times, and we illustrated the mean R^2 and 95% confidence intervals for each model in Figure 6. In addition to simple linear regression models, we also considered more sophisticated models including elasticNet and randomForest, and we obtained similar results shown in Figure S10. Finally, Figure S11 shows the results of using various influencing factors to predict a group's cumulative duration on tasks (i.e., the sum of the group's duration of each of the five task instances of phase two) through different predictive models, and it was shown that compared to normalized scores that a group could obtain, how much time the group would spend on tasks is substantially more difficult to predict using various variables of group composition.

Block Randomization

To illustrate the value of the block randomization scheme used in phase two of the experiment, consider the frequency of HH individuals in the population is $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$; hence, under simple random assignment, the expected frequency of all HH groups would be $\left(\frac{1}{4}\right)^3 = \frac{1}{64}$. Of the 1,200 participants who were qualified for phase two, 828 participants entered the experiment, of which 237 were assigned to the individual condition (the data for 3 individuals was incomplete; hence the effective number of individuals is 234) and 591 were assigned to the group condition (of the

197 groups formed, the data for 1 group was incomplete, hence the effective number of groups is 196). In the absence of block randomization, we would expect to have $\frac{196}{64} = 3$ All-HH groups. With block randomization, we are guaranteed at least 22 All-HH groups (because of random assignment in the MH block it is possible that one or more additional All-HH groups would result). Put another way, to generate 22 All-HH groups with simple random assignment we would have required $22 \times 64 = 1408$ groups or over 4,000 participants just for the groups' condition (6,000 in total). Another benefit brought by the block randomization scheme is to decrease the level of self-selection bias by effectively oversampling the subgroups of participants who were potentially underrepresented in phase two, compared to the pool of participants in phase one (e.g., participants who had a lower skill and social perceptiveness level).

Statistical analysis

Because each group (or individual) completed the five room assignment tasks, we conducted tests for differences across conditions at the task level. Then, to further check the robustness of our results, we modeled the data using a generalized linear mixed model for each standardized outcome (i.e., score, duration, and efficiency) with a random effect for the group or individual identifier. These models account for the nested structure of the data. All statistical tests were two-tailed.

References

1. S. Wuchty, B. F. Jones, B. Uzzi, The increasing dominance of teams in production of knowledge. *Science*. **316**, 1036–1039 (2007).
2. B. F. Jones, S. Wuchty, B. Uzzi, Multi-university research teams: shifting impact, geography, and stratification in science. *Science*. **322**, 1259–1262 (2008).
3. A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, T. W. Malone, Evidence for a collective intelligence factor in the performance of human groups. *Science*. **330**, 686–688 (2010).
4. B. Bahrami, K. Olsen, P. E. Latham, A. Roepstorff, G. Rees, C. D. Frith, Optimally interacting minds. *Science*. **329**, 1081–1085 (2010).
5. S. Mukherjee, Y. Huang, J. Neidhardt, B. Uzzi, N. Contractor, Prior shared success predicts victory in team competitions. *Nature Human Behaviour*, 1 (2018).
6. L. Wu, D. Wang, J. A. Evans, Large teams develop and small teams disrupt science and technology. *Nature* (2019), doi:10.1038/s41586-019-0941-9.
7. N. L. Kerr, R. S. Tindale, Group performance and decision making. *Annu. Rev. Psychol.* **55**, 623–655 (2004).
8. L. L. Thompson, E. R. Wilson, in *Emerging Trends in the Social and Behavioral Sciences*, R. A. Scott, S. M. Kosslyn, Eds. (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2015), vol. 89, pp. 1–14.

9. A. Koriat, When are two heads better than one and why? *Science*. **336**, 360–362 (2012).
10. A. Mao, W. Mason, S. Suri, D. J. Watts, An experimental study of team size and performance on a complex task. *PLoS One*. **11**, e0153048 (2016).
11. P. R. Laughlin, J. Adamopoulos, Social combination processes and individual learning for six-person cooperative groups on an intellectual task. *J. Pers. Soc. Psychol.* **38**, 941–947.
12. J. A. LePine, Team adaptation and postchange performance: effects of team composition in terms of members' cognitive ability and personality. *J. Appl. Psychol.* **88**, 27–39 (2003).
13. D. J. Devine, J. L. Philips, Do Smarter Teams Do Better: A Meta-Analysis of Cognitive Ability and Team Performance. *Small Group Research*. **32**, 507–532 (2001).
14. G. L. Stewart, A Meta-Analytic Review of Relationships Between Team Design Features and Team Performance. *J. Manage.* **32**, 29–55 (2006).
15. S. T. Bell, Deep-level composition variables as predictors of team performance: a meta-analysis. *J. Appl. Psychol.* **92**, 595–615 (2007).
16. M. P. Lillis, Emotional Intelligence, Diversity, and Group Performance: The Effect of Team Composition on Executive Education Program Outcomes. *Journal of Executive Education*. **6**, 4 (2013).
17. D. Engel, A. W. Woolley, L. X. Jing, C. F. Chabris, T. W. Malone, Reading the Mind in the Eyes or reading between the lines? Theory of Mind predicts collective intelligence equally well online and face-to-face. *PLoS One*. **9**, e115212 (2014).
18. Y. J. Kim, D. Engel, A. W. Woolley, J. Y.-T. Lin, N. McArthur, T. W. Malone, in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17* (ACM Press, New York, New York, USA, 2017), pp. 2316–2329.
19. L. Hong, S. E. Page, Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc. Natl. Acad. Sci. U. S. A.* **101**, 16385–16389 (2004).
20. S. E. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies - New Edition* (Princeton University Press, 2008).
21. N. Ellemers, F. Rink, Diversity in work groups. *Current Opinion in Psychology*. **11**, 49–53 (2016).
22. B. K. AlShebli, T. Rahwan, W. L. Woon, The preeminence of ethnic diversity in scientific collaboration. *Nat. Commun.* **9**, 5163 (2018).
23. I. Aggarwal, A. W. Woolley, Team Creativity, Cognition, and Cognitive Style Diversity. *Manage. Sci.* (2018), doi:10.1287/mnsc.2017.3001.
24. J. Bendor, S. E. Page, Optimal team composition for tool-based problem solving: BENDOR and PAGE. *J Econ Manage Strat.* **90**, 34 (2018).
25. R. E. Wood, Task complexity: Definition of the construct. *Organ. Behav. Hum. Decis. Process.* **37**, 60–82 (1986).
26. J. P. Simmons, L. D. Nelson, U. Simonsohn, False-positive psychology undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychol. Sci.* **22**, 1359–1366 (2011).
27. E. Tsang, *Foundations of Constraint Satisfaction: The Classic Text* (BoD – Books on Demand, 2014).
28. A. Barreto, M. de O. Barros, C. M. L. Werner, Staffing a software project: A constraint satisfaction and optimization-based approach. *Computers & Operations Research*. **35** (2008), pp. 3073–3089.
29. J. M. Balmaceda, S. N. Schiaffino, J. A. D. Pace, Using constraint satisfaction to aid group formation in CSCL. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial*. **17**, 35–45 (2014).

30. L. Ingolotti, A. Lova, F. Barber, P. Tormos, M. A. Salido, M. Abril, in *Advances in Applied Artificial Intelligence* (Springer Berlin Heidelberg, 2006), pp. 400–409.
31. J. Shore, E. Bernstein, D. Lazer, Facts and Figuring: An Experimental Investigation of Network Structure and Performance in Information and Solution Spaces. *Organization Science* (2015).
32. D. H. Wolpert, W. G. Macready, No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*. **1** (1997), pp. 67–82.
33. S. Baron-Cohen, S. Wheelwright, J. Hill, Y. Raste, I. Plumb, The “Reading the Mind in the Eyes” test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *J. Child Psychol. Psychiatry*. **42**, 241–251 (2001).
34. B. Baribault, C. Donkin, D. R. Little, J. S. Trueblood, Z. Oravecz, D. van Ravenzwaaij, C. N. White, P. De Boeck, J. Vandekerckhove, Metastudies for robust tests of theory. *Proc. Natl. Acad. Sci. U. S. A.* **115**, 2607–2612 (2018).
35. T. Sasaki, B. Granovskiy, R. P. Mann, D. J. T. Sumpter, S. C. Pratt, Ant colonies outperform individuals when a sensory discrimination task is difficult but not when it is easy. *Proc. Natl. Acad. Sci. U. S. A.* **110**, 13769–13773 (2013).
36. J. M. Hofman, A. Sharma, D. J. Watts, Prediction and explanation in social systems. *Science*. **355**, 486–488 (2017).
37. T. Yarkoni, J. Westfall, Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspect. Psychol. Sci.* **12**, 1100–1122 (2017).
38. D. E. Rigobon, E. Jahani, Y. Suhara, K. AlGhoneim, A. Alghunaim, A. “sandy” Pentland, A. Almaatouq, Winning Models for Grade Point Average, Grit, and Layoff in the Fragile Families Challenge. *Socius*. **5**, 2378023118820418 (2019).
39. T. Martin, J. M. Hofman, A. Sharma, A. Anderson, D. J. Watts, “Exploring Limits to Prediction in Complex Social Systems” (Proceedings of the 25th International Conference on World Wide Web 978–1–4503–4143–1, International World Wide Web Conferences Steering Committee, Montreal, Quebec, Canada, 2016), pp. 683–694.
40. A. H. Eagly, When passionate advocates meet research on diversity, does the honest broker stand a chance? *J. Soc. Issues*. **72**, 199–222 (2016).
41. D. J. Watts, Should social science be more solution-oriented? *Nature Human Behaviour*. **1**, 0015 (2017).
42. N. Paton, A. Almaatouq, Empirica: Open-Source, Real-Time, Synchronous, Virtual Lab Framework (2018), doi:10.5281/zenodo.1488413.

Supplementary Material for

Collective Problem-Solving of Groups Across Tasks of Varying Complexity

Abdullah Almaatouq¹, Ming Yin², Duncan J. Watts^{3,4}

Correspondence to: amaatouq@mit.edu, mingyin@purdue.edu, duncan@microsoft.com

¹Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA, 02139.

²Department of Computer Science, Purdue University, 305 N. University Street, West Lafayette, IN 47907.

³Microsoft Research, 641 Ave of the Americas, 7th Floor, New York, NY 10011.

⁴The University of Pennsylvania, Philadelphia, PA 19104

S1. Access to data, code, and analysis plan	1
1.1 Exceptions to the pre-registration plan	1
S2. Details of Experimental Setup	2
2.1 Room assignment problem	2
2.2 Reading the mind in the eyes	5
2.3 Screenshots of the instructions & comprehension check	6
2.4 Experimental Design	11
2.4.1 Participants recruitment	12
2.4.2 Design of phase one experiment	12
2.4.3 Validity of participant's individual skill measure	16
2.4.4 Design of phase two experiment	17
2.4.5 Comparing participants in phase one and two	22
S3. Details of Analysis	24
3.1 Independent and dependent variables	24
3.2 Performance as a function of the complexity	26
3.3 Comparing the performance of groups and individuals	28
3.4 Comparing effects on collective performance	37
3.5 Testing robustness of the group composition results	42
3.6 Predicting collective performance	52

S1. Access to data, code, and analysis plan

All of the data, analysis code and the pre-registration plan are publicly available at the Open Science Framework (OSF) repository¹. Our main hypotheses, experimental design, and analyses were pre-registered before the collection of the data (AsPredicted #13123)². The study was reviewed by the Microsoft Research Ethics Advisory Board and approved by the Microsoft Research Institutional Review Board (MSR IRB; Approval#: 0000019).

The experiment was developed using Empirica (<https://empirica.ly/>), an open-source “virtual lab” framework and a platform for running multiplayer interactive experiments and games in the browser¹. The source code for the “Room Assignment” task can be found at <https://github.com/amaatouq/room-assignment>, and the source code for “Reading the Mind in the Eye” (RME) test can be found at https://github.com/amaatouq/RME_test.

S1.1 Exceptions to the pre-registration plan

The sample sizes used in this study constitute a minor deviation from our pre-registration, where we intended to get 1200 participants in phase one, and 780 participants in phase two (with 540 of them placed into the group condition forming 180 groups, and the rest 240 of them placed into the individual condition). However, there is stochasticity involved in recruiting participants from Amazon Mechanical Turk. For example, amongst the workers who agreed to take part in our experiment we cannot tell a priori how many would pass the attention checks, or how many would do both tasks required in phase one (i.e., the room assignment task described in section 2.1 and the RME test described in section 2.2).

Therefore, we recruited more than 1250 participants in phase one and we ended up with 1211 valid participants who completed both tasks (+11 from our pre-registered sample size).

However, to honor our pre-registration we only qualified the first 1200 to take part in phase two of the experiment. In phase two, amongst the 828 participants who showed up and passed the attention checks, 591 participants were assigned to a group condition and ended up forming 197 groups (+16 from our pre-registered sample size; data for one group is incomplete, leading to the number of valid groups to be 196), and 237 participants were

¹ https://osf.io/c6sqb/?view_only=6ae8d48a5de74623a631135e8b4b755e

² AsPredicted pre-registration <http://aspredicted.org/blind.php?x=73tw6s>

assigned to an individual condition (-3 from our pre-registered sample size; data from three individuals is incomplete, leading to the number of valid individuals to be 234). The complexity with obtaining exactly the predetermined number of groups (i.e., 180) and individuals (i.e., 240) was our inability to know with certainty how many participants would show up for phase two at the appointed time and the uncertainty associated with random assignment schemes; thus we opted to err in favor of more groups than needed rather than fewer. As the precise number of groups was not important to our hypotheses and as the block randomization was implemented as intended, we do not believe this deviation had any impact on our results.

S2. Details of Experimental Setup

S2.1 Room assignment problem

In our experiments, we asked participants to solve room assignment problems, first individually and then potentially within a group of three members (i.e., based on whether the participants is assigned to “individual” or “group” condition). A room assignment problem is a type of Constraint Satisfaction and Optimization Problem (CSOP, that is, an optimization problem on top of a constraint satisfaction problem), which is an abstraction of many real-world complex problems such as resource allocation and scheduling and thus a subject of intensive research in artificial intelligence and operational research ^{2,3}. In our case, participant(s) were tasked with assigning each of N “students” to one of M “rooms,” while also respecting Q constraints on their choices (e.g., students A and B must be neighbors, must not share a room, etc.). In each room assignment problem, a “utility table” was presented, providing participant(s) with the information on students’ ratings/preferences (between 0 and 100) to each of the M rooms indicating how satisfied they would be if being assigned to the room. The participant(s) was then asked to find a room assignment plan that maximized satisfaction across all students without violating the constraints.

To incentivize the search for an optimal solution (i.e., the optimal room assignment plan), we provided participant(s) with additional bonuses based on how good their submitted solutions for the problem were. In particular, we defined the “score” of a room assignment plan as the following:

*Score = The sum of students' ratings of their assigned rooms - 100 * the number of violated constraints*

By submitting a complete plan (that is, each student got assigned to one room) with a positive score in a room assignment problem, participant(s) could earn a “performance-based bonus” using a 500 points:\$1 USD conversion rate to exchange scores into payments. Moreover, if the submitted plan was indeed the optimal one, an additional \$0.5 USD “optimal assignment bonus” would be given³. We determined these values for the payments by conducting a series of pilot studies and observing how participants responded to different payment schemes. For an interface for the room assignment task that is solved by a group, see Figures S1-S2.

³ These bonus rates are for phase one experiment. In phase two, the performance-based bonus conversion rate is 1000 points:\$1, while the optimal assignment bonus is \$0.7. We set these bonus rates to maintain a similar level of hourly payment between the phase one and two experiment.

A Very low complexity

Timer
09:51

Score
0

Room 101

Room 102

Room 103

Room 104

Payoff

Rooms	101	102	103	104
Student A	10	47	32	20
Student B	78	65	46	37
Student C	35	59	41	53
Student D	40	65	12	43
Student E	18	39	40	78
Student F	22	51	57	40

Constraints

- B and E must be neighbors.
- C and F can't live in the same room or be neighbors.

⋮

B Very high complexity

Timer
09:51

Score
0

Room 101

Room 102

Room 103

Room 104

Room 105

Room 106

Room 107

Room 108

Rooms	101	102	103	104	105	106	107	108
Student A	27	51	31	25	47	45	41	54
Student B	44	40	11	62	46	22	27	53
Student C	24	17	38	42	32	64	58	41
Student D	50	26	57	28	54	11	17	32
Student E	47	38	51	12	16	34	59	48
Student F	37	63	48	61	24	28	57	39
Student G	77	61	12	72	66	27	53	35
Student H	68	44	35	58	61	11	24	45
Student I	31	51	16	47	38	58	35	66
Student J	54	55	61	68	24	46	70	42
Student K	66	29	63	74	35	60	43	50
Student L	15	63	19	43	51	28	34	40
Student M	43	50	36	21	17	56	39	47
Student N	33	50	68	44	29	43	54	36
Student O	17	32	45	68	54	16	33	53
Student P	30	34	17	44	49	25	16	18
Student Q	49	29	61	37	39	20	62	48
Student R	19	25	28	50	54	37	42	20

Constraints

- A and B must be neighbors.
- A and J can't live in the same room or be neighbors.
- B and F must be neighbors.
- ✗ B and H must live in the same room.
- B and P can't live in the same room or be neighbors.
- ✗ C and E must be neighbors.
- C and J can't live in the same room.
- D and F can't live in the same room or be neighbors.
- E and Q can't live in the same room or be neighbors.
- F and I can't live in the same room.
- ✗ G and J must be neighbors.
- I and N can't live in the same room or be neighbors.
- J and K can't live in the same room or be neighbors.
- ✗ K and R must be neighbors.
- L and Q can't live in the same room or be neighbors.
- M and O must live in the same room.
- N and R must live in the same room.
- O and P must be neighbors.

Figure S1. Two instances of the “room assignment” task used in the experiment. (a) An instance with "very low" complexity in which six students are to be assigned to four rooms subject to two constraints. (b) An instance with "very high" complexity in which 18 students are to be assigned to 8 room subject to 18 constraints.

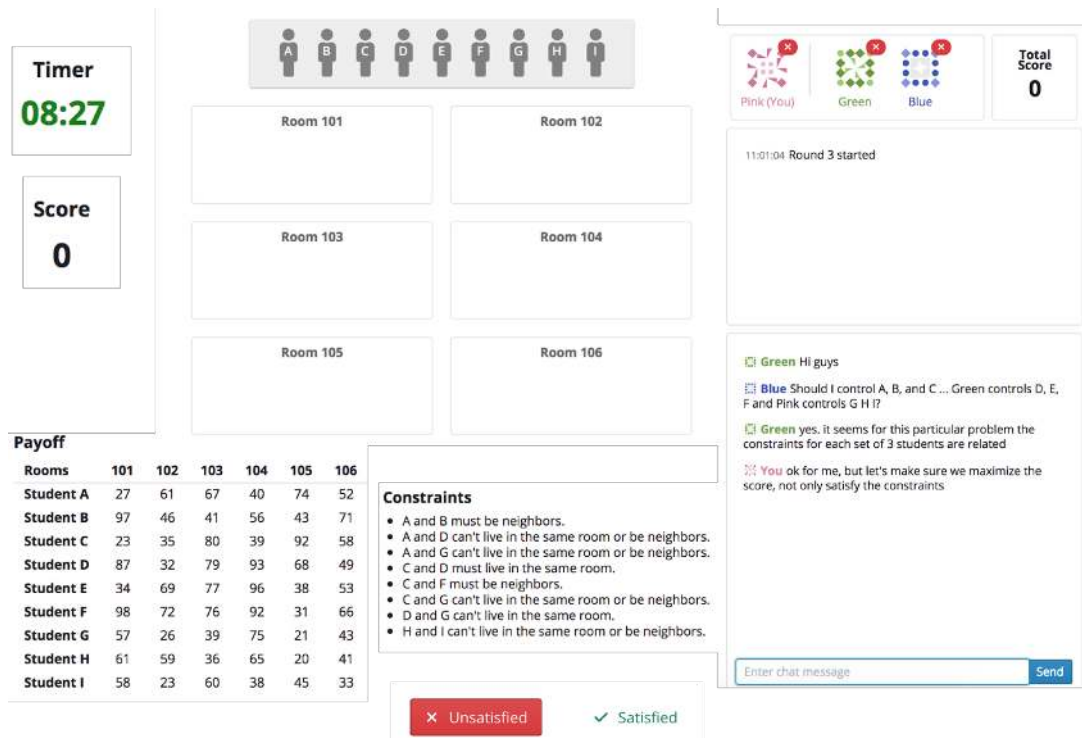


Figure S2. An illustration of phase two “room assignment” task that was done by a group of three individuals

2.2 Reading the mind in the eyes

Each participant also completed a revised version of the “Reading the Mind in the Eyes” test⁴, a widely used test for measuring Social Perceptiveness/Emotional Intelligence. In this test, participants are shown 36 pairs of eyes. For each pair of eyes, they are provided with four words describing emotions. The participant is asked to select one of the four words that best describe what the person in the picture is thinking or feeling. See Figure S3 for an illustration of the test.




Figure S3. An illustration of the “Reading the Mind in the Eye” test used in phase one of the experiment. The participant is shown a pair of eyes and asked to choose the emotion that best describes what the individual in the picture is feeling or thinking of.

2.3 Screenshots of the instructions & comprehension check⁴

Game Overview

In this game, you will be **asked to solve a sequence of resource allocation tasks**. In each task, you are going to **assign a group of students into dorm rooms**. You are asked to find the room assignment plan that maximizes overall satisfaction for the group while respecting certain constraints (e.g., some students can not live together in one room).

You have at most **10 minutes** to work on each task. Completing the entire game may take you as long as 60 minutes. **If you do not have at least 60 minutes available to work on this HIT please return it now.**

You will play this game simultaneously with 2 other participants in real-time. As we will explain in more detail later, in each task, you and your teammates will submit a single room assignment plan.

At the end of the game, you will have the opportunity to earn a bonus payment and the amount is dependent on your accumulated score in all tasks. **Note that “free riding” is not permitted.** *If we detect that you are inactive during a task, you will not receive a bonus for that task.*

The game **must be played on a desktop or laptop**. There is NO mobile support

For the best experience, please maximize the window containing this task or make it as large as possible.

« Previous Next »

⁴ The instruction screens titled “You will be part of a group” and “Event logs and In-Game Chat” were shown only to participants who were assigned to the group condition in Phase 2.

Room Assignment Tasks

In each task (or round), you will be asked to **assign students to dorm rooms**. Students express their degree of satisfaction for living in a room as a number between 0 and 100 (the higher the rating, the more satisfied the student is).

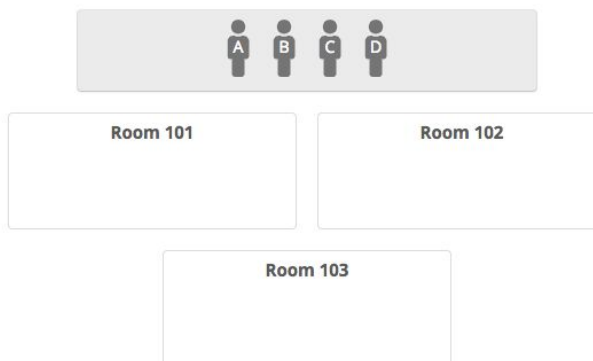
You are provided with a handy **drag and drop** tool to solve the problem. To assign a student into a room, drag the icon of that student and drop it into the room. Try this example:

Payoff

Rooms	101	102	103
Student A	20	80	65
Student B	67	90	76
Student C	85	82	79
Student D	20	75	78

Score

N/A



NOTE: ALL the students HAVE to be assigned to a room in order for your score to count.

« Previous [Next](#) »

Respecting the Constraints

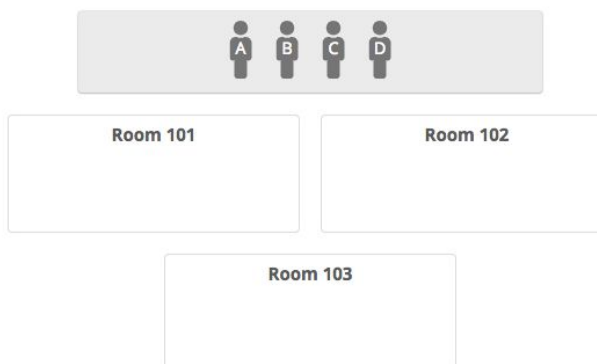
You need to **consider some constraints when assigning students to rooms**. Some students can't live together in the same room and some students must be neighbors.

These constraints vary from task to task, and there are no additional constraints you need to respect other than the ones stated (e.g., feel free to leave one room empty if no constraint requires you to assign at least one student in each room).

Try this example again and see what will happen if a constraint is violated:

Constraints

- A and B can't live in the same room or be neighbors.
- B and C must live in the same room.



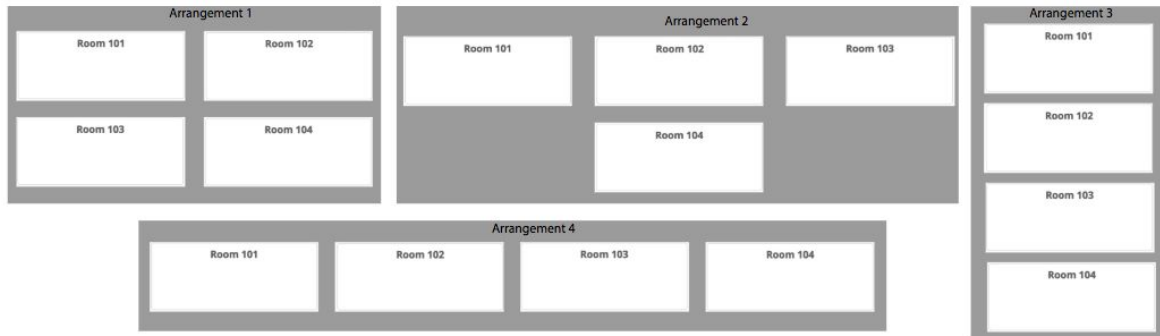
NOTE: Every violated constraint will result in deducting 100 points from your score.

NOTE: It is OK to leave some rooms empty, but you have to assign all the students.

« Previous [Next](#) »

Task Room Arrangements

Depending on the number of rooms, number of students, and your screen/browser size and resolution, the arrangement of the rooms might "look" different on your screen.




In all cases and for any arrangement that appears for you, you only need to consider the numbers on those rooms when addressing constraints in a task. In particular, "**neighbor**" is defined as rooms with consecutive numbers. For example, regardless of the arrangement you have on the screen, Room 102 is next door to both Room 101 and Room 103. On the other hand, Room 101 is only next door to Room 102.

« Previous [Next](#) »

You will be part of a team

In this game, you will **play together with 2 other participants (your teammates)**. They are other MTurk workers who are undertaking the same study simultaneously. Throughout all the tasks, the team will submit only one answer, and therefore, **all members of the team will receive the same score**. To help you identify yourself and differentiate each other in the team, we will assign a color to you when the game starts (as shown in the following example).



Note that the game allows for simultaneous and real-time actions. That means that you will be able to drag students to assign them to rooms while your teammates are doing the same. However, when any member in the team starts dragging a student, that student will be locked (i.e., no one else can move it) until it is assigned to a room. **The student that is being moved will have the color of the participant.**

« Previous [Next](#) »

Event Logs and In-Game Chat

We will log every action taken by you or any of your teammates, and this log will be shown to you to help you keep track of all the actions that have taken place so far.

Also, you may communicate with your teammates through the in-game chat. This chat room is public so whatever you write will appear to the other 2 teammates. You can use this in anyway you want.

Remember, you and your teammates have 10 minutes in each task to find a room assignment plan. You will automatically **progress to the next task when the time is up**.

However, you can always indicate whether you are satisfied with the answer before the timer is up (indicated by the check mark on the avatar). Click on the "Satisfied" button in the following example and see what happens!



If all team members are satisfied with the answer before the timer is up, the answer will be submitted and your team will proceed to the next task. If the "Satisfied" button is unclickable (i.e., inactive) for you for more than 10 seconds, try to refresh the page..

[« Previous](#) [Next »](#)

Scores and Bonuses

In each task, we use "score" to evaluate the quality of the room assignment plan that your team came up with. **Your score starts counting only when you have a complete assignment** (that is, each student has been assigned to a room).

The score of your assignment is calculated as:

$$S = \text{The sum of students' ratings of their assigned rooms} - 100 * \text{the number of violated constraints}$$

That means, **for each constraint you violate, you get 100 points deducted.**

As a team, **you will submit ONE answer per task** and therefore **all team members will have the same score on each task.**

There are two parts of the bonus that you will have opportunity to earn in each task:

1. **"performance-based bonus"**: When your score is positive, no matter whether your answer is the BEST possible assignment or not. The exchange rate is **1000 game points = \$1 bonus.**
2. **"optimal assignment bonus"** : When your answer is the BEST possible assignment, you get **an additional bonus of \$0.7 in that task.**

Therefore, **big part of the bonus is for finding the BEST possible assignment** (i.e., "optimal assignment bonus", which can be up to \$3.5 total). Also, **you can earn more game points (i.e., more performance-based bonuses) from the difficult tasks** compared to the easier ones (more students/rooms means more possible bonus).

Together with your teammates, you should try to find a complete room assignment with a score that is as high as possible to earn more bonus in each task!

Remember, free riding is not permitted. If we detect that you are inactive during a task, you will not receive a bonus for that task.

« Previous [Next »](#)

Game Interface

We are almost there! please take a second to familiarize yourself with the game User Interface shown here:

Round 1 > [Response](#)

Timer: 04:53 | Score: 0

Constraints

- A and D can't live in the same room or be neighbors.
- A and F must be neighbors.
- A and G must live in the same room.

Payoff

Rooms	101	102	103	104	105	106
Student A	33	28	50	65	25	67
Student B	71	23	75	88	36	64
Student C	44	52	66	80	55	35
Student D	36	93	88	61	94	71
Student E	22	69	94	22	43	87
Student F	89	70	28	48	85	53
Student G	67	36	20	24	59	42
Student H	57	74	62	40	82	85
Student I	49	65	33	62	71	46

Room 101, Room 102, Room 103, Room 104, Room 105, Room 106

Blue (You), Pink, Green | Total Score: 0

11:36:34 Round 1 started

No messages yet...

Enter chat message | [Send](#)

UnSatisfied | Satisfied

If the "Satisfied" button for in is unclickable (i.e., inactive) for more than 10 seconds, try to refresh the page. Otherwise, you will have wait for the time run out. This will not effect your bonus.

Now you know where everything goes and ready to take the quiz! Good luck.

« Previous [Next »](#)

Quiz

How many participants will play at the same time, including yourself?

Select the true statement about the score:

- I will score points only based on the assignments that I make
- We will submit only one answer as a team and therefore we will all get the same score.

is it ok to have some rooms empty? (the answer is 'Yes')

- Yes!
- No!

If your team ended up NOT assigning all students to room (i.e., at least one student remained in the deck) then your score in that task will be:

For each unsatisfied (i.e., violated) constraint, how many points will be deducted from you?

Which of the following rooms is a neighbor of Room 101? Please select all that apply.

- Room 101 Room 102
- Room 103
- Room 104 Room 105

Which of the following rooms is a neighbor of Room 103? Please select all that apply.

- Room 101
- Room 102 Room 103
- Room 104
- Room 105

[← Back to instructions](#) [Submit ↩](#)

2.4 Experimental Design

The goal of our experiment was to examine (1) under what conditions, if any, do groups perform better than individuals, and (2) which of several factors (e.g., skill level, skill diversity, social perceptiveness level, cognitive style diversity, etc.) predicts collective performance. To answer these questions, we used a novel “two-phase” experimental design in which we recruited the same group of participants twice to solve the room assignment problems. On the high-level, phase one was used for gathering ex-ante measurements of each participant’s skill level on the room assignment problem, social perceptiveness level, and cognitive style. Then, in phase two, we deployed a block randomization scheme, randomly assigning participants into one of six blocks based on their phase one measurement results (we will describe more details below in Section 2.4.4). Within each block, we then randomly assign participants into the “individual” condition or the “group” condition. While participants assigned to the individual condition solved another set of room assignment

problems on their own, participants in the group condition formed groups of three randomly and each group collectively solved the same set of room assignment problems as those solved in the individual condition. Finally, we used the ex-ante measurements from phase one to construct the independent variables (i.e., different influencing factors of collective performance) and used the actual performance of groups/individuals in phase two as the dependent variables, which together allowed us to compare the performance of groups versus individuals and examine the performance effects of different group compositions.

2.4.1 Participants recruitment

All participants were recruited on Amazon Mechanical Turk (<http://mturk.com>, MTurk), which is an online labor market with a large and diverse pool of people ready to promptly perform tasks for pay (called human intelligence tasks, or HITs). We recruited our participants by posting a HIT for the experiment, entitled “Play games and get up to \$17 in total pay,” a neutral title that was accurate without disclosing the purpose of the experiment. The study was reviewed by the Microsoft Research Ethics Advisory Board and approved by the Microsoft Research Institutional Review Board (Approval#: 0000019). All participants provided explicit consent to participate in this study and MSR IRB approved the consent procedure. All data collected in the experiment could be associated only with the participant's Amazon Worker ID on MTurk, not with any personally-identifiable information. All participants remained anonymous for the entire study. In each phase of the experiment, participants first read instructions and could start the experiment only after they had correctly answered a set of questions testing their comprehension of the instructions (see Section 2.3 for screenshots and examples).

2.4.2 Design of phase one experiment

In phase one of the experiment, participants were asked to complete a sequence of 36 “Reading the Mind in the Eyes” (RME) test questions as well as a sequence of 5 room assignment tasks. More specifically, room assignment tasks and RME test were implemented as two distinct web apps, each of which appeared as a separate link in the MTurk iframe. The order of the links was randomized for each participant but they could choose to click on them in whatever order they wished. For the RME questions, participants were shown in each

question a pair of eyes and were asked to select one of the four words that best describe the emotions shown by the eyes (Figure S3).

For the room assignment task part, we first introduced participants to the problem and each completed one practice task (as per our pre-registration, it is not included in the analysis), in which $N=8$ students need to be assigned to $M=5$ rooms while respecting $Q=4$ constraints. Each participant was then given a sequence of five room assignment tasks, to be completed independently, where the maximum amount of time a participant could spend on a task was 5 minutes. Table S1 summarizes the main properties of the five task instances used in phase one.

Table S1. Main properties of the 5 room assignment tasks used in phase one of our experiment.

Task Order	N (# of students)	M (# of rooms)	Q (# of constraints)	Maximum possible score	Complexity Level
1	6	4	2	343	Very low
2	9	6	8	554	Moderate
3	6	4	2	323	Very low
4	9	6	8	564	Moderate
5	6	4	2	325	Very low

As shown in the table, we intentionally included 3 tasks of very low complexity and 2 moderately complex tasks in the sequence. We did *not* randomize the order of the task instances in phase one to minimize the noise in the measurement of individual skill due to random ordering effects. We included more tasks of a lower level of complexity in phase one to minimize potential self-selection in phase two of our experiments⁵ (i.e. where only participants who did well in phase one would return for phase two), which turned out to be very effective (see section 2.4.5 for more details).

⁵ For example, participants who performed well in phase one may be more likely to participate in the phase-two experiment, implying possible self-selection biases; by having more task instances of very low complexity in phase one, most participants may feel they performed well thus bias is attenuated.

When working on a room assignment task, a participant was presented with a graphical interface where each student was represented as a person icon and each room was shown as a box (see Figure S1). The participant could then drag the icons of students and drop them to different boxes to adjust the room assignment plans. Assistive information such as the score of the current room assignment plan, the list of violated constraints, and the amount of time left in the task was also displayed and updated on the interface while the participant changed the solution. At any time during the allotted 5-minute period for a task, the participant could push a button to submit her solution and move on to the next task (or to the end of the room assignment task sequence), or the participant would be automatically redirected to the next task when the 5-minute timer was up. After the participant solved all five room assignment tasks in phase one, she was asked to complete an exit survey, in which we asked her to self-report the following information:

- Age
- Gender
- Highest Education Received
 - High School
 - US Bachelor's Degree
 - Master's or higher
 - Other
- Were the instructions clear?
- Was the pay fair?
- Was the time limit per task reasonable?
- Did you encounter any problems with the user interface?
- If you had assigned all students to rooms and had no conflicts, which of the following would you be most likely to do?
 - Submit your solution and move on to the next task
 - Try to increase your score by moving students around as long as you didn't generate any new conflicts
 - Try to increase your score by moving students around even if it meant generating new conflicts
- If you had assigned some (but not all) students to rooms and had encountered one or more conflicts, would you:
 - Put off resolving the conflict(s) until all students had been assigned?
 - Stop assigning students to rooms until conflict(s) had been resolved?
 - Continue assigning students as long as no more than one conflict were present?
- When assigning a student to a room, did you focus more on
 - Which room had the highest score?
 - Which room(s) would avoid generating conflicts?
- Any other feedback?

At the end of phase one, we obtained a number of measurements for each participant:

- *Skill*: defined as the sum of the participant's score on the two moderate room assignment tasks. We only use participant's scores on the moderate tasks as moderate tasks are more discriminative and scores on moderate tasks have higher variability. However, we note that a participant's score on the two moderate tasks highly correlates with the participant's score on each of the five-room assignment tasks (see Section 2.4.3 for validity check).
- *Social perceptiveness level*: defined as the number of RME questions the participant correctly answered.
- *Cognitive style*: defined based on the observed behavior and reported problem-solving strategies of different participants:

(1) Speed ("tortoises" vs. "hares"): defined in terms of the total amount of time the participant spent on solving the moderate instances of phase one room assignment tasks: hares take less than median time; tortoises take more than median time. However, this measure is highly correlated with performance metrics.

(2) Preference for efficiency vs. perfection ("pragmatist" vs. "perfectionist"): defined in terms of the participant's self-reported answer for the exit-survey question "If you had assigned all students to rooms and had no conflicts, which of the following would you be most likely to do?": *pragmatist* (i.e., the participant chose "submit your solution and move on to the next task" or "try to increase your score by moving students around as long as you didn't generate any new conflicts") or *perfectionist* (i.e., the participant chose "try to increase your score by moving students around even if it meant generating new conflicts");

(3) Tolerance for constraint violation (low vs. high): defined in terms of participant's self-reported answer to the exit-survey question "If you had assigned some (but not all) students to rooms and had encountered one or more conflicts, what would you do?": *low* (i.e., the participant chose "stop assigning students to rooms until conflict(s)

had been resolved”) or *high* (i.e., the participant chose “put off resolving the conflict(s) until all students had been assigned” or “continue assigning students as long as no more than one conflict were present”); and

(4) Preference for optimization vs. constraint resolution (“optimizer” vs. “resolver”), which is decided by the participant’s self-reported answer for the exit-survey question “When assigning a student to a room, what did you focus more on?”: *optimizer* (i.e., the participants chose “which room had the highest score”) or *resolver* (i.e., the participant chose “which room(s) would avoid generating conflicts”).

Although our measurements of each participant’s skill and social perceptiveness level are continuous, to facilitate the block randomization scheme that we would adopt in phase two of our experiment, we further used a median split to categorize each participant into the high or low class on both measurements. For example, a participant whose skill was above the median skill while social perceptiveness was below the median level would be categorized as “high skill, low social perceptiveness.” We note that in our analysis on the performance effect of group composition, we use the original (continuous) scores for individuals that we obtained from phase one (where a group’s score is the average of the group members’ scores), not the block labels, to differentiate high-skilled/low-skilled (or high social perceptiveness/low social perceptiveness) groups. See Section 3.1 for more details.

2.4.3 Validity of participant’s individual skill measure

In our experiment, we defined an individual participant’s skill score as the sum of her scores on the two moderate tasks in phase one experiment, and we further labeled the participant as “high” or “low” on skill by examining whether her skill score was larger or smaller than the median score obtained among all participants. To illustrate the validity of this measurement of skill level, Figure S4 contrasts the normalized scores (i.e., actual score obtained in a task instance / the maximum possible score for that task instance) obtained by “high skill” participants with those obtained by “low skill” participants, on each of the six tasks in phase one, including one practice task (hard) and five actual tasks (3 tasks of very low complexity and 2 tasks of moderate complexity). Clearly, on all task instances, participants that are determined as “high skill” outperformed those participants that are determined as “low skill.” In other words, participants’ scores on the two task instances with moderate complexity are

highly correlated with their scores on any single task instance, regardless of its complexity, which suggests that it is valid to use participants' scores on the two moderate tasks to measure skill levels.

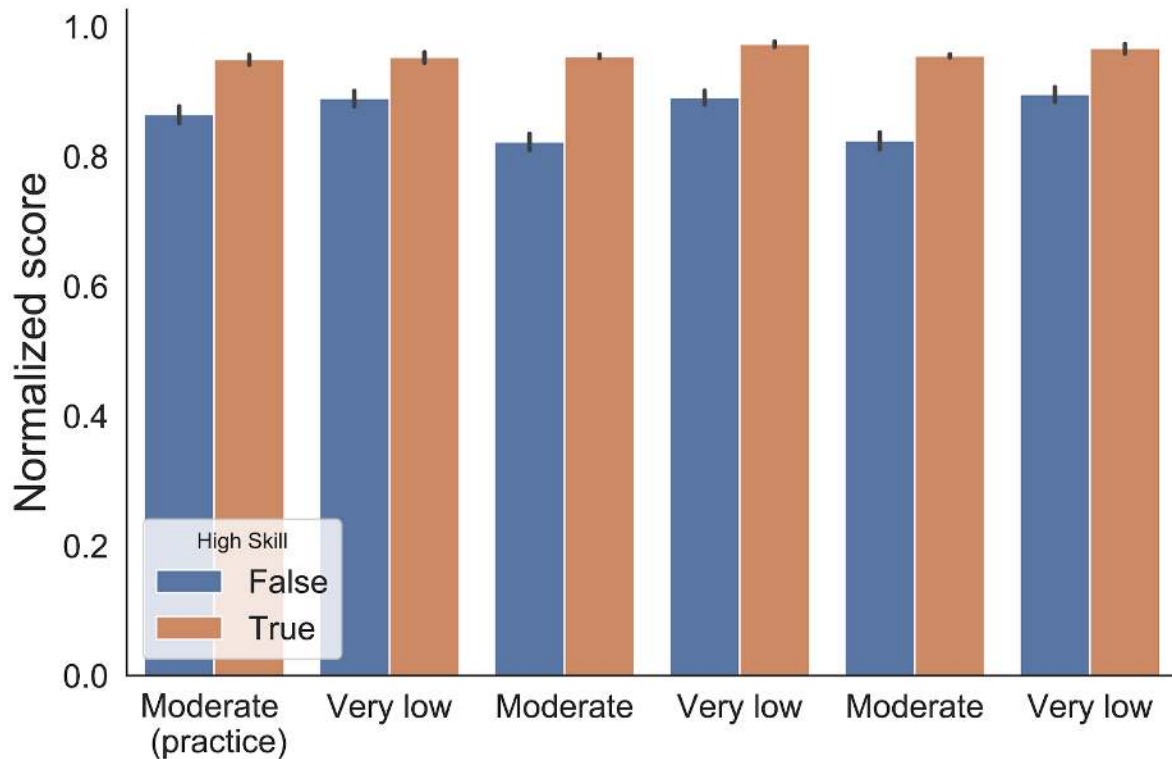


Figure S4. Participants who obtained a higher score on the two moderately complex tasks in the phase one experiment (i.e., “high skill”) outperformed participants who obtained a lower score on those two tasks (i.e., “low skill”) on each single task instance. Error bars represent 95% confidence intervals.

2.4.4 Design of phase two experiment

As per our pre-registration, we included the first 1200 participants who completed our phase one experiment into the second phase of our experiment. Among these 1200 participants, there were 313 “high skill, high social perceptiveness” (HH) individuals, 284 “high skill, low social perceptiveness” (HL) individuals, 249 “low skill, high social perceptiveness” (LH), and 354 “low skill, low social perceptiveness” (LL) individuals.

During a pilot study we conducted prior to our main experiment, we deployed a simple randomization scheme and had individuals of different levels of skills and social perceptiveness to form groups of three members at random in phase two. The majority of the

groups formed in this way contained a mixture of high/low skill (or high/low social perceptiveness) individuals. As a result, the variance of a group’s skill or social perceptiveness level (defined as the average skill or social perceptiveness level of members in that group) across different groups was limited. Practically, this implies that a large sample size would be needed to detect any statistically significant performance effect of group composition (or that the effects are unrealistically large to be detectable).

To address this problem, we adopted a block randomization scheme in phase two of our main experiment. Specifically, prior to the start of phase two, we created six qualifications on Amazon Mechanical Turk, with each qualification corresponded to a “block.” Participants of one particular block could only find and work on the HIT corresponding to their block, but not the other five HITs. Table S2 provides a summary of these six blocks.

Table S2. Summary of the six blocks that we used in phase two of our experiments.

Block name	# of participants assigned to this block	Participant decomposition	# of participants entering Phase two in this block
HH	155	155 participants all labeled as (high skill, high social perceptiveness)	100
MH	285	158 participants labeled as (high skill, high social perceptiveness); 127 participants labeled as (low skill, high social perceptiveness)	213
LH	122	122 participants all labeled as (low skill, high social perceptiveness)	90
HL	147	147 participants all labeled as (high skill, low social perceptiveness)	97
ML	310	137 participants labeled as (high skill, low social perceptiveness); 173 participants labeled as (low skill, low social perceptiveness);	221
LL	181	181 participants all labeled as (low skill, low social perceptiveness)	107

For each individual of a particular type (e.g. “high skill, low social perceptiveness” or HL), with 50% probability we assigned her to the block in which all individuals were of the same type (e.g., the “HL” block), and with 50% probability we assigned her to the block in which all individuals had the same social perceptiveness label as her, but may have different skill labels (e.g., the “ML” block, meaning “mixed skill levels, low social perceptiveness”). Within each block, we further randomly assigned participants either to the individual condition (31% of the time) or to the group condition (69% of the time). The individual condition was identical to phase one except that the five-room assignment tasks were different (and generally more complex) and that the maximum time allotted per task was ten rather than five minutes. In the group condition, participants worked in groups of three randomly selected members from the same block. Table S3 summarizes the main properties of the 5 task instances we used in our phase two experiment (the task sequence used in the individual condition is the same as that used in the group condition; the order of the task instances in the sequence was randomized for each trail).

Table S3. Main properties of the 5 room assignment tasks used in phase two of our experiment. The order of tasks was randomized in the experiment.

Task ID	N (# of students)	M (# of rooms)	Q (# of constraints)	Maximum possible score	Complexity Level
1	6	4	2	340	Very low
2	8	5	5	441	Low
3	9	6	8	672	Moderate
4	12	7	12	673	High
5	18	8	18	996	Very high

The main effect of the block randomization scheme was to oversample statistically less frequent combinations (e.g., all group members had high skills or high social perceptiveness), which helped us to increase the statistical power of our experiments (a secondary benefit was that it allowed us to match the distributions of participant types in phases one and two; see Section 2.4.5). To illustrate, the frequency of HH individuals in the population is $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$

hence under simple random assignment the expected frequency of all HH groups would be $(\frac{1}{4})^3 = \frac{1}{64}$. Of the 1,200 participants who were qualified for phase two, 828 participants entered the experiment and 237 of them placed in the individual condition (the data for 3 of them was incomplete; hence the effective number of individuals is 234) and 591 placed in the group condition. Of the 197 groups formed, the data for 1 group was incomplete, hence the effective number of groups is 196. In the absence of block randomization, therefore, we would expect to have $\frac{196}{64} = 3$ All-HH groups. With block randomization, the number of groups formed in HH, MH, LH, HL, ML, LL blocks was 22, 55, 18, 21, 56, 24, respectively. In this way, we guarantee at least 22 All-HH groups from the HH block (because of random assignment in the MH block it is possible that one or more additional All-HH groups would result). Put another way, to generate 22 All-HH groups with simple random assignment we would have required $22 \times 64 = 1408$ groups or over 4,000 participants just for the group's condition (6,000 in total). Note that we did not block on participants' cognitive styles, as doing so would require a much larger sample size.

During phase two experiment, each individual/group first completed one practice task ($N=9$, $M=6$, $Q=8$). Then, they could proceed to complete the sequence of room assignment tasks of various levels of complexity; each task had a maximum time limit of 10 minutes (unlike phase one, which had a time limit of 5 minutes), and the task order was randomized to account for any ordering effects (recall that the task order was *not* randomized in phase one to eliminate noise in measuring individual skill level). Participants in the individual condition were presented with an identical interface as that used in phase one experiment. At the end of phase two of the experiment, participants in the individual condition were asked to complete the same exit survey they completed in phase one. Therefore, we can confirm that our specific measure of cognitive style (“optimizer” vs “resolver”) has good test-retest reliability of (0.72 overall; 0.78 for the highly-skilled participants). We also confirm that the other measures of cognitive style that we reported in the supplementary materials also pass the test-retest reliability (0.74 for constraint violation tolerance; and 0.71 for preference for efficiency vs. perfection).

Participants in a group, on the other hand, were presented with an interface where all group members can drag any icon of students to any room cell simultaneously as they wish (see

Figure S2 for an example of task interface in the group condition). To avoid conflicts, when one group member was moving a student icon, that particular student icon was “locked” and other group members could not move it until it was released. We provided a chatbox on the task interface, enabling group members to communicate freely with each other during the tasks. We also presented an event log on the task interface to help group members make sense of all movements that had been made within the current task. At any time during a task, each group member could indicate whether she was satisfied with the current solution using a toggle button. Once all three members of a group indicated they were satisfied with the solution, the group would move on to the next task (or to the end of the experiment). If the group had never unanimously suggested they were satisfied with the solution, the group would automatically be redirected to the next task when the 10-minute timer was up.

At the end of phase two of the experiment, participants in the group condition were asked to complete an exit survey⁶, in which we asked the following:

- How would you describe your strategy in the game?
- Do you feel the pay was fair?
- How satisfied are you with the outcome of the game?
 - Extremely satisfied (1) --- Extremely dissatisfied (7)
- Do you think your group worked well together?
 - Strongly agree (1) --- Strongly disagree (7)
- How valuable do you think your perspective was to the end results?
 - Extremely valuable (1) --- Extremely invaluable (7)
- How comfortable were you in sharing your perspective through the chat?
 - Extremely comfortable (1) --- Extremely uncomfortable (7)
- Feedback, including problems you encountered.

⁶ Participants in the individual condition were presented with an exit survey that is identical to the one presented in phase one.

2.4.5 Comparing participants in phase one and two

One natural concern regarding the two-phase experimental design is whether different participants' experience in the phase one experiment will lead to a varying tendency to participating in the phase two experiment, implying potential self-selection that may result in biased experimental results.

To examine whether self-selection bias would be a substantial concern, we first conducted a pilot study, in which 42 participants (these participants were not allowed to participate in the actual study) were recruited from Amazon Mechanical Turk to complete the first version of our two-phase experiment. In this pilot study, we asked each participant to complete a sequence of 5 room assignment tasks of varying complexity levels as well as 36 RME questions in phase one. Two hours later, we invited all participants who had completed phase one to join the second-phase experiment, in which they would be randomly assigned to groups of three members and they were asked to solve another sequence of 5 room assignment tasks together with other members in their group.

Figure S5 (left panel) compares the distributions of participants who completed phase one (i.e., gray bars and curves) and phase two (i.e., blue bars and curves) of the pilot study, with respect to their skill levels (i.e., the cumulative score a worker got in the 5 room assignment tasks of the phase one experiment; top row) and their social perceptiveness levels (i.e., the number of RME questions a worker answered correctly in the phase one experiment; bottom row). It is clear that during the pilot study, participants who decided to take the phase two experiment had both higher skill levels and higher social perceptiveness levels, compared to the entire pool of participants who had completed the phase one experiment. In other words, the experimental design and procedure that we adopted during our pilot study indeed led to a degree of self-selection bias.

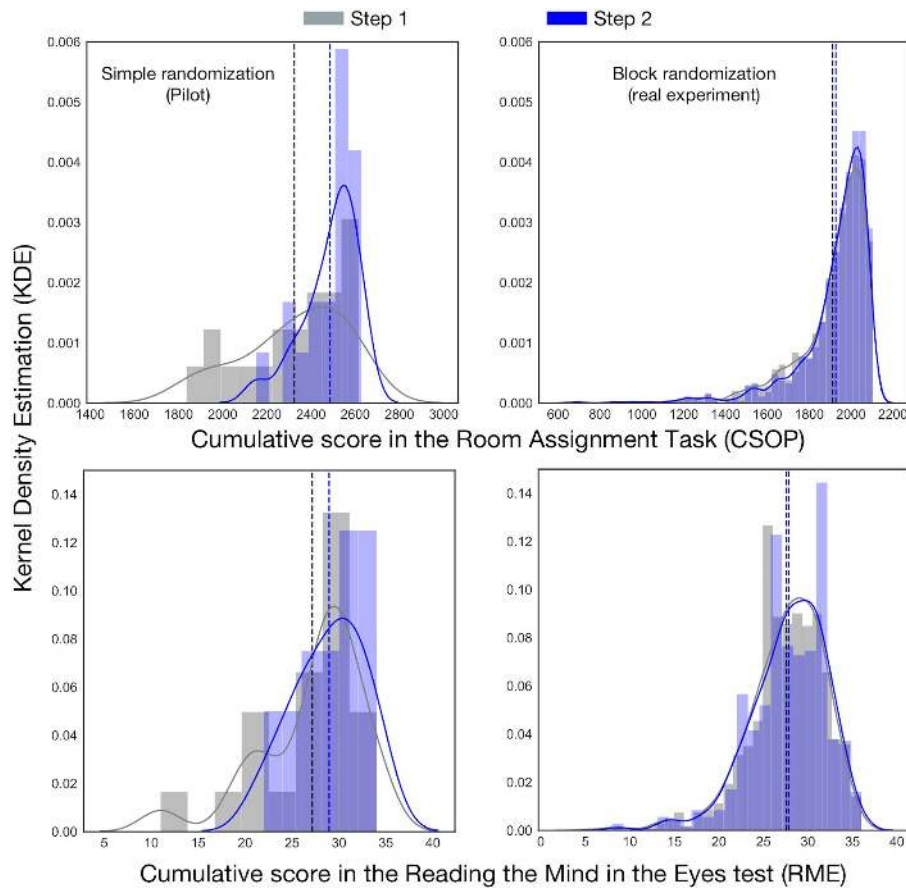


Figure S5. Comparing the distributions of phase one participants and phase two participants with respect to their skill (i.e., scores obtained in room assignment tasks) and social perceptiveness levels (i.e., scores obtained in RME tests). Left: comparison results for the pilot study; Right: comparison results for the main experiment. Gaussian kernels are used for kernel density estimation.

To decrease the level of self-selection bias, we made three changes during our main experiment. First, we altered the mix of tasks that we included in the phase one experiment to 3 tasks of “very low” complexity and 2 tasks of “moderate” complexity. We hypothesized that with a higher fraction of tasks of very low complexity in phase one, participants would have a higher perceived self-efficiency in the room assignment tasks, and thus more likely to come back during phase two to complete more such tasks. Second, we adopted a block randomization scheme rather than a simple randomization scheme during our real experiment. Each block corresponded to a particular mixture of participants with high/mixed/low skill and high/low social perceptiveness (see Section 2.4.4 for more details), and we set the target number of workers to recruit *at the block level*. Doing so allowed us to

effectively oversample the subgroups of participants who were potentially underrepresented in phase two, compared to the pool of participants in phase one (e.g., participants who had a lower skill and social perceptiveness level)⁷. Finally, we extended the gap between the two phases of our experiment from two hours to six days, conjecturing that a longer gap would refresh participants' memory and potentially lead more of them to find it enjoyable to take similar types of tasks again in our phase two experiment. Figure S5 (right panel) shows the distribution comparisons between participants who completed phase one and phase two of the real experiment. Here, we find there is no clear difference between the two groups of participants in terms of either their skill or their social perceptiveness. In other words, with the three changes that we made, we managed to minimize the self-selection biases between the two phases in our real experiment.

3. Details of Analysis

3.1 Independent and dependent variables

Our first objective is to understand under what conditions, if any, do groups perform better than individuals. To answer this question, we considered all participants who took our phase two experiment, regardless of whether they were assigned to the individual condition or the group condition. We adopted two ways to define “reference individuals” to compare group's performance against: First, given a group formed in a particular block, we considered a randomly-drawn individual from that block as a comparable individual; second, we constructed “nominal groups” by randomly drawing three individual participants without replacement from that block, and then defined the performance of the nominal group as the phase two performance of the individual who had the highest *skill score* from phase one among the three members in this group. Such nominal groups simulate a situation in which groups simply nominate their best performer to do all the work while the others contribute nothing; hence the performance of nominal groups reflect the performance of “best” individuals. We compared the performance obtained by real groups with that obtained by random individuals as well as that obtained by random nominal groups. Thus, the

⁷ As we mentioned in Section 2.4.4, another benefit brought up by the block randomization scheme is that we effectively oversampled less frequent combinations of workers (i.e., groups) even if self-selection bias was not present, such as groups with all three members being high on skill and social perceptiveness.

independent variable we used for addressing our first objective is the existence and type of group, which had three levels---no group (random individual), real group, and nominal group (best individual).

The main dependent variable then is each individual/group's performance in the second phase of our experiment. Specifically, we measured collective performance in three ways:

- *Normalized score*: the score an individual or a group obtained in a room assignment task divided by the maximum score of that task
- *Duration*: the amount of time an individual or a group spent on solving a room assignment task
- *Efficiency*: defined as $\frac{\text{Normalized Score}}{\text{Duration}}$, this performance measure is not pre-registered

Our second objective was to compare the effect of several factors (e.g., skill level, skill diversity, social perceptiveness level, cognitive style diversity, etc.) in determining collective performance. To this end, we considered only the data from those participants who were assigned to the group condition in phase two experiment, and defined a number of measures as our independent variables to capture various possible influencing factors of collective performance:

- *(group-level) skill*: the average value of three group members' skills (recall each member's skill was measured in phase one experiment as the sum of scores obtained on the two moderately complex tasks)
- *(group-level) social perceptiveness level*: the average value of three group members' social perceptiveness level (recall each member's social perceptiveness level was measured in the phase one experiment as the number of RME questions correctly answered)
- *Skill diversity*: the variance of the three group members' skills
- *Cognitive style diversity*: Given an operationalization of cognitive style, we label the group as homogeneous or diverse on that cognitive style by checking whether the

three group members in the group belong to the same type (“homogeneous”) or not (“diverse”).

The dependent variable used then is again each group’s phase two performance, as measured by the normalized score, duration, as well as efficiency (i.e., score divided by duration).

3.2 Performance as a function of the complexity

Figure 3 in the main text shows that higher task complexity led both individuals and groups to score a lower fraction of the maximum possible score (3A), work for longer (3B), and produce a lower efficiency (3C). Figures S6-S8 show that all these trends hold within each of the six blocks.

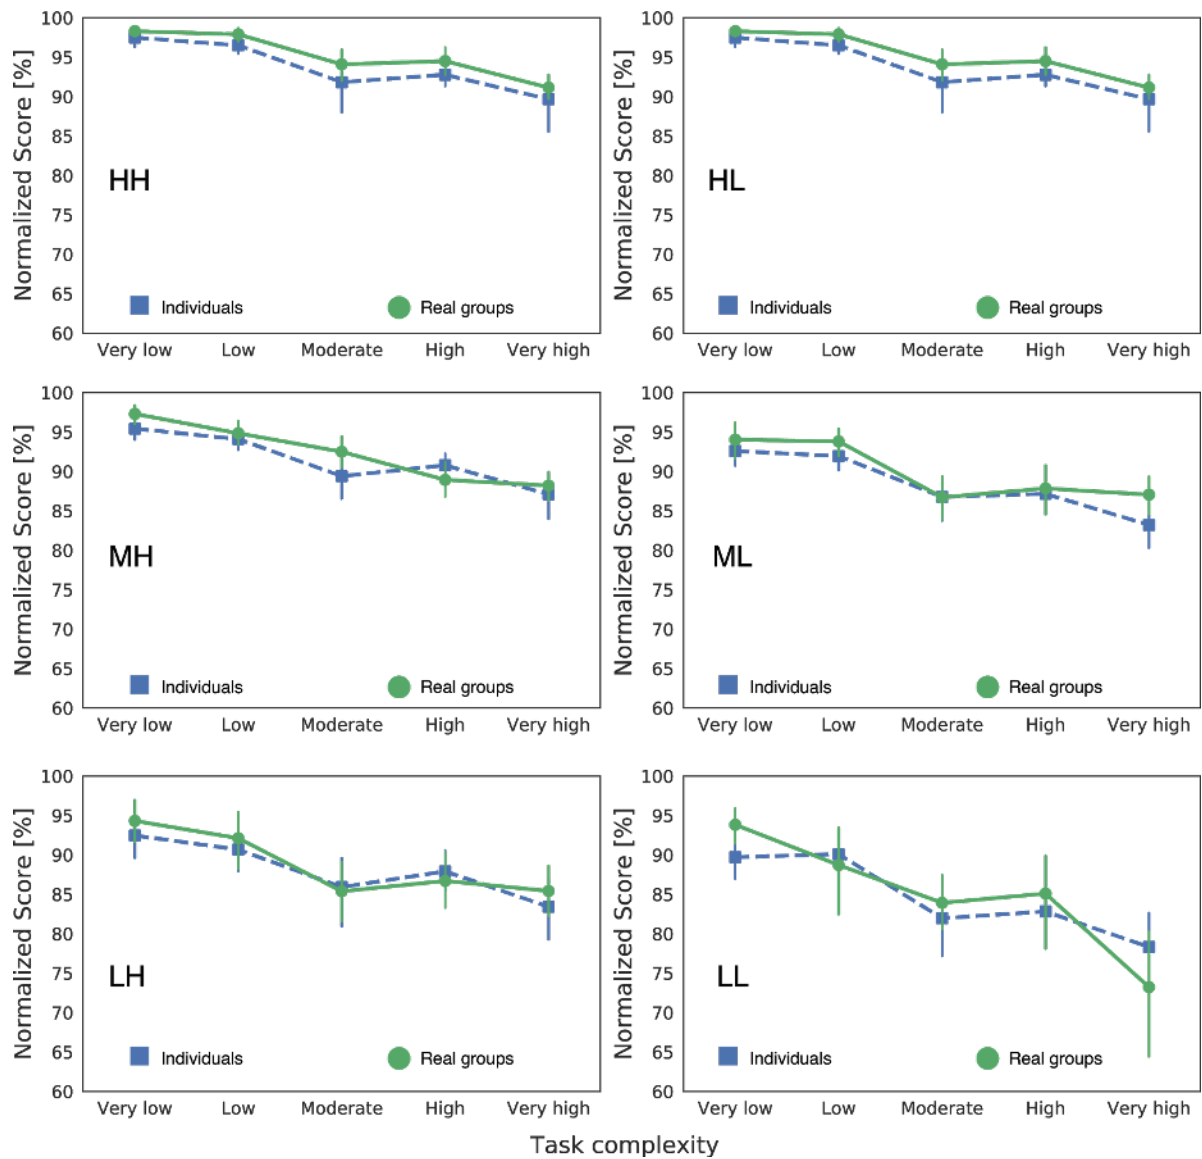


Figure S6. Varying the room assignment task complexity versus normalized score. The five task complexity levels were characterized by the different number of students to be assigned, the number of dorm rooms available, and the number of constraints. Data is analyzed separately for individuals and groups from each of the six blocks. Increasing the task complexity reduces the normalized score for both individuals and groups of all skill levels and social perceptiveness. Error bars indicate 95% confidence intervals.

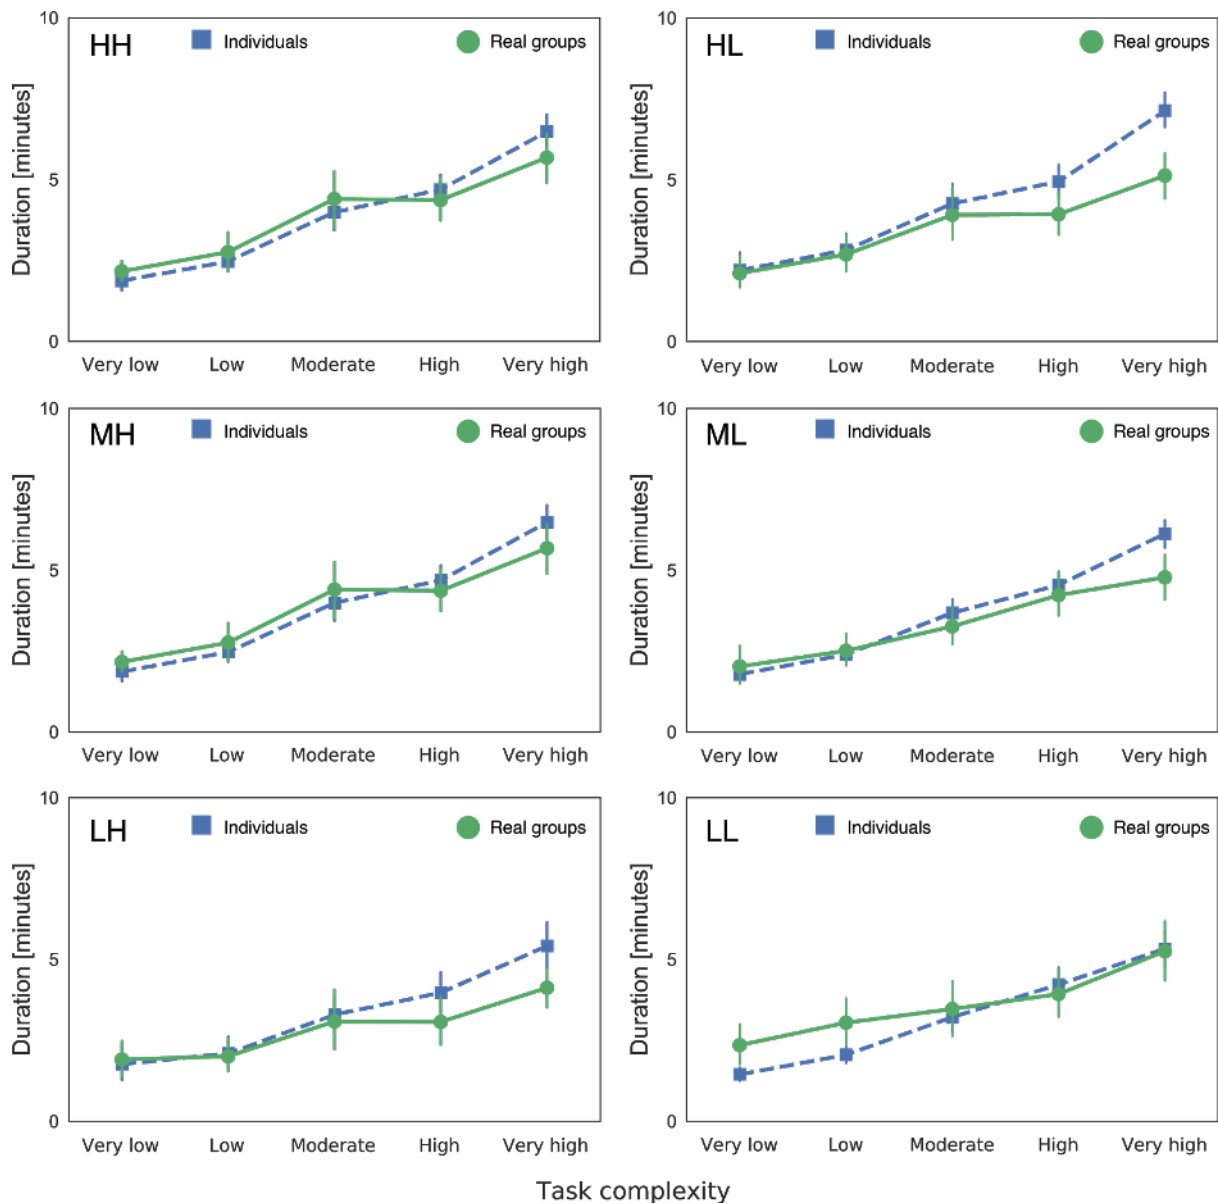


Figure S7. Varying the room assignment task complexity versus duration. The five task complexity levels were characterized by the different number of students to be assigned, the number of dorm rooms available, and the number of constraints. Data is analyzed separately for individuals and groups from each of the six blocks. Increasing the task complexity increases the time it takes participants to submit an assignment plan for both individuals and groups of all skill levels and social perceptiveness. Error bars indicate 95% confidence intervals.

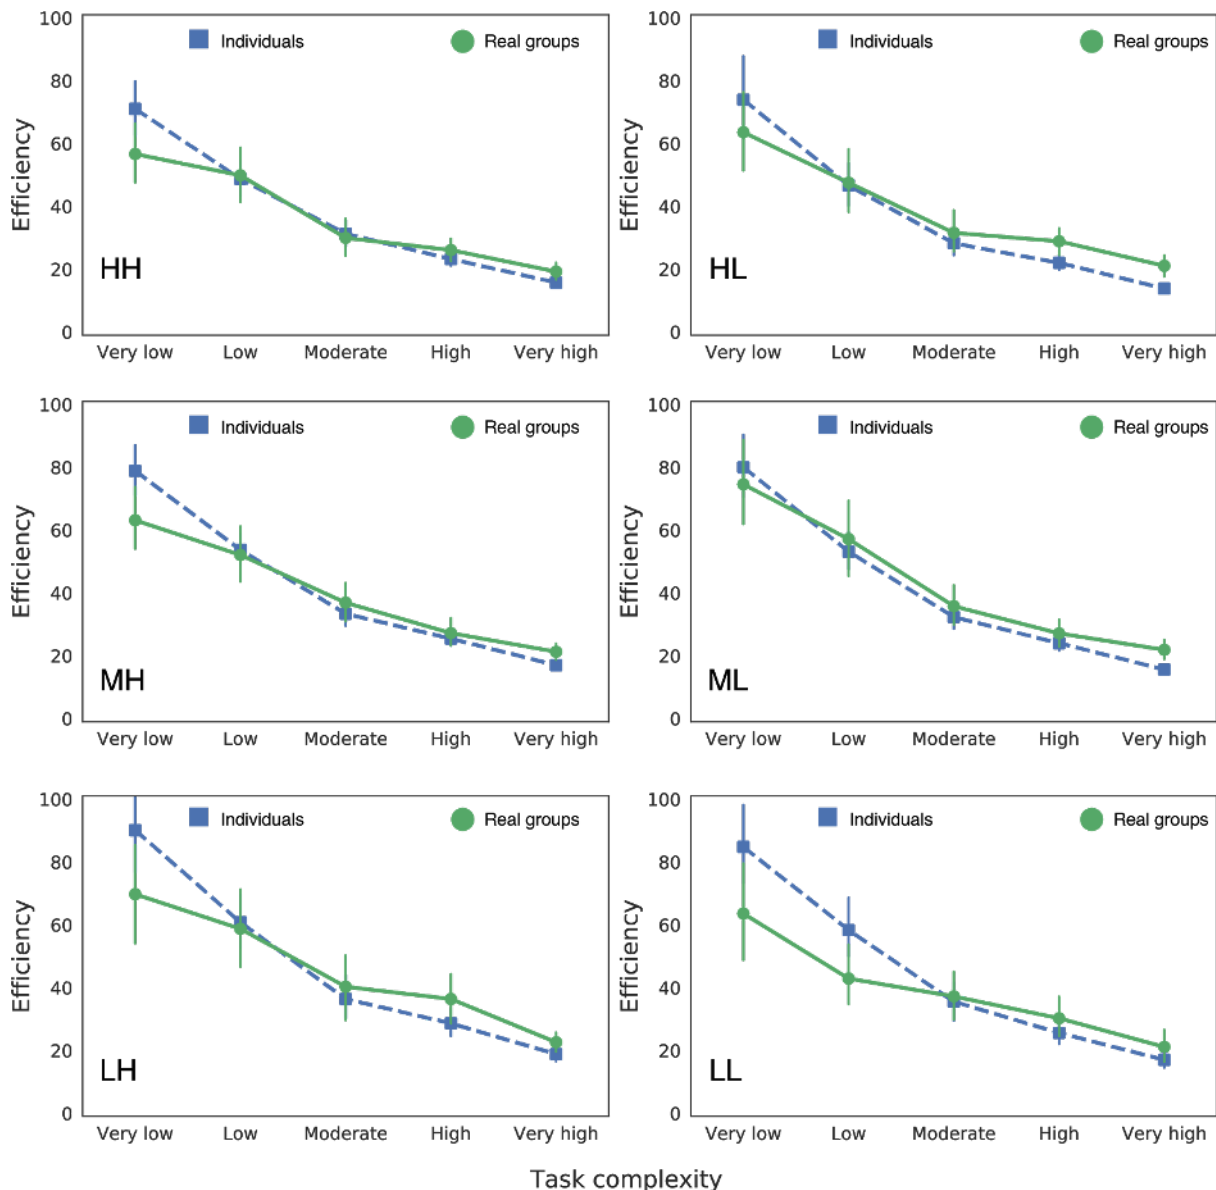


Figure S8. Varying the room assignment task complexity versus efficiency. The five task complexity levels were characterized by the different number of students to be assigned, the number of dorm rooms available, and the number of constraints. Data is analyzed separately for individuals and groups from each of the six blocks. Increasing the task complexity reduces the efficiency for both individuals and groups of all skill levels and social perceptiveness. Error bars indicate 95% confidence intervals.

3.3 Comparing the performance of groups and individuals

To understand whether and when do groups outperform individuals, we compared in each of the five task instances in the phase two experiment, the average performance produced by a real group, a randomly selected individual from the same block, and a nominal group.

Specifically, given a task instance (e.g., the task of “high” complexity in phase two), we first obtained the performance data for real groups, individuals, and nominal groups according to the method that is described in Section 3.1. Then, we standardized the performance for each of the tasks; that is, given a task instance t at a particular complexity level, we pooled the performance data (i.e., normalized scores, durations, or efficiency) obtained by all real groups, randomly selected individuals, and nominal groups on t , and a group/individual’s standardized performance on t can be computed as (the group/individual’s performance on task t - the mean value of all performance scores on task t) / standard deviation of all performance scores on task t . Such standardization enables meaningful comparisons of effect sizes across tasks of different complexity levels.

Figure 4 in the main text shows how real groups’ performance compares with that of average individuals or nominal groups. In Figure 4, given a task instance, we pooled the performance data for real groups/individuals/nominal groups from all six blocks, and it is shown that there is a striking interaction between task complexity and performance --- while on tasks of very low complexity, real groups spend substantially more time on the tasks compared to random individuals or nominal groups and thus lead to a lower efficiency, on more complex tasks, real groups are significantly more efficient than individuals or nominal groups. Figure S9-S11 again show the performance comparisons across real groups, individuals, and nominal groups, on normalized scores, duration, and efficiency, respectively, but this time it is broken down by the six blocks. Here again, we observed the interaction effect between task complexity and the performance comparison between groups and individuals persist for groups/individuals across varying levels of skills or social perceptiveness.

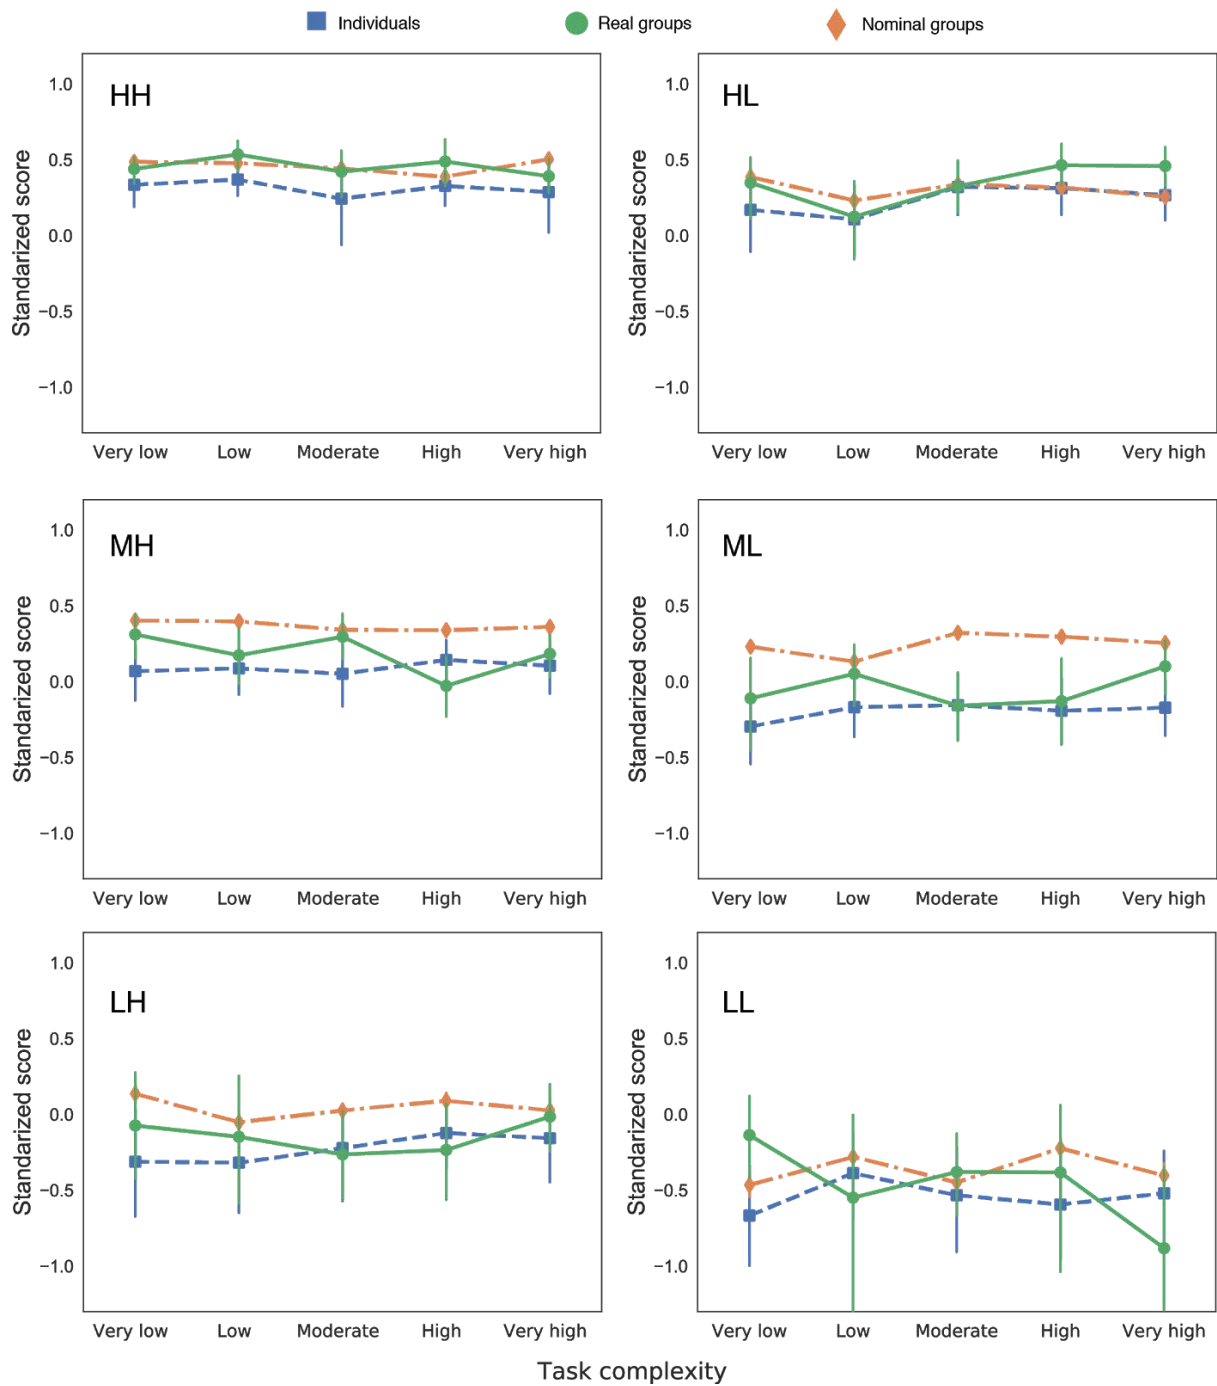


Figure S9. Comparing performance across individuals, real groups, and nominal groups, in terms of the normalized score (i.e., quality of the CSOP solution). Given a block, individual, real group, or nominal group scores are standardized within each task complexity level. Error bars indicate 95% confidence intervals.

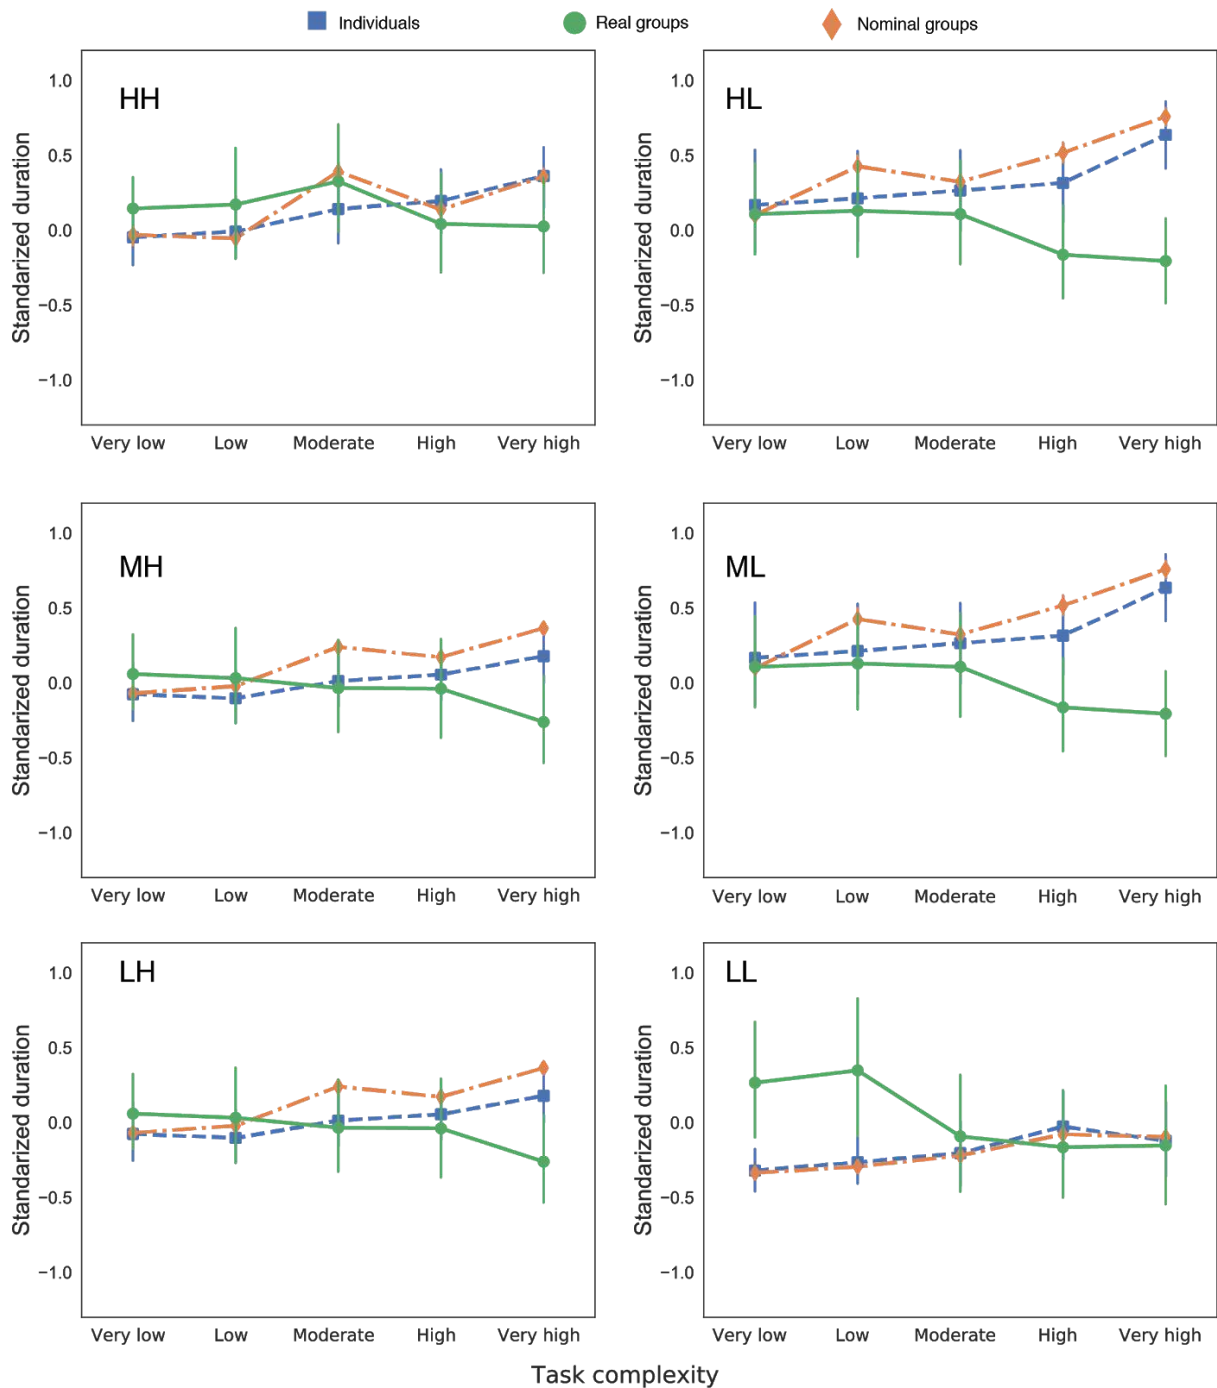


Figure S10. Comparing performance across individuals, real groups, and nominal groups, in terms of the duration (i.e., time to completion). Given a block, individual, real group, or nominal group durations are standardized within each task complexity level. Error bars indicate 95% confidence intervals.

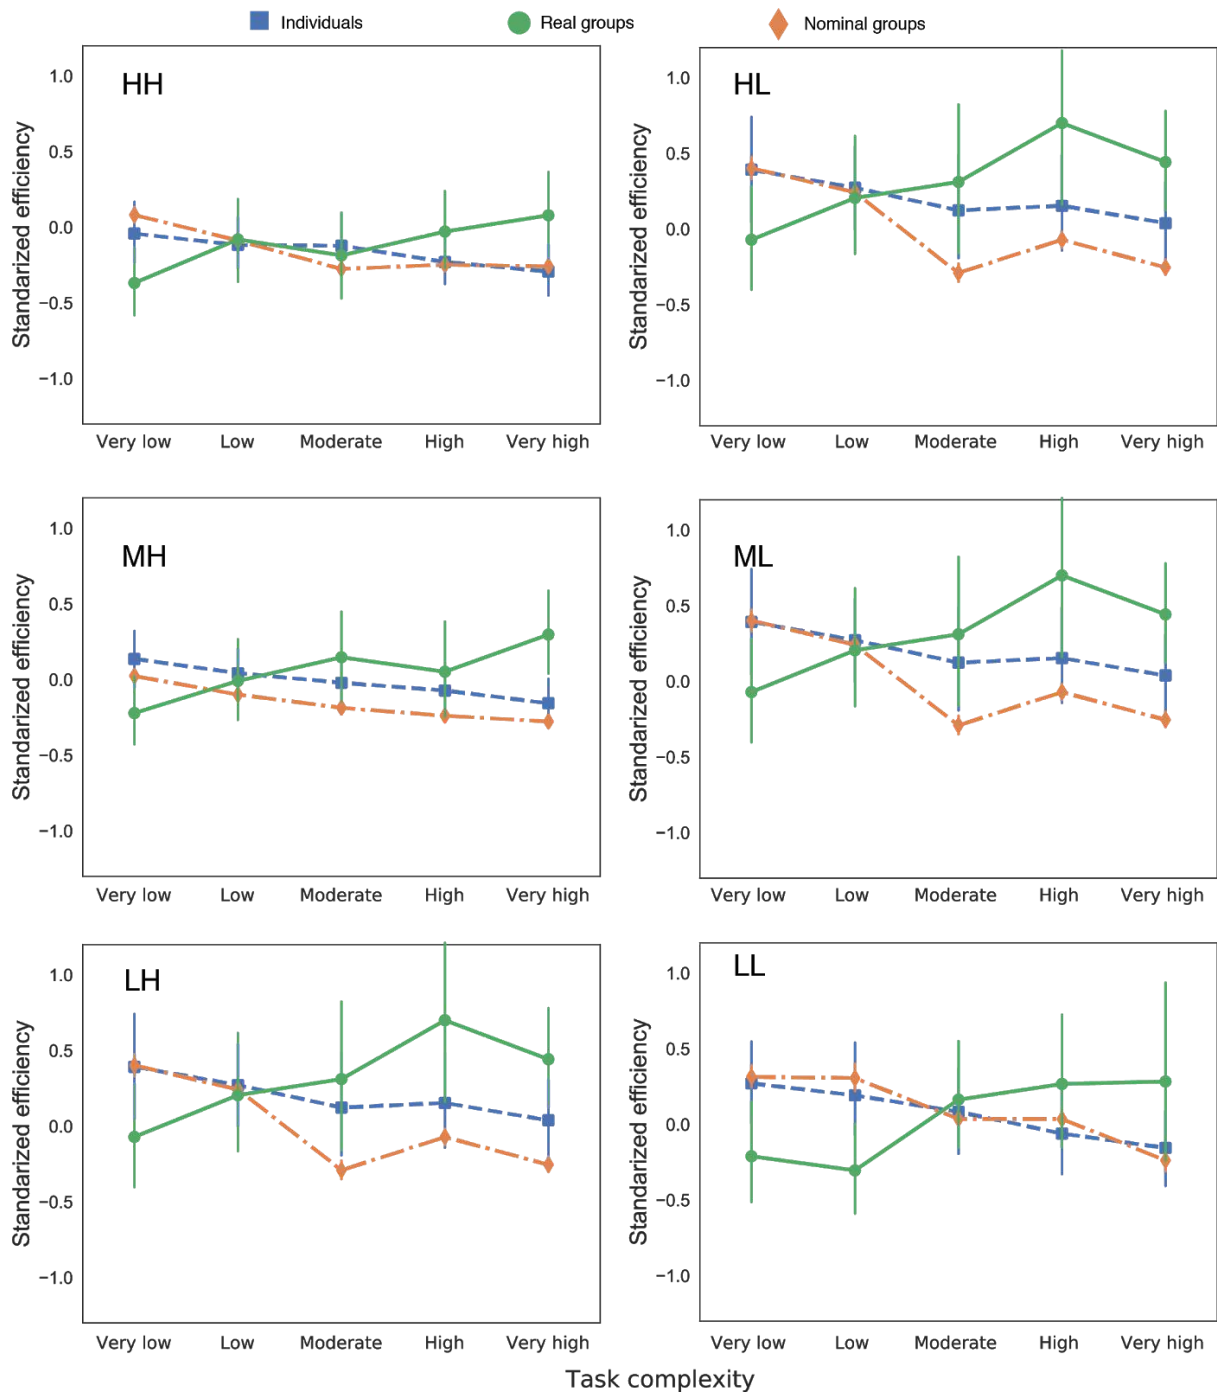


Figure S11. Comparing performance across individuals, real groups, and nominal groups, in terms of efficiency (i.e., normalized score/duration). Given a block, individual, real group, or nominal group efficiencies are standardized within each task complexity level. Error bars indicate 95% confidence intervals.

Finally, to further illustrate the interaction, we conducted linear regressions on the performance obtained by all real groups or real individuals in all tasks of the phase two experiment, where we included the complexity level of the task (reference complexity level is set as “moderate”) and the group indicator variable (reference level is set as “false”; hence the reference is performance of individuals) as the independent variables, and we explicitly considered the interactions between these two variables in our regressions. Table S4-S6 reported the regression results when the performance metric is the normalized score, duration, and efficiency, respectively. Table S4 suggests that on average, real groups consistently score higher than real individuals, though no statistical significance is detected. Table S5-S6, on the other hand, again supported our previous observations that while in general, both groups and individuals spend more time and thus are less efficient on more complex tasks, comparatively speaking, the increase of duration with task complexity for groups is smaller than that of individuals (hence the decrease of efficiency with task complexity for groups is smaller than that of individuals). As a result, groups tend to be less efficient than individuals on tasks of very low complexity but more efficient than individuals on tasks of very high complexity.

Table S4a. Real-groups versus nominal-groups and individuals: the overall effect on score [the reference group is *real-group*]. The mixed effect model includes random effects for group/individual ID to control for correlated error terms within each group.

Model:	MixedLM	Dependent Variable:	Standardized score			
No. Observations:	39960	Method:	REML			
No. Groups:	7992	Scale:	0.4373			
	Coef.	Std.Err.	z	P> z 	[0.025	0.975]
Intercept	0.070	0.044	1.604	0.109	-0.016	0.157
Comparable individuals [True]	-0.130	0.060	-2.174	0.030	-0.246	-0.013
Nominal Group [True]	0.135	0.045*	3.026	0.002	0.047	0.222

Table S4b. Individual versus group score comparison: regression with the group [the reference is false] * complexity [the reference is moderate complexity] interaction term included. The mixed effect model includes random effects for group/individual ID to control for correlated error terms within each group.

Model:	MixedLM	Dependent Variable:	Standardized score			
No. Observations:	2150	Method:	REML			
No. Groups:	430	Scale:	0.6451			
	Coef.	Std.Err.	z	P> z 	[0.025	0.975]
Intercept	-0.056	0.065	-0.857	0.391	-0.184	0.072
group [True]	0.123	0.097	1.270	0.204	-0.067	0.312
Complexity [Very low]	-0.065	0.074	-0.882	0.378	-0.211	0.080
Complexity [Low]	0.010	0.074	0.130	0.896	-0.136	0.155
Complexity [High]	0.024	0.074	0.323	0.747	-0.122	0.170
Complexity [Very high]	0.016	0.074	0.222	0.825	-0.129	0.162
group [True] * Complexity [Very low]	0.144	0.110	1.306	0.192	-0.072	0.359
group [True] * Complexity [Low]	-0.021	0.110	-0.193	0.847	-0.237	0.194
group [True] * Complexity [High]	-0.053	0.110	-0.478	0.632	-0.268	0.163
group [True] * Complexity [Very high]	-0.036	0.110	-0.328	0.743	-0.252	0.179

Table S5. Individual versus group duration comparison: regression with the group [reference is false] * complexity [reference is moderate complexity] interaction term included. The mixed effect model includes random effects for group/individual ID to control for correlated error terms within each group.

	Model:	MixedLM	Dependent Variable:	Standardized duration		
	No. Observations:	2150	Method:	REML		
	No. Groups:	430	Scale:	0.4995		
	Coef.	Std.Err.	z	P> z 	[0.025	0.975]
Intercept	0.008	0.065	0.125	0.901	-0.119	0.135
group [True]	-0.018	0.096	-0.185	0.854	-0.206	0.171
Complexity [Very Low]	-0.099	0.065	-1.512	0.131	-0.227	0.029
Complexity [Low]	-0.085	0.065	-1.300	0.194	-0.213	0.043
Complexity [High]	0.084	0.065	1.283	0.200	-0.044	0.212
Complexity [Very High]	0.189***	0.065	2.889	0.004	0.061	0.317
group [True] * Complexity [Very low]	0.217*	0.097	2.239	0.025	0.027	0.406
group [True] * Complexity [Low]	0.186	0.097	1.926	0.054	-0.003	0.376
group [True] * Complexity [High]	-0.184	0.097	-1.900	0.057	-0.374	0.006
group [True] * Complexity [Very high]	-0.414***	0.097	-4.279	< 0.001	-0.604	-0.224

Table S6. Individual versus group efficiency comparison: regression with the group [reference is false] * complexity [reference is moderate complexity] interaction term included. The mixed effect model includes random effects for group/individual ID to control for correlated error terms within each group.

	Model:	MixedLM	Dependent Variable:	Standardized efficiency			
	No. Observations:	2150	Method:	REML			
	No. Groups:	430	Scale:	0.5758			
		Coef.	Std.Err.	z	P> z 	[0.025	0.975]
Intercept		-0.048	0.065	-0.739	0.460	-0.175	0.079
group [True]		0.105	0.096	1.095	0.273	-0.083	0.293
Complexity [Very low]		0.198**	0.070	2.816	0.005	0.060	0.335
Complexity [Low]		0.082	0.070	1.169	0.242	-0.055	0.220
Complexity [High]		-0.077	0.070	-1.101	0.271	-0.215	0.060
Complexity [Very high]		-0.186**	0.070	-2.652	0.008	-0.324	-0.049
group [True] * Complexity [Very low]		-0.433***	0.104	-4.172	< 0.001	-0.637	-0.230
group [True] * Complexity [Low]		-0.180	0.104	-1.732	0.083	-0.384	0.024
group [True] * Complexity [High]		0.169	0.104	1.631	0.103	-0.034	0.373
group [True] * Complexity [Very high]		0.408***	0.104	3.929	< 0.001	0.205	0.612

3.4 Comparing effects on collective performance

To estimate the effect size of each influencing factor (i.e., skill, social perceptiveness, skill diversity, cognitive style diversity) on collective performance for tasks of varying complexity levels, we conducted linear regressions on the collective performance data obtained on each of the 5 task instances that we included in phase two. Specifically, given a task instance (e.g., the task of “high” complexity in phase two), a particular influencing factor (e.g., social perceptiveness), and a performance metric (e.g., normalized score), we regressed each group’s standardized performance in that task on the group’s standardized value of the chosen influencing factor to estimate how this factor affects collective performance.

Standardization of performance metric values and influencing factor values are conducted within the chosen task. For example, given a task instance t at a particular complexity level, we first pooled the normalized scores obtained by all groups on t . Then, a group’s standardized score on t can be computed as (the group’s normalized score on task t - the mean value of all groups’ normalized scores on task t) / standard deviation of all groups’ normalized scores on task t . Such standardization enables meaningful comparisons of effect sizes across tasks of different complexity levels.

Figure 5 in the main text reports the coefficients estimated from such linear regressions as well as the 95% confidence interval associated with each estimate. As an example, we can get conclusions like the following by inspecting Figure 5: In a task of high complexity, increasing a group’s skill level by one standard deviation leads to about 0.33 standard deviation increases in the normalized score that the group can obtain, while increasing a group’s social perceptiveness by one standard deviation leads to about 0.15 standard deviation increase in score.

Note that in Figure 5, we reported the effect sizes of skill diversity on collective performance using the above procedure as per our pre-registration, and we found that higher skill diversity was associated with lower scores. A closer look at the data suggested that a group’s skill diversity is negatively correlated with the group’s skill level, indicating the observed negative effect of skill diversity on score might actually reflect the positive relationship between a group’s skill level and normalized score that the group can obtain. Although this negative association between skill diversity and the level is consistent with the original claim

regarding skill diversity⁵, we might nonetheless wonder if the effect of diversity might change sign after conditioning on skill level. To check this possibility, for each task complexity level, we further fitted the data into a mixed effect model, where the dependent variables were standardized collective performance metrics, the independent variable was the standardized skill diversity of the group, and we also took into account the random effects associated with the category of the group's skill level (i.e., high, low, or mixed, which is decided by whether the group is composed of three members of high skills, three members of low skills, or three members with mixed levels of skills). Figure S12 shows the effect of skill diversity estimated from this mixed effect model. Again, we found that even after adjusting for the group's skill level, the effects of skill diversity on the group's score is still largely negative, especially on the tasks of lower complexity.

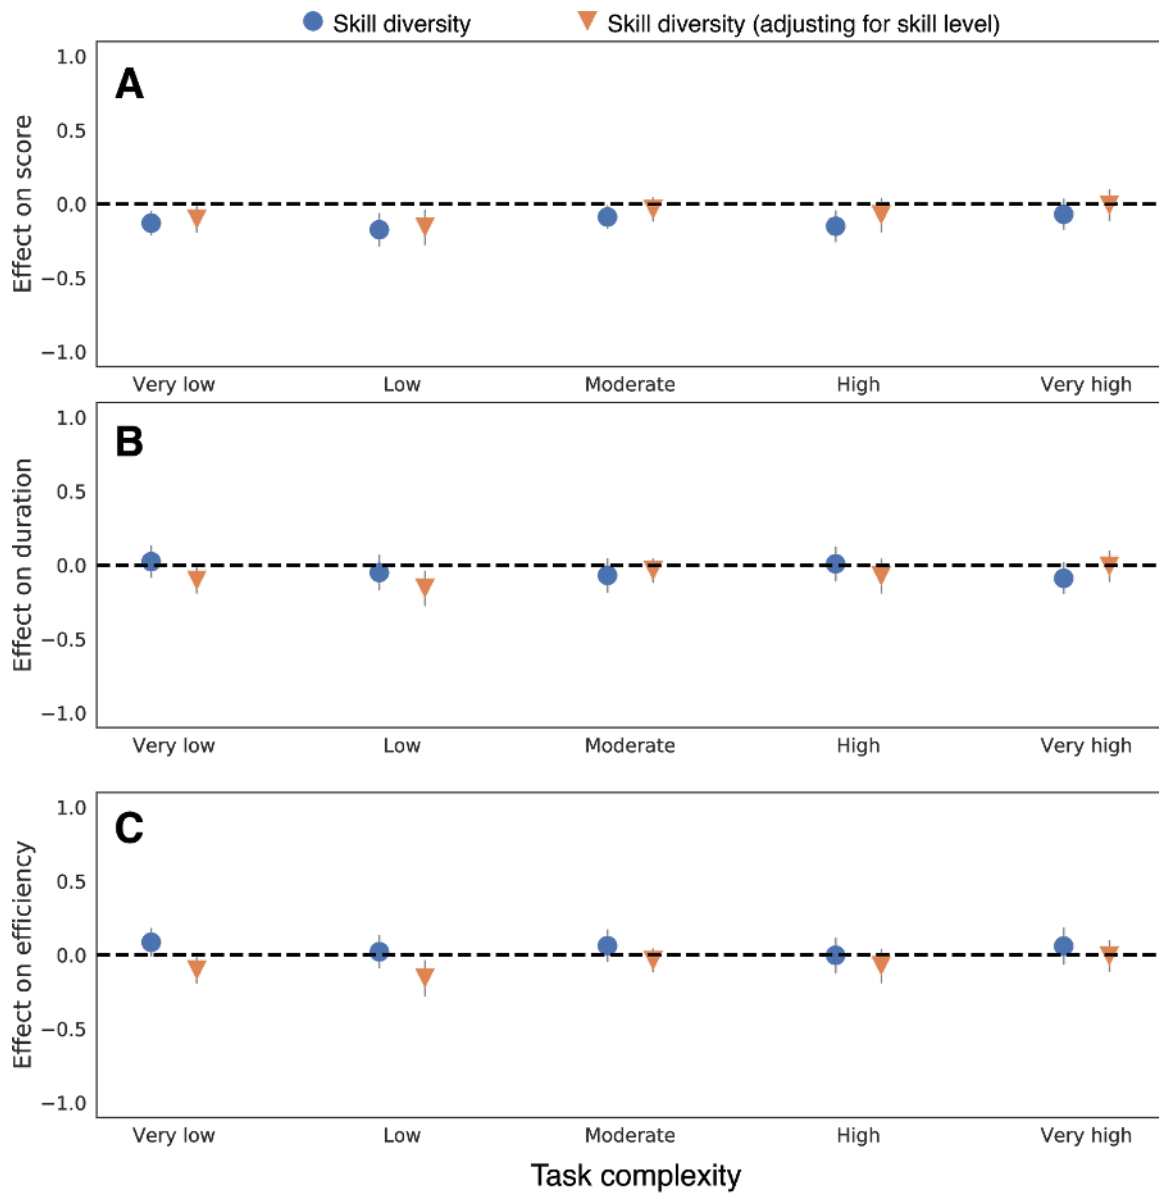


Figure S12. Standardized regression coefficients for skill diversity (pre-registered), and skill diversity while adjusting for the group’s skill level category (i.e., high, low, or mixed; not pre-registered) as a function of task complexity when predicting (A) normalized score, (B) duration, and (C) efficiency. Error bars indicate 95% confidence intervals. All regression coefficients are standardized within each task complexity level as per our pre-registration.

To determine the effect of cognitive style diversity on collective performance, we first conducted a selection of cognitive styles among the 4 candidate operationalizations using the following criteria which are consistent with earlier work ⁶: an ideal operationalization of cognitive style must be persistent for a given individual (i.e., consistent across tasks); it must be heterogeneous across the sample (ideally, roughly equal numbers of participants would have each style); it must not be highly correlated with the individual's skill level. Among the 4 operationalizations, we identified a participant's emphasis on optimization versus constraint satisfaction (i.e., optimizer vs. resolver) as the most appropriate. Notably, all other operationalizations of cognitive styles are correlated with the skill level to some degree (i.e., participants of one style tend to consistently have a higher skill on the room assignment tasks). Therefore, when we refer to cognitive style diversity in the main text it is defined in terms of the mixture of optimizers vs resolvers. As we show in Figure 5, according to that definition, cognitive style diversity does not have any consistent and significant effects on collective performance. To check the robustness of this result with respect to alternative definitions, for each task complexity level we fitted a mixed effect model, where the dependent variable is standardized collective performance (i.e., normalized score, duration, or efficiency), the fixed effects are the standardized cognitive style diversity of the group using different operationalizations of cognitive styles, and we included random effects for the category of the group's skill level (i.e., high, mixed, or low). Figure S13 shows the results, which are consistent with what we observed in Figure 5.

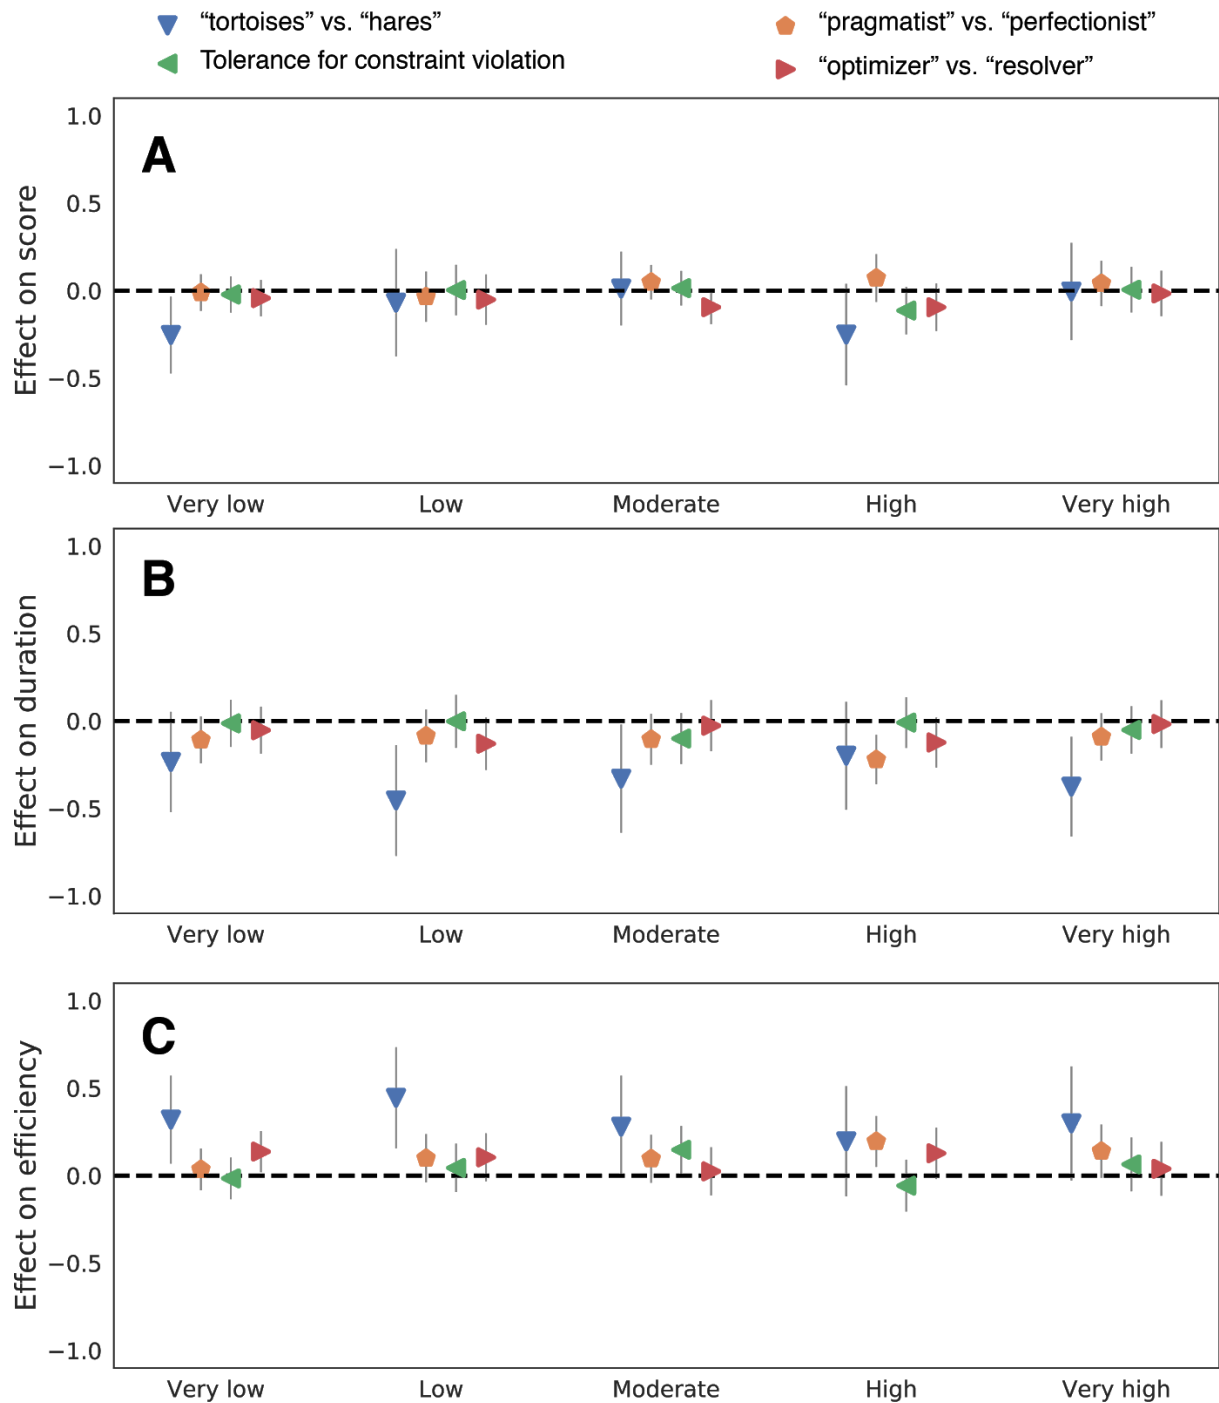


Figure S13. Standardized regression coefficients for different operationalizations of cognitive style diversity (adjusting for the group’s skill level category) as a function of task complexity when predicting (A) normalized score, (B) duration, and (C) efficiency. Error bars indicate 95% confidence intervals. All regression coefficients are standardized within each task complexity level as per our pre-registration.

3.5 Testing robustness of the group composition results

Inspecting Figure 5, we found that the complexity level of a task does not seem to moderate the effects of group composition on its performance. Hence, as an alternative robustness check, we further pooled the collective performance data across all tasks of various complexity levels and estimated the effects of each influencing factor through mixed-effect models. Specifically, Tables S7-S15 report the results of mixed-effect models, where the dependent variables are performance metrics (again standardized within each complexity level), the independent variables are the influencing factors of interests (i.e., S7: skill level; S8: social perceptiveness; S9: skill diversity; S10: skill diversity with skill level adjustment; S11: cognitive style diversity where cognitive style is defined as optimizer vs. resolver; S12: cognitive style diversity where cognitive style is defined as optimizer vs. resolver with skill level adjustment), and we also took the random effects associated with each group into consideration. Tables S13-S15 then report the results of mixed-effect model for cognitive style diversity where the cognitive style is defined using different ways (S13: tortoises vs. hares; S14: constraint violation tolerance low vs. high; S15: pragmatist vs. perfectionist), and again we considered random effects associated with both each group and the group's category of skill level.

Table S7. The relation between a group’s skill level and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s skill (i.e., the average skill level of three members of the group). All models include random effects for groups as an intercept to account for dependence across tasks (i.e., random effects are clustered on each group, using group id as the identifier). Increasing a group’s skill significantly increases the group’s score in solving CSOPs, but has no effect on duration or efficiency.

	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimate	CI (95%)	P	Estimate	CI (95%)	P	Estimate	CI (95%)	P
(Intercept)	0.08*	0.01 – 0.15	0.022	-0.03	-0.14 – 0.08	0.625	0.05	-0.05 – 0.15	0.330
Skill	0.30***	0.22 – 0.38	<0.001	0.10	-0.03 – 0.23	0.126	-0.10	-0.22 – 0.02	0.097
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.13 _{group_id}			0.52 _{group_id}			0.39 _{group_id}		
ICC	0.18 _{group_id}			0.50 _{group_id}			0.40 _{group_id}		
Observations	980			980			980		
R ²	0.252			0.508			0.400		

Table S8. The relation between a group’s social perceptiveness level and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s social perceptiveness level (i.e., the average social perceptiveness level of three members of the group). All models include random effects for groups as an intercept to account for dependence across tasks (i.e., random effects are clustered on each group, using group id as the identifier). Increasing a group’s social perceptiveness level significantly increases the group’s score in solving CSOPs, but has no effect on duration or efficiency.

	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimate	CI (95%)	P	Estimate	CI (95%)	P	Estimate	CI (95%)	P
(Intercept)	0.07	-0.01 – 0.14	0.071	-0.03	-0.14 – 0.08	0.592	0.05	-0.05 – 0.15	0.305
Social Perc.	0.17***	0.09 – 0.25	<0.001	-0.05	-0.17 – 0.07	0.392	0.04	-0.07 – 0.14	0.509
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.17 _{group_id}			0.52 _{group_id}			0.40 _{group_id}		
ICC	0.23 _{group_id}			0.51 _{group_id}			0.40 _{group_id}		
Observations	980			980			980		
R ²	0.252			0.508			0.400		

Table S9. The relation between a group’s skill diversity (i.e., the variance of the group member’s skill levels) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s skill diversity. All models include random effects for groups as an intercept to account for dependence across tasks (i.e., random effects are clustered on each group, using group id as the identifier). Increasing a group’s skill diversity significantly decreases the group’s score in solving CSOPs, but has no effect on duration or efficiency.

	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimate	CI (95%)	P	Estimate	CI (95%)	P	Estimate	CI (95%)	P
(Intercept)	0.14***	0.06 – 0.23	0.001	-0.01	-0.13 – 0.11	0.869	0.03	-0.08 – 0.14	0.642
Skill diversity	-0.12***	-0.18 – -0.06	<0.001	-0.03	-0.12 – 0.05	0.440	0.05	-0.04 – 0.13	0.269
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.17 _{group_id}			0.52 _{group_id}			0.40 _{group_id}		
ICC	0.23 _{group_id}			0.51 _{group_id}			0.40 _{group_id}		
Observations	980			980			980		
R ²	0.252			0.508			0.400		

Table S10. The relation between a group’s skill diversity and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s skill diversity (i.e., the variance of group member’s skill levels). All models include random effects for groups as well as the group’s skill level category as an intercept to account for dependence across tasks. Increasing a group’s skill diversity significantly decreases the group’s score in solving CSOPs, but has no effect on duration or efficiency.

Coefficient	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimates	CI (95%)	P	Estimates	CI (95%)	P	Estimates	CI (95%)	P
(Intercept)	0.08	-0.31 – 0.48	0.685	-0.01	-0.13 – 0.11	0.869	0.03	-0.08 – 0.14	0.642
Skill diversity	-0.07**	-0.13 – -0.01	0.018	-0.03	-0.12 – 0.05	0.440	0.05	-0.04 – 0.13	0.269
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.11 _{group_id}			0.52 _{group_id}			0.40 _{group_id}		
	0.12 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
ICC	0.13 _{group_id}			0.51 _{group_id}			0.40 _{group_id}		
	0.15 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
Observations	980			980			980		
R ²	0.289			0.508			NA		

Table S11. The relation between a group’s cognitive style diversity (in terms of whether all group members are optimizers/resolvers or both types exist in the group) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s cognitive style diversity. All models include random effects for groups as an intercept to account for dependence across tasks (i.e., random effects are clustered on each group, using group id as the identifier). Increasing a group’s cognitive diversity has no effect on the group’s score, duration or efficiency.

Coefficient	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimates	CI (95%)	P	Estimates	CI (95%)	P	Estimates	CI (95%)	P
(Intercept)	0.07	-0.01 – 0.15	0.074	-0.03	-0.14 – 0.08	0.583	0.05	-0.05 – 0.15	0.298
optimizers vs resolvers	-0.06	-0.14 – 0.02	0.128	-0.07	-0.18 – 0.04	0.214	0.09	-0.01 – 0.19	0.091
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.19 _{group_id}			0.52 _{group_id}			0.39 _{group_id}		
ICC	0.25 _{group_id}			0.51 _{group_id}			0.40 _{group_id}		
Observations	980			980			980		
R ²	0.253			0.508			0.400		

Table S12. The relation between the group’s cognitive style diversity (in terms of whether all group members are optimizers/resolvers or both type exist in the group) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s cognitive style diversity. All models include random effects for groups as well as the group’s skill level category as an intercept to account for dependence across tasks. Increasing a group’s cognitive style diversity has no effect on the group’s score, duration or efficiency.

	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimate	CI (95%)	P	Estimate	CI (95%)	P	Estimate	CI (95%)	P
(Intercept)	0.04	-0.38 – 0.47	0.845	-0.03	-0.14 – 0.08	0.583	0.05	-0.05 – 0.15	0.298
optimizers vs resolvers	-0.03	-0.10 – 0.04	0.397	-0.07	-0.18 – 0.04	0.214	0.09	-0.01 – 0.19	0.091
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.11 _{group_id}			0.52 _{group_id}			0.39 _{group_id}		
	0.14 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
ICC	0.14 _{group_id}			0.51 _{group_id}			0.40 _{group_id}		
	0.17 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
Observations	980			980			980		
R ²	0.303			0.508			NA		

Table S13. The relation between the group’s cognitive style diversity (in terms of whether all group members are fast/slow problem solvers or both types exist in the group) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s cognitive style diversity. All models include random effects for groups as well as the group’s skill level category as an intercept to account for dependence across tasks. Increasing a group’s cognitive style diversity has no effect on the group’s score, but reduces the duration and increases efficiency.

Coefficient	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimates	CI (95%)	P	Estimates	CI (95%)	P	Estimates	CI (95%)	P
(Intercept)	0.04	-0.38 – 0.47	0.843	-0.03	-0.14 – 0.08	0.578	0.05	-0.05 – 0.16	0.326
CogStyle. diversity (fast vs. slow)	-0.02	-0.09 – 0.04	0.466	-0.15**	-0.26 – -0.04	0.007	0.14**	0.04 – 0.24	0.005
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.11	group_id		0.50	group_id		0.38	group_id	
	0.14	skill_type		0.00	skill_type		0.00	skill_type	
ICC	0.14	group_id		0.50	group_id		0.39	group_id	
	0.17	skill_type		0.00	skill_type		0.00	skill_type	
Observations	980			980			980		
R ²	0.303			NA			0.400		

Table S14. The relation between the group’s cognitive style diversity (in terms of whether all group members have the same constraint violation tolerance or not) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s cognitive style diversity. All models included random effects for groups as well as the group’s skill level category as an intercept to account for dependence across tasks. Increasing a group’s cognitive style diversity has no effect on the group’s score, but reduces the duration and increases efficiency.

Coefficient	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimates	CI (95%)	P	Estimates	CI (95%)	P	Estimates	CI (95%)	P
(Intercept)	0.04	-0.39 – 0.47	0.846	-0.03	-0.15 – 0.09	0.604	0.05	-0.08 – 0.18	0.421
CogStyle diversity (constraint violation tolerance)	0.01	-0.05 – 0.08	0.688	-0.12*	-0.23 – -0.01	0.028	0.12*	0.02 – 0.22	0.023
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.11 _{group_id}			0.51 _{group_id}			0.38 _{group_id}		
	0.14 _{skill_type}			0.00 _{skill_type}			0.01 _{skill_type}		
ICC	0.14 _{group_id}			0.50 _{group_id}			0.39 _{group_id}		
	0.17 _{skill_type}			0.00 _{skill_type}			0.01 _{skill_type}		
Observations	980			980			980		
R ²	0.305			0.508			0.401		

Table S15. The relation between the group’s cognitive style diversity (in terms of whether all group members are pragmatic/tenacious or both types exist in the group) and collective performance. Data is combined across groups in all six blocks and for all five tasks. Models relate performance measures (standardized within each task) with the group’s cognitive style diversity. All models include random effects for groups as well as the group’s skill level category as an intercept to account for dependence across tasks. Increasing a group’s cognitive style diversity has no effect on the group’s score, duration, or efficiency.

Coefficient	Standardized Score			Standardized Duration			Standardized Efficiency		
	Estimates	CI (95%)	P	Estimates	CI (95%)	P	Estimates	CI (95%)	P
(Intercept)	0.04	-0.39 – 0.47	0.846	-0.03	-0.14 – 0.08	0.585	0.05	-0.07 – 0.17	0.379
CogStyle diversity (conservative/pragmatic)	-0.02	-0.09 – 0.04	0.467	-0.04	-0.15 – 0.08	0.533	0.04	-0.06 – 0.14	0.473
Random Effects									
σ^2	0.57			0.51			0.60		
τ_{00}	0.11 _{game_id}			0.52 _{game_id}			0.40 _{game_id}		
	0.14 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
ICC	0.13 _{game_id}			0.51 _{game_id}			0.40 _{game_id}		
	0.17 _{skill_type}			0.00 _{skill_type}			0.00 _{skill_type}		
Observations	980			980			980		
R ²	0.306			0.508			0.400		

3.6 Predicting collective performance

To see how much various influencing factors can be used to “predict” the collective performance, we took an out-of-sample prediction practice. Specifically, we started by considering the simple linear regression model that uses influencing factors to predict a group’s cumulative normalized score (i.e., the sum of a group’s normalized scores in all five task instances of phase two), where influencing factors are included independently (Figure 6A in the main text; blue diamond symbols), in an increasing order of explanatory power (Figure 6A orange square symbols), or in a decreasing order of explanatory power (Figure 6B in the main text). The out-of-sample R^2 we reported in Figure 6 was obtained through the following procedure: We randomly split the entire set of groups that participated our phase two experiment into a 70% training set on which the predictive model was learned, and we then tested the model on the rest 30% testing set. This procedure was repeated 5 times, and we illustrated the mean R^2 and 95% confidence intervals for each model in Figure 6. In addition to simple linear regression models, we also considered more sophisticated models including a regularized regression model, elasticNet, and random forest, and we obtained similar results shown in Figure S14. Finally, Figures S15-16 show the results of using various influencing factors to predict a group’s cumulative duration and efficiency on tasks (i.e., the sum of the group’s duration/efficiency of each of the five task instances of phase two) through different predictive models, and it is shown that compared to normalized scores that a group could obtain, how much time the group would spend on tasks and how efficient it is are substantially more difficult to predict using various variables of group composition.

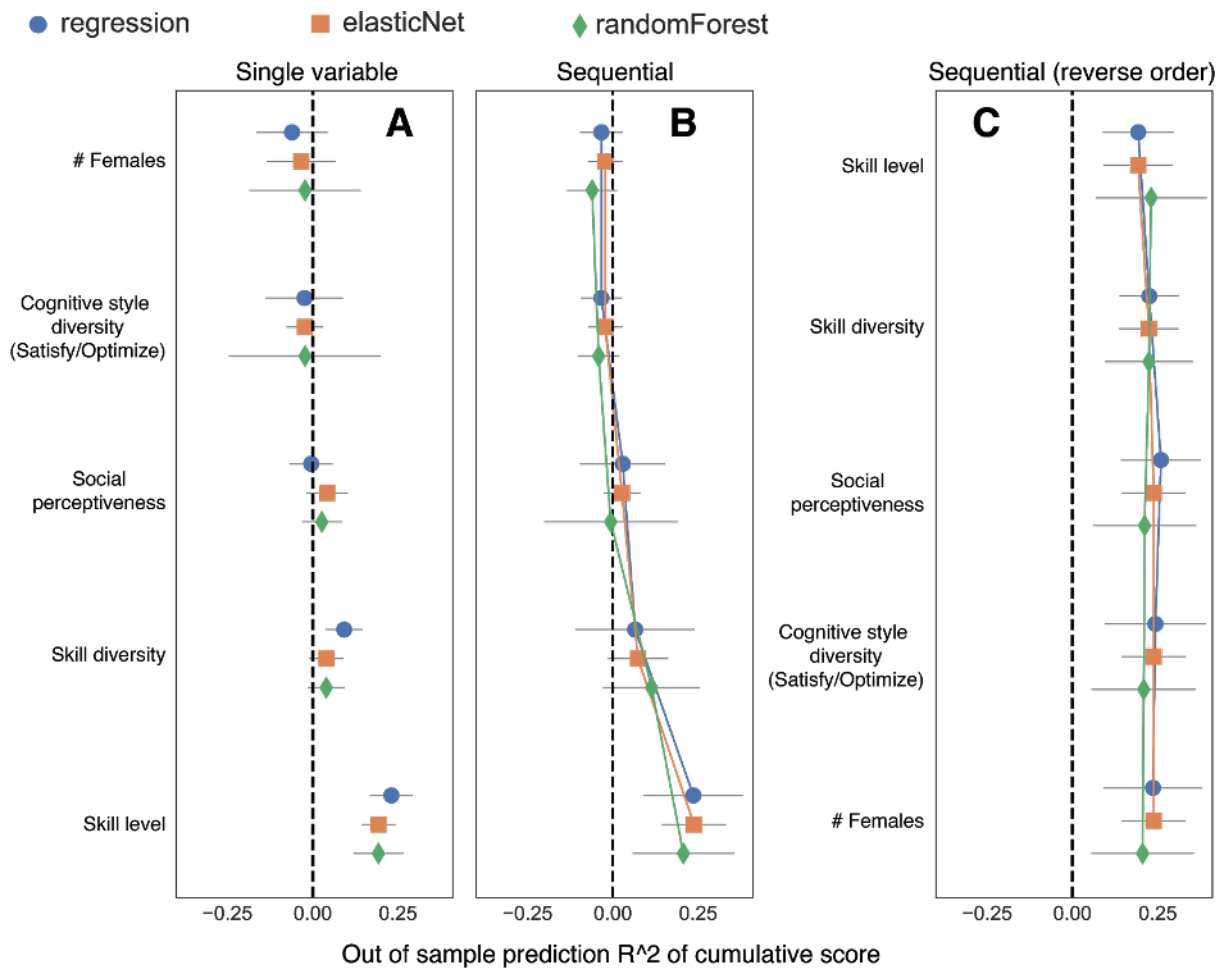


Figure S14. Out of sample predictions on the group's cumulative score. Predict the group's normalized score with the group's skill level, skill diversity, social perceptiveness, cognitive style diversity, and the number of female group members. Three models (i.e., linear regression, elasticNet, and random forest) are used. Models are first learned on 70% of the groups and then tested on the rest 30% of the groups. This procedure is then repeated 5 times. Error bars indicate 95% confidence intervals. In all models, the majority of the explained variance in group's normalized score can be attributed to the group's skill level.

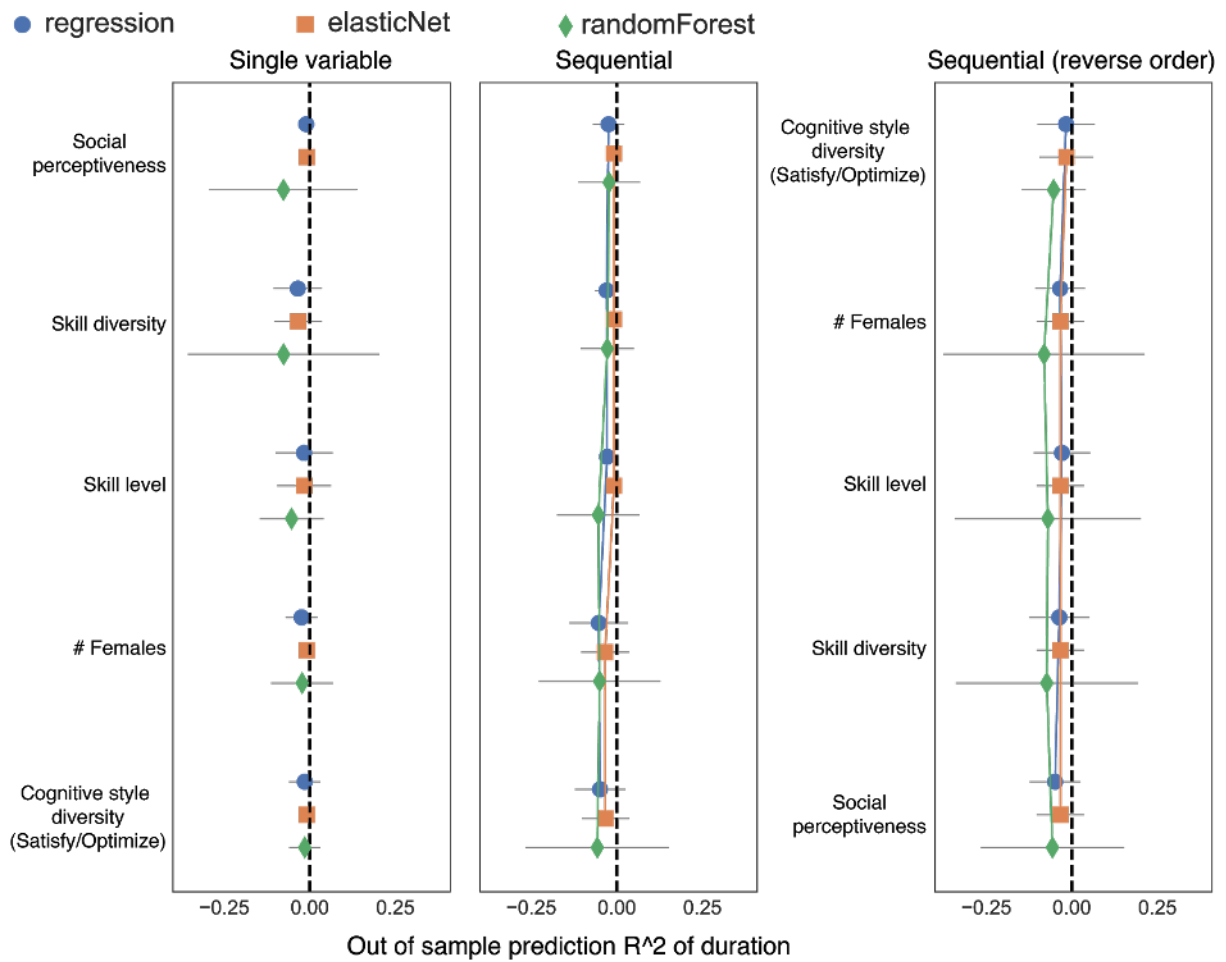


Figure S15. Out of sample predictions on the group’s duration on tasks. Predict the group’s duration on tasks with the group’s skill level, skill diversity, social perceptiveness, cognitive style diversity, and the number of female group members. Three models (i.e., linear regression, elasticNet, and random forest) are used. Models are first learned on 70% of the groups and then tested on the rest 30% of the groups. This procedure is then repeated 5 times. Error bars indicate 95% confidence intervals. The set of independent variables can hardly be used to explain the variance in the group’s duration on tasks.

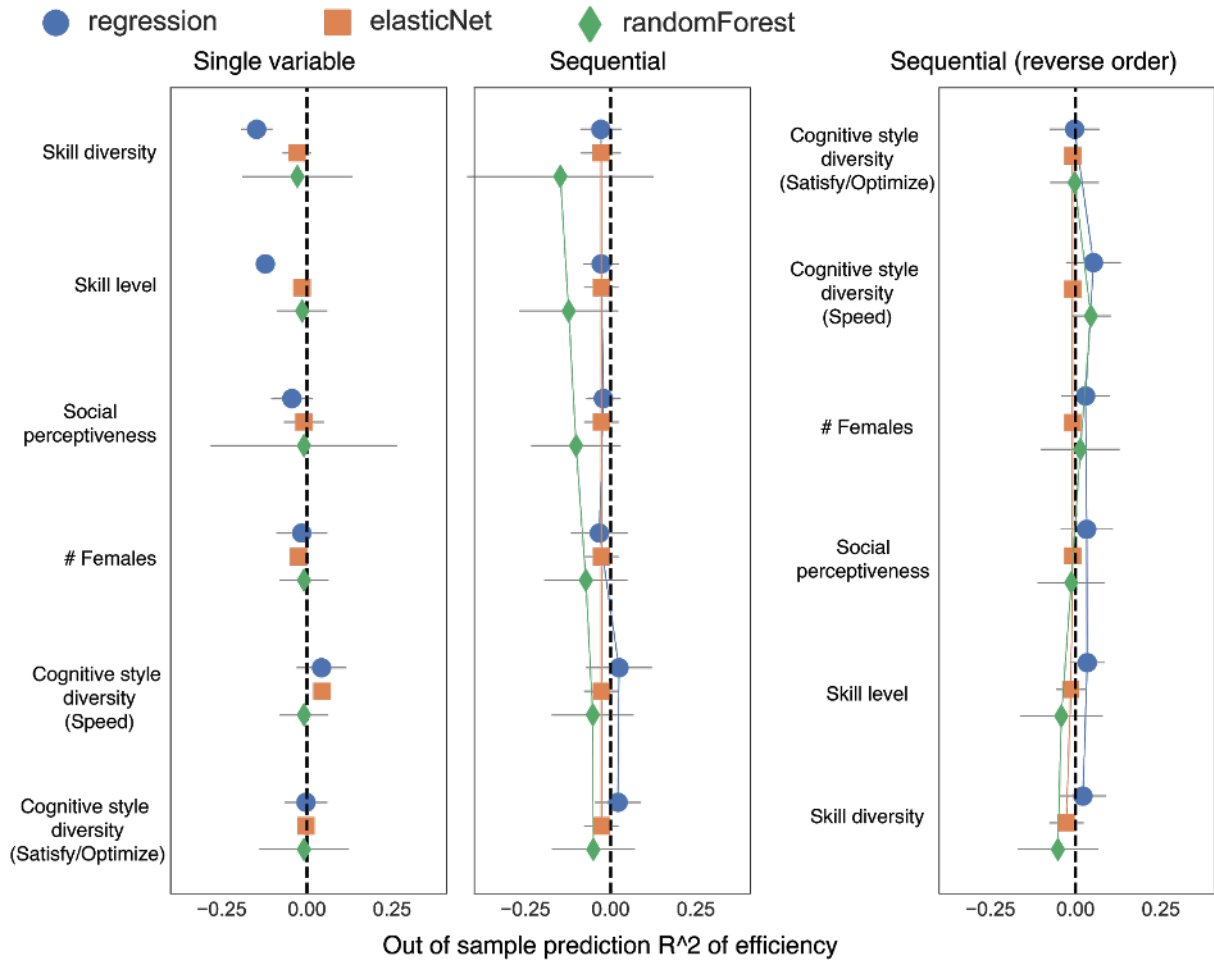


Figure S16. Out of sample predictions on the group's efficiency on tasks. Predict the group's duration on tasks with the group's skill level, skill diversity, social perceptiveness, cognitive style diversity, and the number of female group members. Three models (i.e., linear regression, elasticNet, and random forest) are used. Models are first learned on 70% of the groups and then tested on the rest 30% of the groups. This procedure is then repeated 5 times. Error bars indicate 95% confidence intervals. The set of independent variables can hardly be used to explain the variance in the group's duration on tasks.

References

1. Paton, N. & Almaatouq, A. Empirica: Open-Source, Real-Time, Synchronous, Virtual Lab Framework. (2018). doi:10.5281/zenodo.1488413
2. Ghédira, K. & Dubuisson, B. Constraint Satisfaction and Optimization Problems. in *Constraint Satisfaction Problems* (eds. Ghédira, K. & Dubuisson, B.) 165–180 (John Wiley & Sons, Inc., 2013).
3. Tsang, E. *Foundations of Constraint Satisfaction: The Classic Text*. (BoD – Books on Demand, 2014).
4. Baron-Cohen, S., Wheelwright, S., Hill, J., Raste, Y. & Plumb, I. The ‘Reading the Mind in the Eyes’ test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *J. Child Psychol. Psychiatry* **42**, 241–251 (2001).
5. Hong, L. & Page, S. E. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc. Natl. Acad. Sci. U. S. A.* **101**, 16385–16389 (2004).
6. Aggarwal, I. & Woolley, A. W. Team Creativity, Cognition, and Cognitive Style Diversity. *Manage. Sci.* (2018). doi:10.1287/mnsc.2017.3001