



How intermittent breaks in interaction improve collective intelligence

Ethan Bernstein^{a,1}, Jesse Shore^{b,1,2}, and David Lazer^{c,d,1}

^aOrganizational Behavior Unit, Harvard Business School, Harvard University, Boston, MA 02163; ^bInformation Systems Department, Questrom School of Business, Boston University, Boston, MA 02215; ^cDepartment of Political Science, College of Computer and Information Science, Northeastern University, Boston, MA 02115; and ^dThe Institute for Quantitative Social Science, Harvard University, Cambridge, MA 02138

Edited by Matthew O. Jackson, Stanford University, Stanford, CA, and approved June 21, 2018 (received for review February 8, 2018)

People influence each other when they interact to solve problems. Such social influence introduces both benefits (higher average solution quality due to exploitation of existing answers through social learning) and costs (lower maximum solution quality due to a reduction in individual exploration for novel answers) relative to independent problem solving. In contrast to prior work, which has focused on how the presence and network structure of social influence affect performance, here we investigate the effects of time. We show that when social influence is intermittent it provides the benefits of constant social influence without the costs. Human subjects solved the canonical traveling salesperson problem in groups of three, randomized into treatments with constant social influence, intermittent social influence, or no social influence. Groups in the intermittent social-influence treatment found the optimum solution frequently (like groups without influence) but had a high mean performance (like groups with constant influence); they learned from each other, while maintaining a high level of exploration. Solutions improved most on rounds with social influence after a period of separation. We also show that storing subjects' best solutions so that they could be reloaded and possibly modified in subsequent rounds—a ubiquitous feature of personal productivity software—is similar to constant social influence: It increases mean performance but decreases exploration.

collective intelligence | social influence | social networks

Collective intelligence—the ability of collectives of individuals to solve problems well—has emerged as an important interdisciplinary area of study with applications in understanding and supporting the performance of groups and teams (1), networks (2), crowds (3–6), financial markets (7), prediction markets (8), innovation contests (9), and democracies (10, 11), as well as collectives of nonhuman organisms (e.g., ref. 12). Across these diverse and important settings, a fundamental question is this: How does social influence—exposure of solvers to each other's behavior or solutions through interacting—affect collective intelligence?

In this work, we conduct randomized experiments to study how collective intelligence is affected by two frequently experienced impacts of technology use: changes to the temporal nature of social influence (from intermittent social influence, which is more characteristic of face-to-face communication, to constant social influence, characteristic of “always-on,” transparency-enhancing communication technologies) and storage and quick recall of a solver's current best solution to a problem (which in effect increases the influence of a solver's past solutions on their current solution).

Past research shows that social influence leads individuals to adopt their peers' opinions and copy their solutions to problems (5), especially under conditions of network clustering (13, 14), leading to a loss of aggregate diversity. Under certain conditions—performing simple estimation tasks or searching simple solution spaces with clear performance feedback—more efficiently connected networks of problem solvers can collec-

tively outperform disconnected individuals (12, 15, 16) and inefficient networks (17, 18), as people learn from each other and better solutions spread rapidly. In general, however, and especially for complex and uncertain problems, maintaining and integrating diverse information and perspectives is a critical driver of collective performance (1, 9, 19). Social influence and in particular network clustering can result in too much local copying behavior, driving out beneficial diversity and resulting in a collective convergence on a suboptimal solution (2, 5, 7, 17, 18). For example, sharing ideas in the early stages of a brainstorming task has been shown to reduce the number and quality of ideas produced (20).

Working together in clusters or groups does offer other benefits for handling large problems, despite the tendency for connected individuals to underexplore solution spaces. In particular, groups are capable of handling complex problems that individuals themselves cannot (21, 22). Innovation and invention are also widely thought to be social processes, in which ideas or partial ideas from multiple individuals are recombined (23, 24).

A major unresolved question in collective intelligence in complex tasks is thus whether it is possible to get the benefits of social influence and network clustering (collective learning) without the associated costs (premature convergence on a suboptimal solution). Here, we report on an experimental study that provides evidence that it is indeed possible, and moreover that

Significance

Many human endeavors—from teams and organizations to crowds and democracies—rely on solving problems collectively. Prior research has shown that when people interact and influence each other while solving complex problems, the average problem-solving performance of the group increases, but the best solution of the group actually decreases in quality. We find that when such influence is intermittent it improves the average while maintaining a high maximum performance. We also show that storing solutions for quick recall is similar to constant social influence. Instead of supporting more transparency, the results imply that technologies and organizations should be redesigned to intermittently isolate people from each other's work for best collective performance in solving complex problems.

Author contributions: E.B., J.S., and D.L. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: The data reported in this paper have been deposited in the Harvard Dataverse (<https://doi.org/10.7910/DVN/TSSQLY>).

¹E.B., J.S., and D.L. contributed equally to this work.

²To whom correspondence should be addressed. Email: jccs@bu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1802407115/-DCSupplemental.

Published online August 13, 2018.

conditions typical of real (as opposed to laboratory) face-to-face social networks result in both benefits.

We study the performance of sets of three individuals (hereafter “triads”) completing the Euclidean traveling salesperson problem (TSP), which involves finding the shortest path among symbols representing cities on a synthetic 2D map presented visually. The TSP is NP-hard (nondeterministic polynomial time hard) (25) and characterized by many local optima (26); thus, like other tasks thought to be good models for complex problems (27), solution spaces for the TSP are “rugged” in that simple hill climbing will generally fail to produce a good solution. Although feasible for human subjects (28), finding the globally optimal solution is not trivial and is expected to benefit from more—or more efficient—collective exploration. In our study, each TSP map included 25 different cities; a full path included 25 “legs” of the journey, each connecting a pair of cities. In a single trial, our subjects completed the task 17 times (“rounds”) in a row and thus were able to refine their solution and, depending on the experimental treatment they were assigned to, learn from the other members of their triad.

Our experimental treatments were inspired by the fact that outside of the laboratory real face-to-face communication ties are not constant: Even strong social ties involve intermittent interaction punctuated by time apart (29). Thus, we conducted a three-way randomization with respect to how much network ties within the triad are “on.” One-third of our subjects were assigned to a constant ties (CT) condition, in which they could see the solutions of their neighbors every round of the trial. One-third of our subjects were assigned to an intermittent ties (IT) condition, in which they were able to see their neighbors’ solutions every three rounds (on rounds 4, 7, 10, 13, and 16). The final third were assigned to a no ties (NT) condition, in which subjects could never see their neighbors’ solutions.

In wisdom of the crowd-type tasks (with applications in estimation and prediction), scholars focus on the mean (or other measure of central tendency) of a collective of estimates (3, 5, 8). In complex problem-solving settings such as ours, in addition to the quality of the mean solution, the quality of the best solution produced in a collective is often of critical importance (with applications in, e.g., brainstorming, crowdsourcing, and innovation) (9, 30). In this latter context, scholars have been particularly focused on whether a collective finds the global optimum to a complex problem (2, 17). We therefore consider both performance metrics—best solution and mean solution—in our study.

Results

Main Result. Because NT triads lack social influence among solvers, prior literature predicts that NT triads would generate more diverse solutions and thus find the optimum solution in more trials than CT triads (9, 17), but at the expense of having an inferior mean solution to that of CT triads (16, 17). Our findings bore out these predictions. Strikingly, as we discuss below, we found that IT triads showed the positive features of both CT and NT triads: They found the optimum solution as frequently as NT triads, but with a higher quality mean solution like CT triads.

CT triads found the optimal solution in 33.3% of trials, IT triads found the optimum in 48.3% of trials, and NT triads found the optimum in 44.1% of trials (the difference between IT and CT was significant after controlling for covariates in a logistic regression; see Table 1 for full models and *Materials and Methods* for more details). Whether or not a group found the optimum, the best solution found in IT triads and NT triads was significantly better (shorter) than the best solution found in CT triads [$\log(1 + \text{difference from optimal distance})$ in CT was worse than IT by 0.285, $P < 0.001$ and NT by 0.211, $P < 0.001$]; IT and NT triads were not statistically different.

Although social influence reduces exploration and thus depresses the quality of top solutions, it is expected to improve

Table 1. Effects of treatment on performance

Parameter	Dependent variable:			
	Optimum found logistic (1)	Best solution OLS (2)	No. of unique solutions Poisson (3)	Mean solution OLS (4)
CT	−0.900* (0.418)	0.285*** (0.053)	−0.240*** (0.059)	−0.098 (0.107)
NT	−0.444 (0.426)	0.074 (0.051)	0.103* (0.040)	0.351*** (0.104)
CT with storage	−0.584 (0.350)	0.238*** (0.047)	−0.532*** (0.061)	−0.400*** (0.121)
IT with storage	−0.541 (0.389)	0.133*** (0.048)	−0.344*** (0.054)	−0.271* (0.116)
NT with storage	−0.672* (0.328)	0.125** (0.043)	−0.121* (0.050)	0.114 (0.113)
Problem 2	2.200*** (0.306)	−0.479*** (0.032)	−0.029 (0.034)	−0.571*** (0.063)
Problem 3	1.225*** (0.279)	−0.130** (0.046)	−0.097** (0.034)	−0.038 (0.078)
Problem 4	−0.288 (0.326)	0.345*** (0.042)	−0.019 (0.033)	0.064 (0.064)
Problem 5	−0.531 (0.330)	0.270*** (0.037)	0.006 (0.034)	0.285*** (.062)
Log(prob. order)	−0.093 (0.174)	−0.042 (0.024)	−0.237*** (0.018)	−0.459*** (0.038)
Best pretest	0.102 (0.063)	−0.068*** (0.008)	−0.025** (0.008)	−0.046* (0.020)
Own pretest				−0.078*** (0.016)
Round				−0.191*** 0.009
Round ²				0.005*** 0.000
Constant	−0.865 (0.514)	1.000*** (0.074)	3.719*** (0.065)	4.371*** 0.184
Observations	514	514	514	26,214

Columns 1–3: unit of observation is a whole trial; column 4: unit of observation is a single solution. For columns 2 and 4, the dependent variable is solution distance [measured as $\log(1 + \text{difference from optimal distance})$], so lower numbers correspond to better performance. * $P < 0.05$; ** $P < 0.01$; *** $P < 0.001$.

the quality of the mean solution by allowing players with very poor solutions to adopt better solutions from their neighbors (15, 16). We find that to be the case (see Table 1, column 4). The mean solution (all solutions from all triad members across all 17 rounds of a trial) in IT triads was as good as the mean solution in CT triads. The mean solution in NT triads was worse than in IT triads [$\log(1 + \text{difference from optimal distance})$ was 0.351 longer, $P < 0.001$] and CT triads (0.449 longer, $P < 0.001$).

As expected, more social influence resulted in less diversity of solutions. The mean number of unique solutions found by a triad over all 17 rounds was highest in NT triads (30.5), followed by IT triads (27.5) and CT triads (21.4). However, the greater diversity of NT triads did not result in greater performance. Although NT triads found 1.108 times more unique solutions than IT triads (Poisson, $P = 0.010$), they did not find the optimum more frequently (indeed, NT triads found the optimum less frequently than IT triads, but the difference was not statistically significant).

IT triads displayed a balance between learning from peers (through social influence) and trying diverse new solutions (through independent exploration). In IT triads, answers within a triad alternately became more similar to each other (on rounds in which they could see each other’s answers) and became more

different from each other (on rounds in which they could not see each other's answers), exploring from new starting points (Fig. 1). This contrasts with both other treatments in which the answers within a triad largely became more similar to each other over time on average. In NT triads, answers' becoming more similar to each other reflects only independent convergence on similar answers, while in CT triads becoming more similar to each other is also the result of social influence.

As pure strategies at the individual solver level, both independent exploration and social influence can lead to "getting stuck" at a suboptimal solution. Independent exploration tends to lead to low-quality solutions for most individuals, even if there is a high chance of some single solver finding the optimum (9). Social influence can result in a premature consensus on a good solution before the optimum is found (17, 18). Alternating between independent exploration and social influence may have reduced the chances of both types of getting stuck for IT triads.

Among all three treatments, the greatest improvements in solution quality occurred in IT triads during social-influence rounds—even for leading players with no better solution to copy (Fig. 2). Improvement in the mean solution is not surprising, as low performers were able to copy higher performers on rounds with social influence. However, there was also greater improvement in the quality of the best solution in a triad on social-influence rounds than on rounds without social influence. Social influence is especially beneficial—even for leading players—when it follows independent exploration that generates more diversity.

Fig. 3 plots parts of the correct solution that the leading player could learn from (if they were visible). It shows the number of correct solution legs (legs that were part of the globally optimal solution) in leading players' solutions versus the number of correct solution legs in other solutions that were not also part of the leading solution. Leading players in IT triads were exposed to more correct legs than leading players in CT. Of course, lagging solutions in NT triads had the most correct solution legs that were not part of leading solutions, but these were never visible to leading players to learn from. Fig. 4 shows that leading players in IT made their solutions more similar to those of their neighbors




Fig. 1. Mean number of solution links matching solution links of other players by round. The maximum value this could take is 75, when all 25 links of all three players are the same. Error bars are 95% confidence intervals of the mean.




Fig. 2. Fitted values for improvement in solution distance, by round, from model specifications selected and fit by least absolute shrinkage and selection operator (LASSO) regression (31). (Left) Improvement for leading players (subject-round pairs in which there was no better solution in the triad in the previous or focal rounds). (Right) Improvement for lagging players (subject-round pairs in which there was a superior solution in the triad in the previous round).

during social-influence rounds—apparently taking advantage of that beneficial diversity.

Effects of Storing Best Solution. In addition to the above results, we ran a second set of trials evaluating the effect of another realistic condition: including a "storage" feature, in which individuals were reminded of their own best solution previous to the current round and could load it with a single mouse click. Overall, storing a solver's best solution produced results that were qualitatively similar to social influence: Relative to our first set of trials, adding storage substantially decreased exploration (the number of unique solutions was 0.748 times the number without storage for CT, 0.706 for IT, and 0.799 for NT; Poisson, $P < 0.001$ for all comparisons) but resulted in an improvement in mean performance [with storage, $\log(1 + \text{difference from optimal distance})$ was 0.303 higher in CT, $P = 0.010$; IT: 0.271, $P = 0.020$; NT: 0.237, $P = 0.009$].

The chance of finding the optimum solution is related to both mean performance (and thus the number of individuals with good solutions) and the level of exploration (thus the relative chance of improving from an already good solution). Because storage improved one precursor to finding the optimum but decreased the other, storage had different effects on the raw rate at which the different treatments found the optimum. Without storage, CT and IT had a high mean performance, and IT and NT had high exploration. Storage reduced exploration and thus eliminated a major source of high performance in IT and NT. However, storage also increased the mean, creating a simultaneous improvement for all treatments. Taken together, CT, IT, and NT triads found the optimum in 39.1, 39.3, and 38.1% of trials, respectively, representing an increase for CT but decreases for IT and NT.

To simplify, we can think of finding the optimum as most likely when a subject's solution is "in range"—that is, having a solution that can be tweaked to result in an optimum solution—and the subject continues to explore from there. Table 2 shows the raw rates of in-range rounds for each treatment condition along with the rate of improvement from in-range rounds. Rounds are considered in range if the optimum solution has not been found by a member of the triad, and the subject's current solution has 22 or 23 correct solution legs (it is impossible to have 24

correct solution legs without violating the rules of the TSP); the table presents this number as a fraction of all rounds. The rate of improvement is calculated as the fraction of in-range rounds from which the focal subject's solution improved.

Interestingly, the effect of greater improvement by top performers in IT triads on rounds with social influence (Fig. 2) is greatly reduced with the storage feature. Without the greater diversity from high exploration during rounds with independent exploration, the interplay between social influence and independent exploration did not yield any substantial benefit.

Conclusion

Intermittent breaks in interaction improve collective intelligence. Being exposed to diverse answers boosts performance, even if the answers one sees are worse than one's own. To achieve this performance boost within a triad, there is a requirement for both independent exploration (to generate diversity) and interaction (to allow social influence). Only IT triads without storage have the necessary conditions for this boost to top performance. In CT triads, leaders are exposed to others' answers, but they are not as diverse as IT triads without storage on average due to limited exploration. In NT triads, leaders are not exposed to others' answers at all.

Like constant access to others' answers, when one's own past answers can be stored such storage reduces the additional boost to performance by leaders within IT triads. For the interplay between independent exploration and social ties to be beneficial, there must be sufficient exploration during the independent phases of the problem-solving task to generate diverse solutions that lead to learning. Storage works directly against this requirement by suppressing exploration and instead encouraging relative stasis at known solutions. Without the phase of exploration, we would not expect the overall performance to be substantially different from CT triads. Indeed, the coefficients in Table 1 show broad convergence between IT and CT when storage is present.

By shaping subjects' behavior to take advantage of both independent exploration and social learning, intermittent interaction caused subjects to perform better on our complex problem-solving task. That implies, however, that task type represents a




Fig. 3. Possibility of leaders learning from others' solutions by treatment: fitted values (LASSO) for number of correct legs in leading players' solutions versus number of correct legs in other players' solutions that are not present in the focal leading player's solution. Labels indicate round numbers.




Fig. 4. Evolution of leading players' solutions: fitted values (LASSO) for number of solution legs newly matching neighbors' solutions from the previous rounds (from either copying or independent convergence on the same answers) versus legs not present in any solutions from the last round.

likely boundary condition for our results. In tasks where exploration or learning is unnecessary or impossible, we do not expect our results to hold. For example, pure coordination tasks [also known as "additive" tasks (32)], in which the quantity of distinct solutions or contributions is more important than their quality, would not necessarily reward learning. Similarly, some problem spaces are simple or "smooth" and do not require or reward extensive exploration. At the other extreme, other problem spaces may be so rugged that even arbitrarily similar solutions can be dissimilar in their quality; for such problems it would not be helpful to borrow and adapt part of a neighbor's solution.

Our results suggest new avenues for research on the importance of interaction frequency for performance. For example, how does optimal frequency change with problem complexity, social network structure, the type of outcome sought, or the baseline collective intelligence factor of the group (1)? Might our results be moderated by different forms of interaction [such as the active consensus-oriented deliberation used in the second phase of the "hybrid structure" in the brainstorming literature (30)] or different approaches to using storage? Finally, might frequency of interaction differentially affect the various component mechanisms of social influence [e.g., free riding, evaluation apprehension, and production blocking (33)]? In short, our study suggests the importance of refocusing future research on the frequency and pattern of interaction, rather than its absence or presence.*

Our main manipulation (NT, IT, or CT with storage off) reveals that intermittently present social influence achieves the beneficial aspects of both constant social influence and independence when searching complex solution spaces. Prior results showing the benefits of social influence in "wisdom of the crowd" tasks (15, 16) are due to less-confident low performers revising their solutions toward the mean after peer influence. Our results show something more: Triads find the optimum more and high performers do even better with intermittent ties, suggesting the presence of beneficial social learning for all participants, not just low performers. Indeed, intermittent social influence may mitigate the dangers inherent in both independent exploration (spending time on poor solutions) and social influence (premature consensus). Importantly, although past laboratory

*Beyond guiding future work, our finding may also permit a reinterpretation of some prior results. For example, when looked at through the lens of interaction frequency, the pattern of interaction studied in ref. 24 may be closer to IT than CT, presenting a possible explanation for the benefits of interaction described in that paper.

Table 2. Rounds in range of optimum solution

Treatment	In-range rounds	Rate of improvement
CT, storage off	0.185	0.029
IT, storage off	0.106	0.089
NT, storage off	0.080	0.071
CT, storage on	0.216	0.022
IT, storage on	0.178	0.022
NT, storage on	0.110	0.043

experimental work has focused on constant structures of social influence, real online and offline social ties are intermittent (29, 34), like our top-performing treatment.

In general this is a reassuring finding about collective intelligence in the wild but raises many questions about the design of always-on technologies that support collaborative and crowd work. Broadly speaking, productivity tools encourage people to build off of their own previous best work, and transparency-enhancing collaboration and networking tools encourage people to be in constant contact with one another. Extrapolating from our results, one could say that such technology use increases mean performance but depresses maximum performance in complex problem solving. Although much is gained from keeping people connected, even greater problem-solving performance could be achieved by redesigning technologies to intermittently turn on and off the influence that people feel from social ties and their own previous work.

Materials and Methods

Real complex problems of interest involve multiple interacting dimensions and cannot be solved by simple “hill climbing” or local search. Because of the way each aspect of a problem depends on other aspects of that problem, so-called rugged solution spaces have been widely used as models of complex problem solving and innovation (2, 17, 27). Following prior literature, we adopt a rugged solution space for our experiments.

Subjects solved examples of the Euclidean (i.e., 2D) TSP, presented visually. The TSP requires the solver to find the shortest path among a number of “cities” on a map, visiting each city except the first exactly once (see Fig. 5 for an example). Each path concludes with a return to the first city; thus the first city is visited twice.

Solution spaces for TSPs are NP-hard and characterized by many local optima (25, 26). They are also rugged in the sense used by prior social science research on problem solving (2, 17, 27)—they are impossible to solve by local hill climbing, changing one part of the solution at a time—by construction. It is impossible to modify an existing solution by changing only one leg of the journey without violating the rules of the TSP. At a minimum, two pairs of cities must be changed in tandem to move from one valid solution to another.

The task was presented via a web browser-based computer interface in a university experimental laboratory. Subjects were recruited from the university’s experimental subject pool. Informed consent was obtained from subjects before participation. This study was approved by the Harvard University Committee on the Use of Human Subjects.

Each subject completed each of six different TSP maps. Due to a programming error, results from map 6 were not comparable and thus we analyze only results from maps 1 to 5. Each map consisted of 25 cities and thus required 25 separate legs of the overall journey to complete (the last leg connects the final city back to the city the player started at).

To complete the task, subjects were asked to click with the computer’s mouse on the city icons in the sequence corresponding to the path they wished to submit as their answer. For example, in Fig. 5, the subject would have clicked on city A, then S, then B, then L, and so on. The computer program would draw a line segment connecting each pair of cities as soon as the second city in each leg of the journey had been clicked on.

A single trial consisted of 17 tries (rounds) to solve a single TSP problem. Thus, subjects could try up to 17 different solutions to the same problem. In rounds 2–17, subjects could see the solution they entered in the previous round, in a smaller window below the main task window, along with its distance. Each round was limited to 50 s, as indicated by a countdown timer on the left-hand side of the interface. If the subject had not submitted an answer by the end of the 50 s, the interface automatically moved on

to the next round. Any solution or partial solution that had already been entered before the timeout was recorded, but the answer was considered incomplete for the purposes of payment. An entire trial lasted at most 14 min 10 s. In a given trial, all three subjects had completed the same number of previous trials; for example, if a subject had completed two trials already, on the next trial she would have been placed in a triad with two others who had also done exactly two trials already.

Subjects were randomly assigned in sets of three individuals (triads)—a minimal but representative model of problem solving in groups (35)—to one of three treatment conditions. In the CT treatment, subjects could see the previous round’s solutions from the other subjects in their triad (“neighbors” for short) along with the distances of those solutions below the main task window (in addition to their own previous solution and its distance). In the IT condition, subjects could see their neighbors’ solutions only on rounds 4, 7, 10, 13, and 16. In the NT condition, neighbors’ solutions were never visible. Subjects were randomly matched with and anonymous to each other; the order of treatments was also randomly assigned and subjects were exposed to multiple treatments. In regressions, we used cluster-robust standard errors as follows: In Table 1, models 1 and 2, we clustered on the identity of the best performer in each triad, where best is defined as the person who found the most correct legs, or in the case of a tie, the person who found the best-quality solution earliest, or in case of another tie, the person with the highest score on the pretest assessment. In model 3, we cluster by identity of the person who changes their solution most frequently on average across all trials. In model 4, we cluster on subject and group: In CT and IT, a group was equivalent to a triad; in the NT condition, subjects did not interact at all, and thus a group was coded as consisting only of each individual person.

In the storage condition, subjects had an image corresponding to their own previous best solution, along with its distance, in addition to information about their own previous solution and information about their neighbors’ solutions, if applicable. Subjects could load their previous best solution by clicking on it. After loading, subjects could edit it or simply submit it as is.

Before any full trial’s beginning, each subject took a nine-problem pretest, in part to train them on the TSP and in part to assess their individual abilities with respect to solving TSPs. Subjects were paid \$10 for showing up to the experiment, \$1 for each pretest problem they found the optimum for, and 50 cents per round during an experimental trial in which they found the optimum. The maximum total payment that was theoretically possible was \$70, but in practice the interquartile range of payment was \$15 to \$23 and the maximum payment to a single subject was \$36.50. Three of the five problems had more than one optimal solution (i.e., more than one solution achieved the minimal distance). The quality of solutions was recorded as a distance (to be minimized) and as a number of correct legs. A leg was considered correct if it was part of an optimal solution.




Fig. 5. An example TSP from the experiment with the optimal solution filled in. To the left is a timer (showing 32 s remaining) along with the reset and submit buttons. The last leg of the journey (from city J to city A) was filled in automatically by the computer to create a closed loop.

When using solution distance as a dependent variable (DV), we subtract the optimal solution distance from the subject's answer to facilitate comparison across problems. Incomplete solutions did not have a well-defined distance and were replaced with a 99.9th percentile (i.e., very long) distance. Models 2 and 4 use solution distance as the DV, but our conclusions also hold when using the number of correct solution legs as the DV.

To better understand the mechanisms underlying our treatment effects, we fit more detailed models of individual performance using LASSO regression (31). LASSO is a method of penalized regression that both fits and selects parameters subject to the constraint that the sum of fitted parameter values is less than or equal to a regularization term consisting of a

tunable parameter (λ) times the ℓ_1 norm of the fitted parameters. We chose the value of λ such that cross-validation error was minimized. LASSO has the tendency to select a reduced-form model consisting only of variables with high predictive value by setting less-important parameter values to zero. This has the desirable consequences of reducing overfitting and researcher degrees of freedom to choose models that fit preconceived notions of what is going on in the data.

ACKNOWLEDGMENTS. This research was supported by Office of Naval Research Grant G00005072 and funding from the Harvard Business School Division of Research and Faculty Development. The viewpoints expressed in this article are the authors' alone.

1. Woolley AW, Chabris CF, Pentland A, Hashmi N, Malone TW (2010) Evidence for a collective intelligence factor in the performance of human groups. *Science* 330: 686–688.
2. Mason W, Watts DJ (2012) Collaborative learning in networks. *Proc Natl Acad Sci USA* 109:764–769.
3. Galton F (1907) Vox populi (the wisdom of crowds). *Nature* 75:450–451.
4. Malone TW, Laubacher R, Dellarocas C (2010) The collective intelligence genome. *MIT Sloan Manage Rev* 51:21.
5. Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proc Natl Acad Sci USA* 108:9020–9025.
6. Valentine MA, et al. (2017) Flash organizations: Crowdsourcing complex work by structuring crowds as organizations. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (ACM, New York), pp 3523–3537.
7. Pan W, Altshuler Y, Pentland A (2012) Decoding social influence and the wisdom of the crowd in financial trading network. *Proceedings of the 2012 ASE/IEEE International Conference on Social Computing and 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust* (IEEE, Piscataway, NJ), pp 203–209.
8. Arrow KJ, et al. (2008) The promise of prediction markets. *Science* 320:877–878.
9. Boudreau KJ, Lacetera N, Lakhani KR (2011) Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Sci* 57:843–863.
10. Sunstein CR (2005) *Why Societies Need Dissent* (Harvard Univ Press, Cambridge, MA), Vol 9.
11. Noveck BS (2009) *Wiki Government: How Technology Can Make Government Better, Democracy Stronger, and Citizens More Powerful* (Brookings Institution, Washington, DC).
12. Berdahl A, Torney CJ, Ioannou CC, Faria JJ, Couzin ID (2013) Emergent sensing of complex environments by mobile animal groups. *Science* 339:574–576.
13. Centola D (2010) The spread of behavior in an online social network experiment. *Science* 329:1194–1197.
14. Shore J, Bernstein E, Lazer D (2015) Facts and figuring: An experimental investigation of network structure and performance in information and solution spaces. *Organ Sci* 26:1432–1446.
15. Madirolas G, de Polavieja GG (2014) Wisdom of the confidant: Using social interactions to eliminate the bias in wisdom of the crowds. *Proceedings of the Collective Intelligence Conference* (ACM, New York), pp 10–12.
16. Becker J, Brackbill D, Centola D (2017) Network dynamics of social influence in the wisdom of crowds. *Proc Natl Acad Sci USA* 114:E5070–E5076.
17. Lazer D, Friedman A (2007) The network structure of exploration and exploitation. *Administrative Sci Q* 52:667–694.
18. Mason WA, Jones A, Goldstone RL (2008) Propagation of innovations in networked groups. *J Exp Psychol Gen* 137:422–433.
19. Hong L, Page SE (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc Natl Acad Sci USA* 101:16385–16389.
20. Paulus PB, Putman VL, Dugosh KL, Dzindolet MT, Coskun H (2002) Social and cognitive influences in group brainstorming: Predicting production gains and losses. *Eur Rev Soc Psychol* 12:299–325.
21. Liang DW, Moreland R, Argote L (1995) Group versus individual training and group performance: The mediating role of transactive memory. *Personal Social Psychol Bull* 21:384–393.
22. Hackman JR (2011) *Collaborative Intelligence: Using Teams to Solve Hard Problems* (Berrett-Koehler Publishers, Oakland, CA).
23. Schumpeter J, Backhaus J (2003) *The theory of economic development. Joseph Alois Schumpeter: Entrepreneurship, Style and Vision*. The European Heritage in Economics and the Social Sciences, ed Backhaus J (Springer, New York), Vol 1, pp 61–116.
24. Boudreau KJ, Lakhani KR (2015) “Open” disclosure of innovations, incentives and follow-on reuse: Theory on processes of cumulative innovation and a field experiment in computational biology. *Res Policy* 44:4–19.
25. Papadimitriou CH (1977) The euclidean travelling salesman problem is np-complete. *Theor Computer Sci* 4:237–244.
26. Černý V (1985) Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *J optimization Theor Appl* 45:41–51.
27. Levinthal DA (1997) Adaptation on rugged landscapes. *Management Sci* 43:934–950.
28. MacGregor JN, Ormerod T (1996) Human performance on the traveling salesman problem. *Percept Psychophys* 58:527–539.
29. Sekara V, Stopczynski A, Lehmann S (2016) Fundamental structures of dynamic social networks. *Proc Natl Acad Sci USA* 113:9977–9982.
30. Girotra K, Terwiesch C, Ulrich KT (2010) Idea generation and the quality of the best idea. *Management Sci* 56:591–605.
31. Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc* 58: 267–288.
32. Steiner ID (2007) *Group Process and Productivity* (Academic, New York).
33. Diehl M, Stroebe W (1987) Productivity loss in brainstorming groups: Toward the solution of a riddle. *J Personal Soc Psychol* 53:497–509.
34. Riedl C, Woolley AW (2017) Teams vs. crowds: A field test of the relative contribution of incentives, member ability, and emergent collaboration to crowd-based problem solving performance. *Acad Management Discoveries* 3:382–403.
35. Hackman JR, Vidmar N (1970) Effects of size and task type on group performance and member reactions. *Sociometry* 33:1–37.