Understanding the group dynamics and success of teams

Michael Klug[1] and James P. Bagrow[1,2,3,*]

[1]Department of Mathematics & Statistics, The University of Vermont

[2]Vermont Complex Systems Center, The University of Vermont

[3]Vermont Advanced Computing Core, The University of Vermont, Burlington, VT, USA

[*]To whom correspondence should be addressed. E-mail: james.bagrow@uvm.edu

September 16, 2015

Abstract

Complex problems often require coordinated group effort and can consume significant resources, yet our understanding of how teams form and succeed has been limited by a lack of large-scale, quantitative data. We analyze activity traces and success levels for ∼150,000 self-organized, online team projects. While larger teams tend to be more successful, workload is highly skewed across the team, with only a few members performing most work. Focused work activity may indicate that larger teams succeed by acting as a support system for a smaller set of core members. However, it may also simply be due to social loafing, where individuals perform less work in a group than they would on their own. We find that highly successful teams are both significantly more focused than average teams of the same size and their members, even non-core members, are more likely to themselves be core members of other teams. This mixture of size, focus and overlapping team experience cannot be explained by social loafing and points to mechanisms that can maximize the success of collaborative endeavors.

Introduction

Massive datasets describing the activity patterns of large human populations now provide researchers with rich opportunities to quantitatively study human dynamics [1, 2], including the activities of groups or teams [3, 4]. With the rise in prominence of network science [5, 6], much effort has gone into discovering meaningful groups within social networks [7, 8, 9, 10, 11, 12, 13, 14] and quantifying their evolution [15, 14]. Teams are increasingly important in research and industrial efforts [16, 17, 18, 3, 4, 19, 20] and small, coordinated groups are a significant component of modern human conflict [21, 22]. Recently, there has been much debate on the "group size hypothesis", that larger groups are more robust or perform better than smaller ones [23, 24, 25, 26]. Scholars of science have noted for decades that collaborative research teams have been growing in size and importance [27, 28, 29, 19]. At the same time, however, social loafing, where individuals apply less effort to a task when they are in a group than when they are alone, may counterbalance the effectiveness of larger teams [30, 31, 32]. Meanwhile, case studies show that leadership [33, 3, 34, 35] and experience [36, 37] are key components of successful team outcomes, while specialization and multitasking are important but potentially error-prone mechanisms for dealing with complexity and cognitive overload [38, 39]. New tools, including electronic sensor systems, can quantify team activity and performance [40, 4]. In all of these areas, large-scale, quantitative data can push the study of teams forward.

Users of the GitHub web platform can form teams to work on real-world projects, primarily software development but also music, literature, design work, and more. A number of important scientific computing resources are developed through GitHub, including astronomical software, genetic sequencing tools, and key components of the Compact Muon Solenoid experiment's data pipeline.[1] A "GitHub for science" initiative has been launched[2] and GitHub is becoming the dominant service for open scientific development.

GitHub provides rich public data on team activities, including when new teams form, when members join existing teams, and when a team's project is updated. These open collaborations evolve in an entirely self-organized manner and are not driven by hierarchical management as one may encounter at a commercial organization. GitHub also provides social media tools for the discovery of interesting projects. Users who see the work of a team can choose to flag it as interesting to them by "starring" it. The number of these "stargazers" S allows us to accurately quantify the **success** of the team, in a manner analogous to the use of citations of research literature as a proxy for "impact" [41]. As with bibliometric impact, one should be cautious and not consider success to be a perfectly accurate measure of *quality*, something that is far more

[1] For examples, see `https://github.com/showcases/science`.
[2] See `https://github.com/blog/1840-improving-github-for-science`.

difficult to objectively quantify.

In this study, we analyze the memberships and activities of approximately 150,000 self-organized teams, as they perform real-world tasks, to uncover the blend of features that leads to success. To the best of our knowledge this is the largest study of real-world team performance to date. We present results that demonstrate (i) how teams distribute work activity across their members, (ii) the mixture of experiential diversity and collective leadership roles in teams, and (iii) how successful teams are different from other teams. By using a subset of the teams to control for team lifetime we also show that successful teams tend to gain new members before gaining success, i.e., the team brings success more than success attracts the team.

Results

To understand teams and the collaboration network between them, we began by characterizing their sizes and degrees of overlap. The distributions of team size M and teams per person k (Fig. 1A) both display truncation and significant rollover. Teams were relatively small, with only 0.46% of teams having $M > 10$ members. A dashed line in Fig. 1A indicates a pure power law $P(x) \sim x^{-3.5}$ and provides a guide for the eye. Such a high exponent emphasizes the degree of truncation we observed: while broader than a gaussian distribution, both M and k reflect underlying constraints on time and effort that prevent scaling. In contrast, the distribution of success S, the number of times users flag a project as interesting (see Methods), was very broadly distributed. The shallower guide shows a power law with exponent 1.5. This broad distribution of S indicates the presence of significant preferential attachment [42] or rich-get-richer effects, as may be captured by a Simon model [43]. In the Supporting Information (SI) we reinforce these observations with a full statistical analysis, comparing multiple heavy-tailed models to the data.

Despite the apparent drive towards smaller teams, there was a positive and significant relationship ($p < 10^{-10}$, rank correlation $\rho = 0.0845$) between the size of a team and its success, with 300% greater success on average for teams of size $M = 10$ compared to solos with $M = 1$ (Fig. 1B). Larger teams tend to also have more success but, while the trend was highly significant, ρ indicates that there remains considerable variation in S that is not captured by team size alone.

Our next analysis reveals an important relationship between team focus and success. Unlike bibliographic studies, where teams can only be quantified as the listed coauthors of a paper, the data here allow us

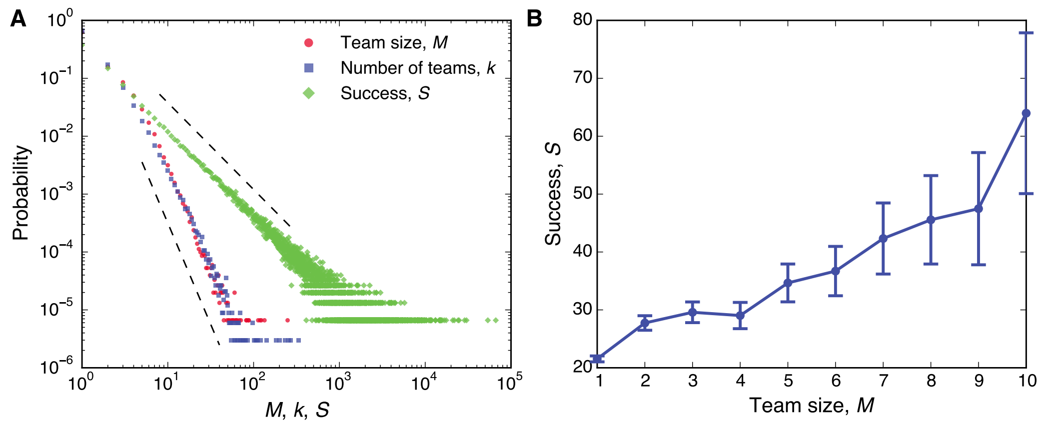

Figure 1: Summary of team data. (**A**) The distributions of: the number of individuals per team M, the number of teams k a person simultaneously belongs to, and success S. Both M and k are truncated; only 0.46% of teams have $M > 10$. In contrast, impact, measured by the popularity of the team (see Materials and Methods), is heavily skewed. Dashed lines denote power laws $P(x) \sim x^{-\alpha}$ with $\alpha = 3.5$ near the steep distributions $P(M)$ and $P(k)$, and $\alpha = 1.5$ near the broad distribution $P(S)$. (**B**) Despite the overabundance of small teams, larger teams have significantly more success on average, with a 300% increase in S as M goes from 1 to 10. Errorbars here and throughout denote ±1.96 s.e.

to measure the intrinsic work or volume of contributions from each team member to the project. For each team we measured the contribution w_r of a member to the team's ongoing project, how many times that member updated the project (see Methods). Team members were ranked by contribution, so w_1 counts the work of the member who contributed the most, w_2 the second heaviest contributor, and so forth. The total work of a team is $W = \sum_{r=1}^{M} w_r$.

We found that the distribution of work over team members showed significant skew, with w_1 often more than 2–3 times greater than w_2 (Fig. 2A and SI). This means that the workloads of projects are predominantly carried by a handful of team members, or even just a single person. Larger teams perform more total work, and the heaviest contributor carries much of that effort: the inset of Fig. 2A shows that w_1/W, the fraction of work carried by the rank one member, falls slowly with team size, and is typically far removed from the lower bound of equal work among all team members. See SI for more details.

This focus in work activity indicates that the majority of the team serves as a support system for a core set of members. Does this arrangement play a role in whether or not teams are successful? We investigated this in several ways. First, we asked whether or not a team was **dominated**, meaning that the lead member contributed more work than all other members combined ($w_1/W > 1/2$). Highly successful "top" teams, those in the top 10% of the success distribution, were significantly more likely to be dominated than average teams, those in the middle 20% of impact (Fig. 2B). Given that it requires more effort to dominate a larger team than a smaller one, the fact that dominated teams were so prevalent in the high-success regime emphasizes a strong need for workload focus.

Next, we moved beyond the effects of the heaviest contributor by performing the following analysis. For each team we computed its **effective** team size m, directly accounting for the skew in workload (see Methods for full details). This effective size can be roughly thought of as the average number of unique contributors per unit time and need not be a whole number. For example, a team of size $M = 2$ where both members contribute equally will have effective size $m = 2$, but if one member is responsible for 95% of the work the team would have $m \approx 1.22$.

Figure 2C shows that (i) teams are effectively much smaller than their total size would indicate, for all sizes $M > 1$, and (ii) top teams are significantly smaller in effective size (and therefore more focused in their work distribution) than average teams with the same M. Further, success is significantly, negatively correlated with m, for all M (Fig. 2D). More focused teams have significantly more success that less focused teams, regardless of total team size. Focused work activity appears to be a crucial factor of success.

Further analyses revealed the importance of team composition and its role in team success. The teams being studied are self-organized, and therefore it seems reasonable for successful teams to have evolved the key ingredients of success in an emergent fashion. Team members do not perform their work in a vacuum, they each bring experiences from their work histories. Often members of one team belong to other teams at the same time. We investigated these facets of a team's composition by exploring (i) how many projects team members have worked on, (ii) how diverse are the other teams that members belonged to, and (iii) how many team members were "leads" of other teams.

Experience E, the average number of other projects that team members have worked on (see Methods), was significantly related to success. However, the trend was not particularly strong (see SI) and, as we later show via combined modeling efforts, this relationship with success was entirely explainable by the teams' other measurable quantities.

It may be that the past volume of work does not contribute much to the success of a team, but this seems

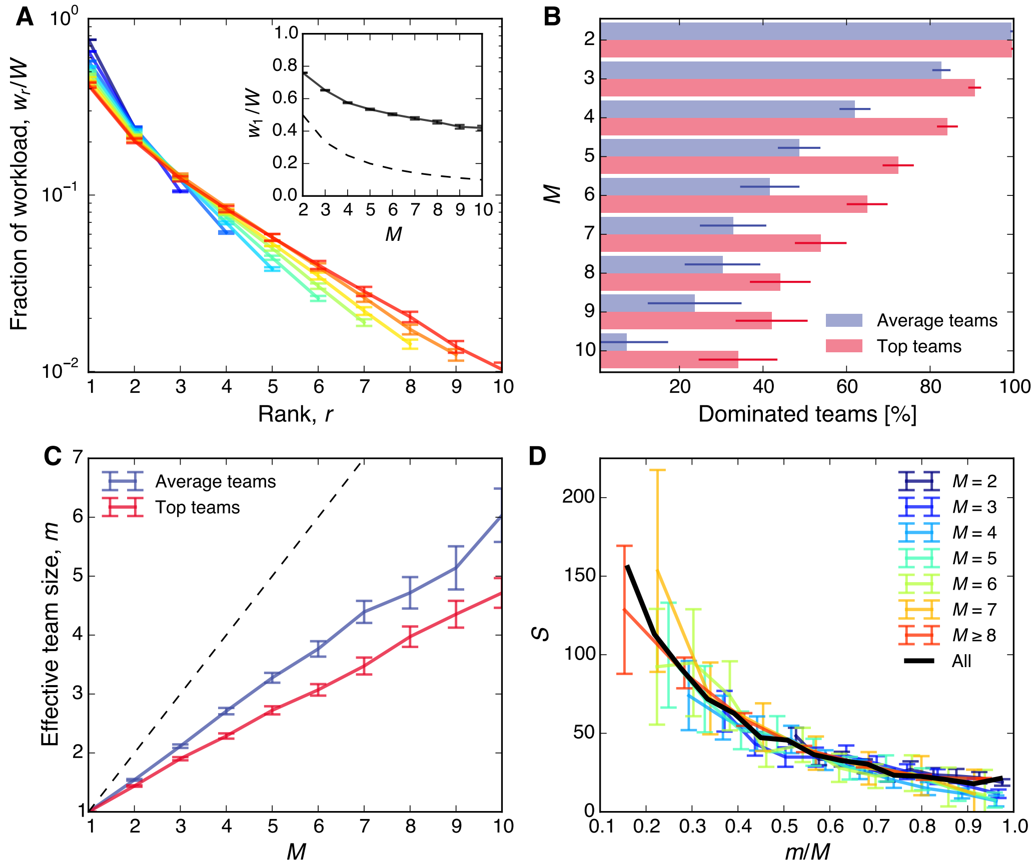

Figure 2: Teams are focused, and top teams are more focused than average teams of the same size. (**A**) The average fraction of work w_r/W performed by the r-th most active member, where W is the total work of the team, for different size teams. Larger teams perform more work overall, but the majority of work is always done by a small subset of the M members (note the logarithmic axis). Inset: The fraction of work performed by the most active team member is always high, often larger than half the total. The dashed line indicates the lower bound of uniform work distribution, $w_r/W = 1/M$. (**B**) A team is **dominated** when the most active member does more work than all other members combined. Top teams, the 10% most successful teams, are significantly more likely to be dominated than average teams, those in the middle 20%. (**C**) The effective team size m (see Methods), a measure that accounts for the skewed distribution of work in A, is significantly smaller than M. Moreover, top teams are significantly more focused, having smaller effective sizes, than average teams at all sizes $M > 1$. The dashed line denotes the upper bound $m = M$. (**D**) Success is universally higher for teams with smaller m/M, independent of M, further supporting the importance of focused workloads. The solid lines indicates the average trend for all teams $2 \leq M \leq 10$.

to contradict previous studies on the importance of experience and wisdom [36, 37]. To investigate, we turned to a different facet of a team's composition, the diversity of the team's background. Successful teams may tend to be comprised of members who have frequently worked together on the same projects in the

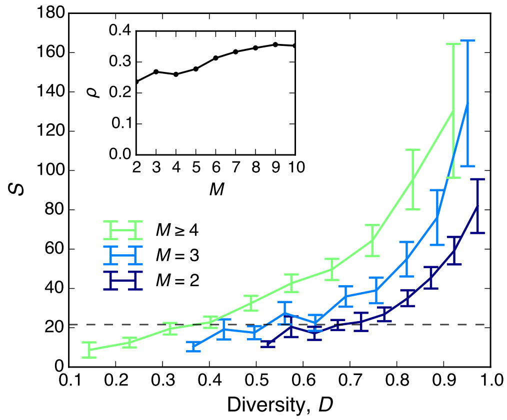

Figure 3: Teams whose members belong to more diverse sets of other teams tend to be more successful, regardless of team size. (Inset) The rank correlation ρ between diversity and success grows with team size. Larger teams benefit more from diversity than smaller teams.

past, perhaps developing an experiential shorthand. Conversely, successful teams may instead have multiple distinct viewpoints, solving challenges with a multi-disciplinary perspective [44].

To capture the distinctness of team member backgrounds, the diversity D was measured as the fraction of projects that team members have worked on that are unique (see Methods). Diversity is low when all M members have worked on the same projects together ($D = 1/M$), but D grows closer to 1 as their backgrounds become increasingly diverse. A high team diversity was significantly correlated with success, regardless of team size (Fig. 3). Even small teams seem to have benefited greatly from diversity: high-D duos averaged nearly eight times the success of low-D duos. The relationship between D and S was even stronger for larger teams (Fig. 3 inset), implying that larger teams can more effectively translate this diversity into success. Even if the raw volume of experience a team has does not play a significant role in the team's success, the diversity of that experience is a key component of team success.

Considerable attention has been paid recently to collective leadership, where decision-making structures emerge from the mass of the group instead of being imposed via a top-down hierarchy [33, 35]. The open collaborations studied here have the potential to display collective leadership due to their volunteer-driven, self-organized nature. The heaviest contributor to a team is most likely to occupy such a leadership role. Further, since teams overlap, a secondary member of one team may be the "lead," or heaviest contributor to another. This poses an interesting question: Even though teams are heavily focused, are teams more suc-

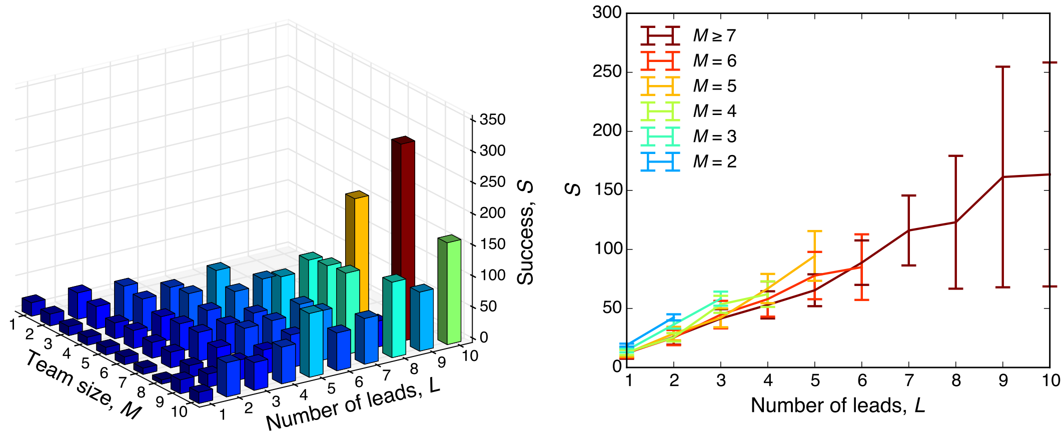

Figure 4: Teams with more leads have higher success than teams of the same size with fewer leads. A lead is someone who contributes more work to at least one team he or she belongs to than any other members of that team.

cessful when they contain many leads, or few? A team with many leads will bring considerable experience, but its members may also lack the time and resources necessary to dedicate their attention to the team.

To answer this, we measured L, the number of team members who are the lead of at least one team ($1 \leq L \leq M$, see Methods), and found that teams with many leads have significantly higher success than teams of the same size with fewer leads (Fig. 4). Despite the limited support a team member can provide when he or she is occupied with other work, successful teams tend to arrange their members in exactly this fashion. Of course, the strong focus in work activity (Fig. 2) is likely interrelated with these observations. However, we will soon show that both remain significantly related to success in combined models.

Expanding on this observation, Table 1 illustrates the extreme case of teams with a single lead ($L = 1$) compared with teams comprised entirely of leads ($L = M$). The latter always displayed significantly higher success than the former (MWU test, see table), independent of team size, underscoring the correlations displayed in Fig. 4. Often the difference was massive: teams of size $M = 7$, for example, averaged more than 1200% higher success when $L = 7$ than when $L = 1$.

These results on experiential diversity and leadership amplify our findings on team focus, and augment important existing research [36, 3, 4, 35, 44]. We have found that teams tend to do best when optimized along all three of these dimensions. Of course, it is necessary to explore the joint effects of quantities, which we will do with linear and nonlinear statistical modeling.

Table 1: Teams composed entirely of leads ($L = M$) are significantly more successful (MWU test on S) than teams with one lead ($L = 1$), regardless of team size M.

	N		S		
M	$L = 1$	$L = M$	$L = 1$	$L = M$	p
2	14823	8894	18.9	42.5	$< 10^{-213}$
3	6171	2261	14.5	58.3	$< 10^{-210}$
4	3063	717	12.8	62.1	$< 10^{-112}$
5	1489	289	12.1	94.5	$< 10^{-55}$
6	740	124	12.3	85.0	$< 10^{-36}$
7	350	46	9.8	120.5	$< 10^{-15}$
8	179	19[a]	7.5	224.1	$< 10^{-8}$
9	125	9	22.2	316.8	< 0.008
10	66	6	17.8	163.5	< 0.005

[a]When $M \geq 8$, the number of teams with $L = M$ is too small ($N < 20$) for us to conclude the difference in S is significant.

Combined models

One important aspect of the individual team measurements is that they do not exist in isolation. For example, successful teams also have high work activity (high W). This can correlate with effective team size m since the potential inequality between team members can grow as their total activity grows. In other words, we need to see how our team measures relate to success together.

To understand the relative effects of these team composition measures, we performed two analyses. First, a linear regression model of success as a function of all explored measures is shown in Table 2. Not only did this regression allow us to determine whether a variable was significant or if it was confounded by the other measures, but the coefficients (on the standardized variables) let us measure the relative strengths of each variable. Team size M, effective team size m, and lead number L play the strongest roles in team success, and all three are significant in the presence of the other variables. The coefficient on m was negative while for M it was positive, further underscoring our result that, while teams should be big, they effectively should be small. Next, the total work W done on the project, followed by the diversity D of the team, were also significant measures related to success. Finally, overall team experience E was not significant in this model ($p > 0.1$). We conclude that, while S and E are correlated by themselves, the effects of E are entirely explained by the other quantities.

Second, we applied a nonlinear modeling technique called symbolic regression [45]. Symbolic re-

Table 2: Combined robust regression model on team success, $S = \alpha + \beta_M M + \beta_m m + \beta_W W + \beta_E E + \beta_D D + \beta_L L$.

Variable x	Coefficient β_x[a]		p
constant	2.553×10^{-13}	$\pm\, 0.002553$	1
Team size, M	0.1212	$\pm\, 0.007662$	$< 10^{-55}$
Eff. team size, m	−0.1266	$\pm\, 0.006063$	$< 10^{-95}$
Total work, W	0.1025	$\pm\, 0.002659$	0
Experience, E	−0.004185	$\pm\, 0.002576$	0.1043
Diversity, D	0.07516	$\pm\, 0.004427$	$< 10^{-63}$
Num. of leads, L	0.1283	$\pm\, 0.003642$	$< 10^{-270}$

[a] Variables are standardized for comparison such that a coefficient β_x implies that increasing a variable x by one standard deviation σ_x corresponds to a $\beta_x \sigma_S$ increase in S, holding other variables fixed.

gression evolves families of equations to best fit data. Our goal was to find equations of the form $S = f(M, m, W, E, D, L)$. The overall best model, in terms of balancing complexity and accuracy, was

$$S = \frac{M + DL^{3.50D}}{m} + (0.000868W - 0.0907)\, L^{3.03D}. \tag{1}$$

In this equation, success increases when M, W, D, and L increase, and decreases when m increases; experience E has no effect. In fact, E never appeared in any discovered model (see SI). This aligns well with the linear model. (How the second term contributes to changes in S depends on the magnitudes of W vs. L and D, so we performed a variable sensitivity study of this and all other discovered equations; see SI.)

Which comes first: the team or its success?

The collaborations under study are open and ongoing in that users can monitor the course of a team's project. The success of the team may grow over time, as it makes improvements or progresses towards its goals. It is then natural to ask, are skilled team members attracted to join a highly successful project, or does their presence cause the team to later attract success?

To understand this, we conducted an experiment using a subset of the teams under study. Each of these $n = 962$ teams formed at approximately the same time, and each was actively updated with multiple members ($M > 1$) at the end of our data window. This ensures that these teams have comparable lifespans, and their entire evolution was available to us.

We tracked the times when new members joined the team, and when S increased (someone from the

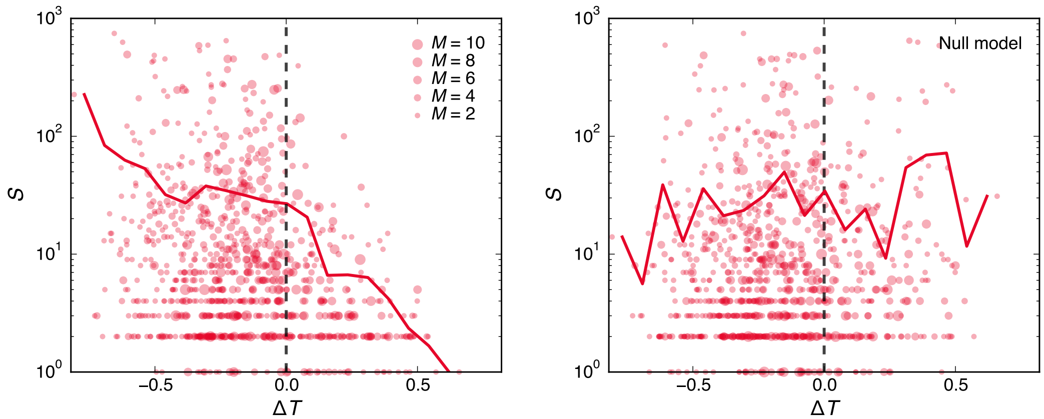

Figure 5: Teams that primarily form before they have success ($\Delta T < 0$) tend to be significantly more successful (MWU test, $p < 10^{-17}$) than teams that form after achieving success ($\Delta T > 0$), independent of team size. The solid line denotes the average trend. Also shown is a null model where the success values were randomly shuffled between teams while preserving all other features including the distributions of S and ΔT. The strong linear trend observed in the real data (left) is entirely absent in the randomized control (right).

greater community flagged the project as interesting). For cross-comparison, these times were linearly normalized on a per-team basis to fall in $[0, 1]$, where $t = 0$ corresponds to the creation of the team, and $t = 1$ the most recent update made by the team in our data window. For each team we measured ΔT, the average of the times when members joined minus the average of the times when success was received. If $\Delta T > 0$, the team tends to receive success earlier than new members; for $\Delta T < 0$, the team tends to grow in membership before success.

Teams showed a strong trend towards membership formation predating success, with high success teams almost always gaining members well before they gain their success (Fig. 5). Teams with $\Delta T < 0$ have a significantly higher median success $S = 8$ compared to $S = 3$ for teams with $\Delta T > 0$ (MWU test, $p < 10^{-17}$). The average trend present in the real data is absent in a randomized null model (Fig. 5).

One explanation for this phenomenon is that the members of high success teams are more capable of estimating a team's future potential than other users, and so can anticipate the best teams to join. Of course, establishing causality is difficult: a confounding factor may have caused the team members to join the team and then later caused the project to become popular. But this result does indicate that one argument is less likely: success does not attract the team.

Discussion

There has been considerable debate concerning the benefits of specialization compared to diversity in the workplace and other sectors [38]. Our discoveries here show that a high-success team forms a diverse support system for a specialist core, indicating that both specialization and diversity contribute to innovation and success. Team members should be both specialists, acting as the lead contributor to a team, and generalists, offering ancillary support for teams led by another member. This has implications when organizations are designing teams and wish to maximize their success. Teams tend to do best on average when they maximize M (Fig. 1B) while minimizing m (Fig. 2D) and maximizing D (Fig. 3) and L (Fig. 4).

The negative relationship between effective team size m and success S (as well as the significantly higher presence of dominated teams among high success teams) further belies the myth of multitasking [38] and supports the "surgical team" arguments of Brooks [16]. Focused work activity, often by even a single person, is a hallmark of successful teams. This focus both limits the cognitive costs of task switching, and lowers communication and coordination barriers, since so much work is being accomplished by one or only a few individuals. Such focus could also be explained by social loafing, yet loafing does not explain the correlation between leads and success (Fig. 4).

Some tasks are too large for a single person or small team to handle, necessitating the need for mega-teams of hundreds or even thousands of members. Our results indicate that such teams may be most effective when broken down into large numbers of small, overlapping groups, where all individuals belong to a few teams and are the lead of at least one. Doing so will help maximize the experiential diversity of each sub-team, while ensuring each team has someone "in charge". An important open question is what are the best ways to design such pervasively overlapping groups [13], a task that may be project- or domain-specific but which is worth further exploration.

Methods and Materials

Dataset

Public GitHub data covering 1 January 2013 to 1 April 2014 was collected from githubarchive.org in April 2014. These activity traces contain approximately 110M unique events, including whenever users create, join, or update projects. Projects on GitHub are called "repositories". A team is the set of users who can directly update ("push to") a repository. Activity or workload W is measured by the number of pushes. Users on GitHub can bookmark projects they find interesting. This is called "stargazing". We take the maximum number of stargazers for a team as its measure of success S. To avoid abandoned projects, studied teams have at least one stargazer ($S > 0$) and at least two updates per month on average (151,542 teams). To mitigate outlier effects when averaging and modeling (excluding Fig. 1A), projects above the 99th percentile of S and the 99.99th of W were removed. Since very few teams have $M > 10$, those teams were not studied. This leaves $N = 148,272$ teams. A total of 275,387 unique users participated in these teams.

Effective Team Size

The number of team members, M, does not fully represent the size of a team since workloads are highly skewed across team members. To capture the *effective* team size m, accounting for the relative contribution levels of members, we use $m = 2^H$, where $H = -\sum_{i=1}^{M} f_i \log_2 f_i$, and $f_i = w_i/W$ is the fraction of work performed by team member i. This gives $m = M$ when all $f_i = 1/M$, as expected. This simple, entropic measure is known as perplexity in linguistics and is closely related to species diversity indices used in ecology and the Herfindahl-Hirschman Index used in economics.

Experience, diversity, and leads

Denote with R_i the set of teams that user i belongs to. (Teams in R_i need at least twice-monthly updates on average, as before, but may have $S = 0$ so as to capture i's full background, not just successful projects.) We measure the experience E of a team of size M as

$$E = \frac{1}{M} \sum_i |R_i| - 1$$

and the diversity D as

$$D = \frac{\left|\bigcup_i R_i\right|}{\sum_i |R_i|}.$$

where the sum and union run over the M members of the team. Note that $D \in [1/M, 1)$. Someone is a *lead* when, for at least one team they belong to, they contribute more work to that team than any other member.

A non-lead member of team j may be the lead of team $k \neq j$. The number of leads L_k in team k is:

$$L_k = \sum_{i=1}^{M_k} \min\left(\sum_j L_{ij}, 1\right),$$

where $L_{ij} = 1$ if user i is the lead of team j, and zero otherwise. The first sum runs over the M_k members of team k, the second runs over all teams j.

Acknowledgments

We thank Josh Bongard, Brian Tivnan, Paul Hines, Michael Szell, and Albert-László Barabási for useful discussions, and we gratefully acknowledge the computational resources provided by the Vermont Advanced Computing Core, supported by NASA (NNX-08AO96G).

References

[1] David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert-László Barabási, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, et al. Life in the network: the coming age of computational social science. *Science*, 323(5915):721, 2009.

[2] Alessandro Vespignani. Modelling dynamical processes in complex socio-technical systems. *Nature Phys.*, 8(1):32–39, 2012.

[3] John RG Dyer, Anders Johansson, Dirk Helbing, Iain D Couzin, and Jens Krause. Leadership, consensus decision making and collective behaviour in humans. *Phil. Trans. R. Soc. B*, 364(1518):781–789, 2009.

[4] Alex Pentland. The new science of building great teams. *Harvard Bus. Rev.*, 90(4):60–69, 2012.

[5] Albert-László Barabási. *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume, 2003.

[6] M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.

[7] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*. Cambridge University Press, 1994.

[8] Raymond T Sparrowe, Robert C Liden, Sandy J Wayne, and Maria L Kraimer. Social networks and the performance of individuals and groups. *Acad. Manage. J.*, 44(2):316–325, 2001.

[9] Michelle Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proc. Natl. Acad. Sci. U. S. A.*, 99(12):7821–7826, 2002.

[10] Duncan J Watts, Peter Sheridan Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296(5571):1302–1305, 2002.

[11] M. E. J. Newman and Juyong Park. Why social networks are different from other types of networks. *Phys. Rev. E*, 68(3):036122, 2003.

[12] Christian Bird, David Pattison, Raissa D'Souza, Vladimir Filkov, and Premkumar Devanbu. Latent social structure in open source projects. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, SIGSOFT '08/FSE-16, pages 24–35. ACM, 2008.

[13] Yong-Yeol Ahn, James P Bagrow, and Sune Lehmann. Link communities reveal multiscale complexity in networks. *Nature*, 466(7307):761–764, 2010.

[14] Peter J Mucha, Thomas Richardson, Kevin Macon, Mason A Porter, and Jukka-Pekka Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *Science*, 328(5980):876–878, 2010.

[15] Gergely Palla, Albert-László Barabási, and Tamás Vicsek. Quantifying social group evolution. *Nature*, 446(7136):664–667, 2007.

[16] Frederick P Brooks Jr. *The Mythical Man-Month*. Pearson Education, 1995.

[17] Samer Faraj and Lee Sproull. Coordinating expertise in software development teams. *Management science*, 46(12):1554–1568, 2000.

[18] Linda Rising and Norman S Janoff. The scrum software development process for small teams. *IEEE software*, (4):26–32, 2000.

[19] Staša Milojević. Principles of scientific research team formation and evolution. *Proc. Natl. Acad. Sci. U. S. A.*, 111(11):3984–3989, 2014.

[20] Yves-Alexandre de Montjoye, Arkadiusz Stopczynski, Erez Shmueli, Alex Pentland, and Sune Lehmann. The strength of the strongest ties in collaborative problem solving. *Sci. Rep.*, 4, 06 2014.

[21] Aaron Clauset, Maxwell Young, and Kristian Skrede Gleditsch. On the frequency of severe terrorist events. *J. Confl. Resolut.*, 51(1):58–87, 2007.

[22] Juan Camilo Bohorquez, Sean Gourley, Alexander R Dixon, Michael Spagat, and Neil F Johnson. Common ecology quantifies human insurgency. *Nature*, 462(7275):911–914, 2009.

[23] Joseph Henrich. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses: the tasmanian case. *Am. Antiqu.*, 69(2):197–214, 2004.

[24] Maxime Derex, Marie-Pauline Beugin, Bernard Godelle, and Michel Raymond. Experimental evidence for the influence of group size on cultural complexity. *Nature*, 503(7476):389–391, 2013.

[25] Claes Andersson and Dwight Read. Group size and cultural complexity. *Nature*, 511(7507):E1–E1, 07 2014.

[26] Maxime Derex, Marie-Pauline Beugin, Bernard Godelle, and Michel Raymond. Derex et al. reply. *Nature*, 511(7507):E2–E2, 07 2014.

[27] Derek de Solla Price. *Little science, big science*. New York: Columbia University Press, 1963.

[28] Miranda Lee Pao. Global and local collaborators: a study of scientific collaboration. *Inf. Process. Manage.*, 28(1):99–109, 1992.

[29] John Hudson. Trends in multi-authored papers in economics. *J. Econ. Perspect.*, 10(3):153–158, 1996.

[30] Bibb Latane, Kipling Williams, and Stephen Harkins. Many hands make light the work: The causes and consequences of social loafing. *Journal of personality and social psychology*, 37(6):822, 1979.

[31] Stephen G Harkins. Social loafing and social facilitation. *Journal of Experimental Social Psychology*, 23(1):1–18, 1987.

[32] Steven J Karau and Kipling D Williams. Social loafing: A meta-analytic review and theoretical integration. *Journal of personality and social psychology*, 65(4):681, 1993.

[33] Jean-Louis Denis, Lise Lamothe, and Ann Langley. The dynamics of collective leadership and strategic change in pluralistic organizations. *Acad. Manage. J.*, 44(4):809–837, 2001.

[34] Rufus A Johnstone and Andrea Manica. Evolution of personality differences in leadership. *Proc. Natl. Acad. Sci. U. S. A.*, 108(20):8373–8378, 2011.

[35] Noshir S Contractor, Leslie A DeChurch, Jay Carson, Dorothy R Carter, and Brian Keegan. The topology of collective leadership. *Leadersh. Q.*, 23(6):994–1011, 2012.

[36] Jon R Katzenbach. *The wisdom of teams: Creating the high-performance organization*. Harvard Business Press, 1993.

[37] Frédéric Delmar and Scott Shane. Does experience matter? The effect of founding team experience on the survival and sales of newly founded ventures. *Strateg. Organ.*, 4(3):215–247, 2006.

[38] Assar Lindbeck and Dennis J Snower. Multitask learning and the reorganization of work: From Tayloristic to holistic organization. *J. Lab. Econ.*, 18(3):353–376, 2000.

[39] Steven Postrel. Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams. *Organ. Sci.*, 13(3):303–320, 2002.

[40] Tanzeem Choudhury and Alex Pentland. Sensing and modeling human networks using the sociometer. *2012 16th International Symposium on Wearable Computers*, 0:216, 2003.

[41] Dashun Wang, Chaoming Song, and Albert-László Barabási. Quantifying long-term scientific impact. *Science*, 342(6154):127–132, 2013.

[42] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[43] Herbert A Simon. On a class of skew distribution functions. *Biometrika*, 42(3/4):425–440, 1955.

[44] Sujin K Horwitz and Irwin B Horwitz. The effects of team diversity on team outcomes: A meta-analytic review of team demography. *J. Manag.*, 33(6):987–1015, 2007.

[45] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 2009.