

Understanding the group dynamics and success of teams

Michael Klug¹ and James P. Bagrow^{1,2,3,*}

¹Department of Mathematics & Statistics, The University of Vermont

²Vermont Complex Systems Center, The University of Vermont

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

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

January 5, 2016

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 focused across the team, with only a few members performing most work. We find that highly successful teams are significantly more focused than average teams of the same size, that their members have worked on more diverse sets of projects, and the members of highly successful teams are more likely to be core members or ‘leads’ of other teams. The relations between team success and size, focus and especially team experience cannot be explained by confounding factors such as team age or external contributions from non-team members nor by group mechanisms such as social loafing. Taken together, these features point to organizational principles that may 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]. New tools, including electronic sensor systems, can quantify team activity and performance [5, 4]. With the rise in prominence of network science [6, 7], much effort has gone into discovering meaningful groups within social networks [8, 9, 10, 11, 12, 13, 14, 15] and quantifying their evolution [16, 15]. Teams are increasingly important in research and industrial efforts [17, 18, 19, 3, 4, 20, 21] and small, coordinated groups are a significant component of modern human conflict [22, 23]. There are many important dimensions along which teams should be studied, including their size, how work is distributed among their members, and the differences and similarities in the experiences and backgrounds of those team members. Recently, there has been much debate on the “group size hypothesis”, that larger groups are more robust or perform better than smaller ones [24, 25, 26, 27]. Scholars of science have noted for decades that collaborative research teams have been growing in size and importance [28, 29, 30, 20]. 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 [31, 32, 33]. Meanwhile, case studies show that leadership [34, 3, 35, 36] and experience [37, 38] 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 [39, 40]. In all of these areas, large-scale, quantitative data can push the study of teams forward.

Teams are important for modern software engineering tasks, and researchers have long studied the digital traces of open source software projects to better quantify and understand how teams work on software projects [41]. Researchers have investigated estimators of work activity or effort based on edit volume, such as different ways to count the number of changes made to a software’s source code [42, 43, 44, 45]. Various dimensions of success of software projects, such as popularity, timeliness of bug fixes, or other quality measures have been studied [46, 47, 48]. Successful open source software projects show a layered structure of primary or core contributors surrounded by lesser, secondary contributors [49]. At the same time, much work is focused on case studies [50, 44] of small numbers of highly successful, large projects [41]. Considering these studies alone runs the risk of survivorship bias or other selection biases, so large-scale studies of large quantities of teams are important complements to these works.

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 now 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 launched² 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. 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 quantify one aspect of the **success** of the team, in a manner analogous to the use of citations of research literature as a proxy for “impact” [51]. Of course, 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 difficult to objectively quantify. Instead this is a measure of popularity as would be other statistics such as web traffic, number of downloads, and so forth [46].

In this study, we analyze the memberships and activities of approximately 150,000 teams, as they perform real-world tasks, to uncover the blend of features that relate to success. To the best of our knowledge this is the largest study of real-world team success to date. We present results that demonstrate (i) how teams distribute or focus 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 while accounting for confounds such as team size.

The rest of this paper is organized as follows: In Sec. 1 we describe our GitHub dataset; give definitions of a team, team success, and work activity/focus of a team member; and introduce metrics to measure various aspects of the experience and experiential diversity of a team’s members. In Sec. 2 we present our results relating these measures to team success. In Sec. 3 we present statistical tests on linear regression models of team features to control for potential confounds between team features and team success. Lastly, we conclude with a discussion in Sec. 4.

¹For examples, see <https://github.com/showcases/science>.

²See <https://github.com/blog/1840-improving-github-for-science>.

1 Methods and Materials

Dataset and team selection

Public GitHub data covering 1 January 2013 to 1 April 2014 was collected from githubarchive.org in April 2014. In their own words, “GitHub Archive is a project to record the public GitHub timeline, archive it, and make it easily accessible for further analysis”. These activity traces contain approximately 110M unique events, including when users create, join, or update projects. Projects on GitHub are called “repositories”. For this work we define a **team** as the set of users who can directly update (“push to”) a repository. These users constitute the **primary** team members as they have either created the project or been granted autonomy to work on the project. The number of team members was denoted by M . Activity or workload W was estimated from the number of pushes. A push is a bundle of code updates (known as commits), however most pushes contain only a single commit (see SI; see also ref. [45]). As with all studies measuring worker effort from lines-of-code metrics, this is an imperfect measure as the complexity of a unit of work does not generally map to the quantity of edits. 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 . This is a popularity measure of success, however the choice to bookmark a project does imply it offers some value to the user. To avoid abandoned projects, studied teams have at least one stargazer ($S > 0$) and at least two updates per month on average within the githubarchive data. These selection criteria leave $N = 151,542$ teams. We also collect T , the time of creation on GitHub, for each team project. This is useful for measuring confounds: for example, older teams may tend to both have more members and have more opportunities to increase success. Of the teams studied, 67.8% were formed within our data window. Otherwise we do not consider temporal features of teams such as team formation in this work. A small number of studied teams (1.08%) have more than ten primary members ($M > 10$); those teams were not shown in figures, but they were present in all statistical analyses.

Secondary Team GitHub provides a mechanism for external, non-team contributors to propose work that team members can then choose to use or not. These proposals are called pull requests. (Other mechanisms, such as discussions about issues, are also available to non-team contributors.) These secondary or external team contributors are not the focus of this work and have already been well studied by OSS researchers [41].

However, it is important to ensure that they do not act as confounding factors for our results, since more successful teams will tend to have more secondary contributions than other teams. So we measure for each team M_{ext} , the number of unique users who submit at least one pull request, and W_{ext} the number of pull requests. We will include these measures in our combined regression models. Despite their visibility in GitHub, pull requests are rare [52]; in our data, 57.7% of teams we study have $W_{\text{ext}} = 0$, and when present pull requests are greatly outnumbered by pushes on average: $\langle W/W_{\text{ext}} \mid W_{\text{ext}} > 0 \rangle = 42.3$ (median 16.0), averaged over all teams with at least one pull request.

Effective team size

The number of team members, M , does not fully represent the size of a team since the distribution of work may be 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 projects that user i works on (has pushed to). (Projects in R_i need at least twice-monthly updates on average, as before, but may have $S = 0$ so as to better capture i 's background, not just successful projects.) We estimate the **experience** E of a team of size M as

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

and the experiential **diversity** D as

$$D = \frac{|\bigcup_i R_i|}{\sum_i |R_i|},$$

where the sums and union run over the M members of the team. Note that $D \in [1/M, 1)$. Experience measures the quantity of projects the team works on while diversity measures how many or how few projects the team members have in common, the goal being to capture how often the team has worked together.

Lastly, someone is a **lead** when, for at least one project they work on, they contribute more work to that project than any other member. A non-lead member of team j may be the lead of project $k \neq j$. The number of leads L_k in team k of size M_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 projects j . Of course, the larger the team the more potential leads it may contain so when studying the effects of leads on team success we only compare teams of the same size (comparing L while holding M fixed). Otherwise, E and D already account for team size.

2 Results

We began our analysis by measuring team success S as a function of team size M , the number of primary contributors to the team’s project. Since S is at least partially popularity measure, we expect larger teams to also be more successful. Indeed, 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. 1). This strong trend holds for the median success as well (inset). While this observed trend was highly significant, the rank correlation 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 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,




Figure 1: Larger teams have significantly more success on average, with a 300% increase in S as M goes from 1 to 10. This correlation may be due to more team members driving project success or success may act as a mechanism to recruit team members. Errorbars here and throughout denote ± 1.96 s.e. (inset) Using the median instead of the mean shows that this trend is not due to outliers.

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 result is in line with prior studies [50], supporting the plausibility of our definition of a team and our measure of work.

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 success (Fig. 2B).

Can this result be due to a confounding effect from success? More successful projects will tend to have more external contributors, for example, which can change the distribution of work. For example, in one scenario a team member may be a “community manager” merging in large numbers of external

contributions from non-team members. To test this we examined only the 57.7% of teams that had no external contributions ($W_{\text{ext}} = 0$) and tested among only those teams whether dominated teams were more successful than non-dominated teams. Within this subset of teams dominated teams had significantly higher S than non-dominated teams (Mann-Whitney U test with continuity correction, $p < 10^{-8}$). The Mann-Whitney U test (MWU) is non-parametric, using ranks of (in this case) S to mitigate the effects of skewed data, and does not assume normality. We conclude from this that external contributions do not fully explain the relationship between workload focus and team success.

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 $m \approx 1.22$. Note that M and m are positively correlated ($\rho = 0.985$).

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 than less focused teams of the same size, regardless of total team size.

Further analyses revealed the importance of team composition and its role in team success.

Team members do not perform their work in a vacuum, they each bring experiences from their other work. Often members of a team will work on other projects. We investigated these facets of a team's composition by exploring (i) how many projects the team's members have worked on, (ii) how diverse are the other projects (do the team members have many or few other projects in common), and (iii) how many team members were "leads" of other projects.

An estimate of 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.

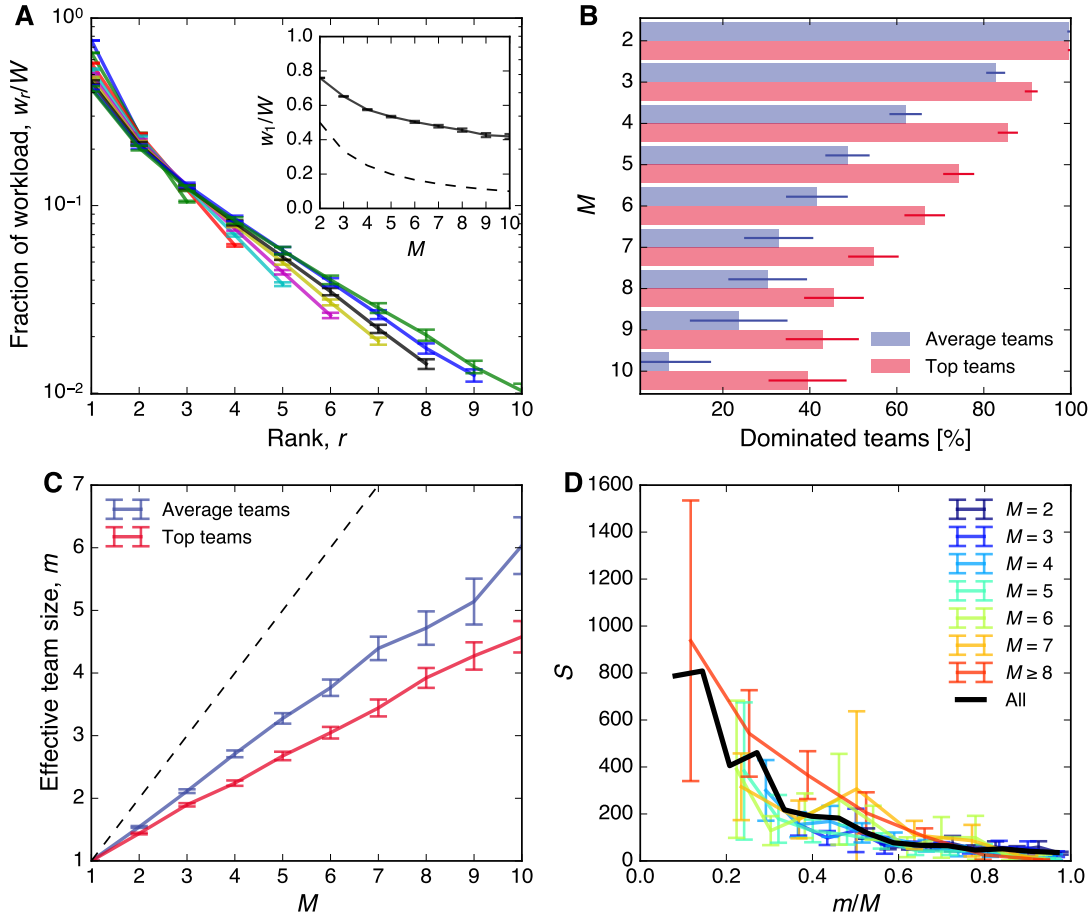


Figure 2: Teams are focused, and top teams are more focused than average teams of the same size. **(A)** The average fraction of work w_r/W performed by the r -th most active member, where W is the total work of the team, for different size teams. Larger teams perform more work overall, but the majority 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 $2 \leq M \leq 10$. These results are not due to outliers in S ; see SI.

It may be that the volume of experience does not contribute much to the success of a team, but this seems to contradict previous studies on the importance of experience and wisdom [37, 38]. 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 past, perhaps developing an experiential shorthand. Conversely, successful teams may instead have multiple distinct viewpoints, solving challenges with a multi-disciplinary perspective [53].

To estimate 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 was significantly correlated team success. See also our combined modeling efforts.

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 [34, 36]. 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 successful 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 project ($1 \leq L \leq M$, 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.

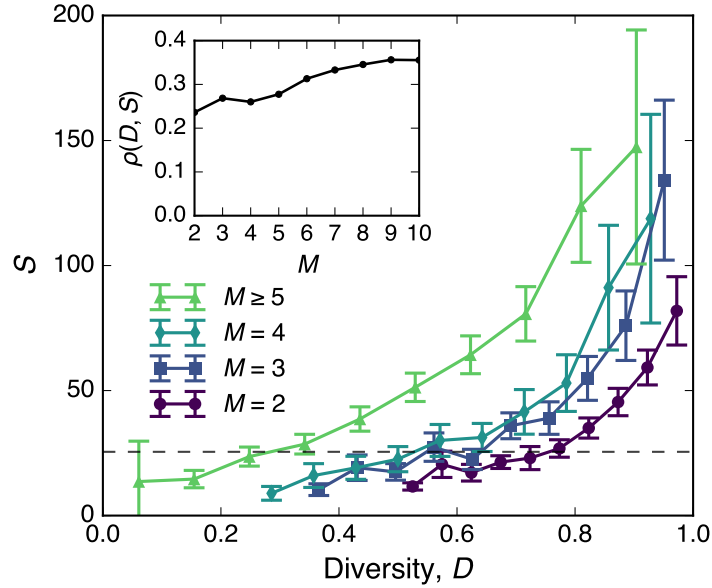


Figure 3: Teams whose members belong to more diverse sets of other teams tend to be more successful, regardless of team size. The dashed line denotes the average success of all teams. (Inset) The rank correlation ρ between diversity and success grows with team size. Teams above the 99th percentile in S were excluded to ensure the trend is not due to outliers.

Expanding on this observation, Table 1 illustrates the extreme case of teams of size M with a single lead ($L = 1$) compared with teams of the same size 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 team composition cannot be easily explained as a confound with success or secondary contributions as they study specific features and projects of the individuals who comprise a team, those features are not related to the successes of other projects an individual may work on, and they strictly control for total team size M (for example, we only compare teams with different values of L when they have the same value M). These results further amplify our findings on team focus, and augment important existing research [37, 3, 4, 36, 53]. 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, to see if one relationship can be explained by another, which we will do with multivariate statistical models.

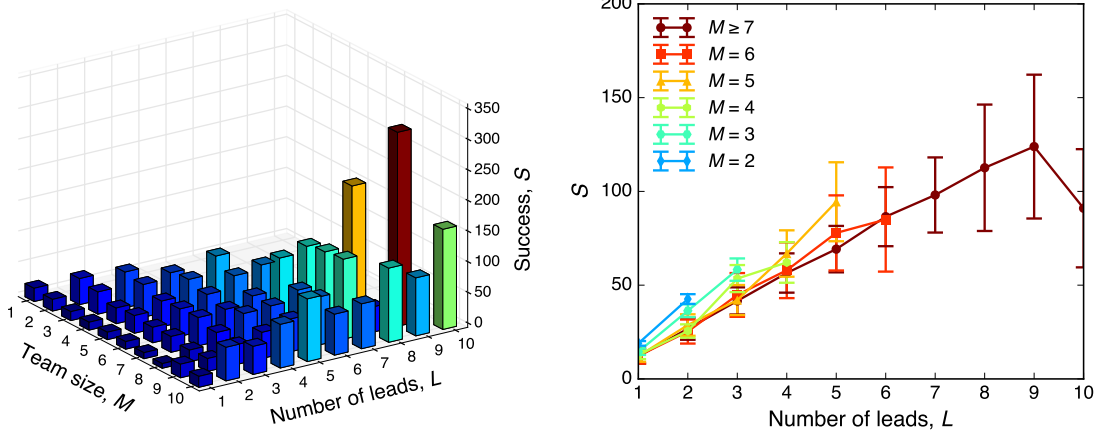


Figure 4: Teams with more leads have higher success than teams of the same size with fewer leads. A lead is someone who contributes more work to at least one team he or she belongs to than any other members of that team. Outliers in S were removed as before.

3 Combined models and confounds

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 fitted a linear regression model of success as a function of all explored measures. See 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. We also included the age of a project T (see Methods) as this may be a potential confounding factor as well (older projects have had more time to gain members and to gain success).

The number of leads L most strongly correlated with team success. Team age T , effective team size m , and team size M play the strongest roles after L in team success, and all three were also 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 diversity D of the team, followed by the total work W done on the project, were also significant measures related to

Table 1: Teams composed entirely of leads ($L = M$) are significantly more successful (MWU test on S) than teams of the same size with one lead ($L = 1$), regardless of team size M . Teams above the 99th percentile in S were excluded to ensure the differences were not due to outliers.

M	Num. teams N		Mean success S		MWU p -value
	$L = 1$	$L = M$	$L = 1$	$L = M$	
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

^aWhen $M \geq 8$, the number of teams with $L = M$ is too small ($N < 20$) for us to reasonably conclude the difference in S is significant, despite the small p -values.

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, any effects of E are explained by the other quantities.

What about secondary contributions, those activities made by individuals outside the primary team? We already performed one test showing that dominated teams are more successful than non-dominated teams even when there are no secondary contributions. Continuing along these lines, we augmented this linear model with two more dependent variables, M_{ext} and W_{ext} . Regressing on this expanded model (see SI for details) did not change the significance of any coefficients at the $p = 0.05$ level; E remained insignificant ($p > 0.1$). Both new variables were significant ($p < 0.05$). Note that there were no multicollinearity effects in either regression model (condition numbers ≤ 10). We conclude that secondary contributions cannot alone explain the observations relating team focus, experience, and lead number to team success.

4 Discussion

There has been considerable debate concerning the benefits of specialization compared to diversity in the workplace and other sectors [39]. 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,

Table 2: OLS regression model on team success, $S = \alpha + \beta_M M + \beta_m m + \beta_W W + \beta_E E + \beta_D D + \beta_L L + \beta_T T$. Outliers (above the 99th percentile in S) were filtered out to ensure they do not skew the model.

Variable x	Coefficient β_x^a	p
constant α	$1.351 \times 10^{-14} \pm 0.004951$	1
Team size, M	0.0848 ± 0.013963	$< 10^{-31}$
Eff. team size, m	-0.0989 ± 0.012140	$< 10^{-56}$
Total work, W	0.0323 ± 0.004997	$< 10^{-35}$
Experience, E	0.0004068 ± 0.004985	0.8729
Diversity, D	0.04099 ± 0.006357	$< 10^{-35}$
Num. of leads, L	0.1388 ± 0.006921	0
Age, T	0.1273 ± 0.005014	0

^aVariables 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.

offering ancillary support for teams led by another member. This has implications when organizations are designing teams and wish to maximize their success, at least as success was measured in these data. 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).

Of course, 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 imply 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 [14], a task that may be project- or domain-specific but which is worth further exploration.

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 [39] and supports the “surgical team” arguments of Brooks [17]. 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. We have provided statistical tests demonstrating that the relationship between focus and success cannot be due to secondary/external team contributions alone.

Work focus could possibly be explained by **social loafing** where individual members of a group con-

tribute less effort as part of the group than they would alone, yet loafing does not explain the correlation between, e.g., leads and success (Fig. 4). Likewise, our team composition results on group experience, experiential diversity, and the number of leads cannot be easily explained as a confound with success or secondary contributions: they study specific features of the individuals who comprise a team, those features are not related to the successes of other projects an individual may work on, and they strictly control for total team size M (except for the number of leads L , so for that measure we only compared teams with the same M). The measures we used for external team contributions, M_{ext} and W_{ext} , may be considered measures of success themselves, and studying or even predicting their levels from team features may prove a fruitful avenue of future work.

Lastly, there are two remaining caveats worth mentioning. We do not specifically control for automatically mirrored repositories (where a computer script copies updates to GitHub). Accurately detecting such projects at scale is a challenge beyond the scope of this work. However, we expect most will either be filtered out by our existing selection criteria or else they will likely only have a single (automated) user that only does the copying. The second concern is work done outside of GitHub or, more generally, mismatched assignments between usernames and their work. This is also challenging to fully address (one issue is that the underlying git repository system does not authenticate users). We acknowledge this concern for our workload focus results, but even it cannot explain the significant trends we observed on team composition such as the density of leads. Noise due to improperly recorded or “out-of-band” work has in principle affected all quantitative studies of online software repositories.

Acknowledgment

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] Tanzeem Choudhury and Alex Pentland. Sensing and modeling human networks using the sociometer. *2012 16th International Symposium on Wearable Computers*, 0:216, 2003.
- [6] Albert-László Barabási. *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume, 2003.
- [7] M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- [8] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*. Cambridge University Press, 1994.
- [9] 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.
- [10] 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.
- [11] Duncan J Watts, Peter Sheridan Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296(5571):1302–1305, 2002.
- [12] M. E. J. Newman and Juyong Park. Why social networks are different from other types of networks. *Phys. Rev. E*, 68(3):036122, 2003.
- [13] 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.
- [14] Yong-Yeol Ahn, James P Bagrow, and Sune Lehmann. Link communities reveal multiscale complexity in networks. *Nature*, 466(7307):761–764, 2010.
- [15] 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.
- [16] Gergely Palla, Albert-László Barabási, and Tamás Vicsek. Quantifying social group evolution. *Nature*, 446(7136):664–667, 2007.
- [17] Frederick P Brooks Jr. *The Mythical Man-Month*. Pearson Education, 1995.
- [18] Samer Faraj and Lee Sproull. Coordinating expertise in software development teams. *Management science*, 46(12):1554–1568, 2000.

- [19] Linda Rising and Norman S Janoff. The scrum software development process for small teams. *IEEE software*, (4):26–32, 2000.
- [20] Staša Milojević. Principles of scientific research team formation and evolution. *Proc. Natl. Acad. Sci. U. S. A.*, 111(11):3984–3989, 2014.
- [21] 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.
- [22] Aaron Clauset, Maxwell Young, and Kristian Skrede Gleditsch. On the frequency of severe terrorist events. *J. Confl. Resolut.*, 51(1):58–87, 2007.
- [23] 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.
- [24] Joseph Henrich. Demography and cultural evolution: how adaptive cultural processes can produce maladaptive losses: the tasmanian case. *Am. Antiqu.*, 69(2):197–214, 2004.
- [25] 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.
- [26] Claes Andersson and Dwight Read. Group size and cultural complexity. *Nature*, 511(7507):E1–E1, 07 2014.
- [27] Maxime Derex, Marie-Pauline Beugin, Bernard Godelle, and Michel Raymond. Derex et al. reply. *Nature*, 511(7507):E2–E2, 07 2014.
- [28] Derek de Solla Price. *Little science, big science*. New York: Columbia University Press, 1963.
- [29] Miranda Lee Pao. Global and local collaborators: a study of scientific collaboration. *Inf. Process. Manage.*, 28(1):99–109, 1992.
- [30] John Hudson. Trends in multi-authored papers in economics. *J. Econ. Perspect.*, 10(3):153–158, 1996.
- [31] 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.
- [32] Stephen G Harkins. Social loafing and social facilitation. *Journal of Experimental Social Psychology*, 23(1):1–18, 1987.
- [33] 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.
- [34] 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.
- [35] Rufus A Johnstone and Andrea Manica. Evolution of personality differences in leadership. *Proc. Natl. Acad. Sci. U. S. A.*, 108(20):8373–8378, 2011.
- [36] 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.

- [37] Jon R Katzenbach. *The wisdom of teams: Creating the high-performance organization*. Harvard Business Press, 1993.
- [38] 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.
- [39] 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.
- [40] Steven Postrel. Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams. *Organ. Sci.*, 13(3):303–320, 2002.
- [41] Kevin Crowston, Kangning Wei, James Howison, and Andrea Wiggins. Free/libre open-source software development: What we know and what we do not know. *ACM Computing Surveys (CSUR)*, 44(2):7, 2012.
- [42] Allan J Albrecht and John E Gaffney Jr. Software function, source lines of code, and development effort prediction: a software science validation. *Software Engineering, IEEE Transactions on*, (6):639–648, 1983.
- [43] Jarrett Rosenberg. Some misconceptions about lines of code. In *Software Metrics Symposium, 1997. Proceedings., Fourth International*, pages 137–142. IEEE, 1997.
- [44] Stefan Koch and Georg Schneider. Effort, co-operation and co-ordination in an open source software project: Gnome. *Information Systems Journal*, 12(1):27–42, 2002.
- [45] Abdulkareem Alali, Huzefa Kagdi, Jonathan Maletic, et al. What’s a typical commit? a characterization of open source software repositories. In *Program Comprehension, 2008. ICPC 2008. The 16th IEEE International Conference on*, pages 182–191. IEEE, 2008.
- [46] Kevin Crowston, James Howison, and Hala Annabi. Information systems success in free and open source software development: Theory and measures. *Software Process: Improvement and Practice*, 11(2):123–148, 2006.
- [47] Chandrasekar Subramaniam, Ravi Sen, and Matthew L Nelson. Determinants of open source software project success: A longitudinal study. *Decision Support Systems*, 46(2):576–585, 2009.
- [48] Amir Hossein Ghapanchi, Aybuke Aurum, and Graham Low. A taxonomy for measuring the success of open source software projects. *First Monday*, 16(8), 2011.
- [49] Kevin Crowston and James Howison. The social structure of free and open source software development. *First Monday*, 10(2), 2005.
- [50] Audris Mockus, Roy T Fielding, and James D Herbsleb. Two case studies of open source software development: Apache and mozilla. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(3):309–346, 2002.
- [51] Dashun Wang, Chaoming Song, and Albert-László Barabási. Quantifying long-term scientific impact. *Science*, 342(6154):127–132, 2013.

- [52] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M German, and Daniela Damian. The promises and perils of mining github. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, pages 92–101. ACM, 2014.
- [53] 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.