The Benefits, Barriers, and Risks of Big Team Science

Authors: Patrick S. Forscher[1,2], Eric-Jan Wagenmakers[3], Nicholas A. Coles[4], Miguel Alejandro Silan[5,6], Natália Dutra[7], Dana Basnight-Brown[8], & Hans IJzerman[2,9]

Affiliations:

[1]Busara Center for Behavioral Economics

[2]Université Grenoble Alpes

[3]University of Amsterdam

[4]Stanford University

[5]Université Lumière Lyon 2

[6]Annecy Behavioral Science Lab

[7]Universidade Federal do Pará

[8]United States International University-Africa

[9]Institut Universitaire de France

Abstract: Progress in psychology has been frustrated by challenges concerning replicability, generalizability, strategy selection, inferential reproducibility, and computational reproducibility. Although often discussed separately, these five challenges may share a common cause: insufficient investment of intellectual and non-intellectual resources into the typical psychology study. We suggest that the emerging emphasis on *big team science* can help address these challenges by allowing researchers to pool their resources together to increase the amount available for a single study. However, the current incentives, infrastructure, and institutions in academic science have all developed under the assumption that science is conducted by solo Principal Investigators and their dependent trainees, an assumption that creates barriers to sustainable big team science. We also anticipate that big team science carries unique risks, such as the potential for big team science organizations to be co-opted by unaccountable leaders, become overly conservative, and make mistakes at a grand scale. Big team science organizations must also acquire personnel who are properly compensated and have clear roles, raising risks related to mismanagement and a lack of financial sustainability. If researchers can manage its unique barriers and risks, big team science has the potential to spur great progress in psychology and beyond.

Despite decades of investment, progress in psychology has been slower than many scholars would like (Fanelli, 2010; Meehl, 1978; Newell, 1973). The lack of progress is empirically supported by the results of replication studies. Whether with the same population, setting, and materials (the *replicability challenge*; Klein et al., 2014; Open Science Collaboration, 2015) or after a change to one or more of these features (the *generalizability challenge*; Henrich et al., 2010; Tiokhin et al., 2019; Yarkoni, 2019), replicated results often differ meaningfully from original results. Meaningful differences also occur in other forms of replication, such as: when separate teams develop research strategies to address the same research question (the *strategy selection challenge*; Landy et al., 2020) when separate teams develop analysis plans for the same dataset (the *inferential reproducibility challenge*; Botvinik-Nezer et al., 2019; Silberzahn et al., 2018), and even when separate teams write code to execute the same analysis (the *computational reproducibility challenge*; Donoho et al., 2008; Hardwicke et al., 2018; Obels et al., 2019).

These five challenges have complex proximal causes. Yet they may share a common ultimate cause: insufficient resource investment in the typical psychology study. Insofar as this premise is true, a particular method of collaboration, *big team science*, may help address this ultimate cause by efficiently scaling up the resources that can be invested in any given study. We define *big team science* as a method involving a relatively large number of collaborators who may be dispersed across labs, institutions, disciplines, cultures, and continents. We contrast this system of science with *small team science*, which is usually organized around a single Principal Investigator and their dependent trainees. If the unique risks and challenges of big team science are properly understood and managed, this method of collaboration may have great potential to improve the efficiency and information value of psychological science.

The challenges in psychology share a common cause

Discussion of the slow progress in psychology has a long history. Scholars have argued persuasively that each of the five challenges facing psychology has a pernicious and potentially destructive influence on scientific progress (Donoho et al., 2008; Landy et al., 2020; Meehl, 1978; Sears, 1986; Silberzahn et al., 2018). These scholars have also presented a dizzying array of remedies for these challenges, ranging from increasing study sample sizes (Button et al., 2013), which ought to improve replicability, to implementing version control (Vuorre & Curley, 2018), which ought to improve computational reproducibility. Although varied, these remedies share a common feature: they ask researchers to incur additional costs to improve a particular aspect of study rigor (Finkel et al., 2017; LeBel et al., 2017; see Table 1 for a list of proposed remedies and their attendant costs).

Because the individual researcher incurs costs for each remedy, discussions of problems in psychology tend to assume that implementations of these remedies are zero-sum in the following sense. The scientific resources that are available to a given researcher, such as time, labor, and money, are limited. This means that, as long as a solution to improve, say, replicability does not also improve all the other aspects of study rigor, resources devoted to implementing that solution to replicability must necessarily take away resources that could be spent on solutions to other aspects of rigor. Replicability, generalizability, inferential reproducibility, strategy selection, and computational reproducibility thus become qualities that must be traded off each

other because devoting resources to improve one of these five qualities means sacrificing resources that could be invested in improving the others.

Challenge	Suggested remedy	Costs of implementation	Who bears costs?
Replicability	Increase sample size	Labor and money for extra recruitment	Researchers; participants
	Preregister analysis plans	Labor creating preregistration and making it accessible on a platform	Researchers; platform maintainers
	Improve documentation of materials	Labor creating and archiving documentation materials	Researchers; platform maintainers
Generalizability	Use special platforms to recruit culturally different participants	Labor developing, maintaining, and learning to use new platforms; money using and maintaining the platforms	Researchers; platform maintainers
	Collaborate with colleagues from different settings & backgrounds	Labor developing relationships with new colleagues	Researchers
	Improve generality of materials	Labor and money developing and validating improved stimuli and measures	Researchers; participants
Strategy selection	Examine many construct operationalizations	Labor and money examining different operationalizations	Researchers; participants
	Consensus design	Labor developing relationships with relevant experts; labor implementing consensus process	Researchers
Inferential reproducibility	Machine-readable hypothesis tests	Labor learning and implementing machine-readable systems; money developing and maintaining platforms	Researchers; platform maintainers
	Many analyst design	Labor making connections with relevant analysts	Researchers
	Examine a multiverse of tests	Labor developing relationships with relevant experts; labor implementing multiverse analysis	Researchers
Computational reproducibility	Share code and data	Labor to put code and data into a shareable form; money developing and maintaining sharing platforms	Researchers; platform maintainers
	Implement code checking	Labor and money implementing code checking	Researches; journals
	Improve documentation of codebases	Labor making code readable and hygienic	Researchers
	Use capsules, version control, and markdown	Labor learning and implementing new workflows; money maintaining capsule and version control platforms	Researchers; platform maintainers

Table 1. Five challenges in psychology and some proposed remedies for these challenges. The challenges all ask researchers to bear extra costs to improve some aspect of study rigor.

We suggest that the views that prioritize one aspect of rigor over the others all share an important, mistaken assumption: the pool of resources available for investment in a given study is fixed. If the resources devoted to a single study are not fixed, these resources may be raised to a level that is high enough that no tradeoffs are necessary among the various types of research rigor. Instead, a researcher need merely select as many of the rigor-enhancing strategies from Table 1 as they like and invest the money, human resources, and skills required to deploy those strategies.

Thus, although they are often discussed separately, all five of psychology's major problems may share a common cause: under-investment of resources, whether those resources are money, person-hours, or specialized expertise, in psychology studies (see also Cuccolo et al., 2021; Uhlmann et al., 2019). Instead of asking *which* aspects of study rigor ought to be prioritized, reframing psychology's problems from this perspective allows us to ask a different question: how can we increase the resources invested in psychology studies and ascertain that the extra investment is used efficiently?

Increasing resource investment through big team science organizations

Increasing the per-study resource investment across an entire research ecosystem faces a collective action problem: as long as scientists are rewarded for publishing more studies, any single scientist who decides to invest more resources in fewer studies will be outcompeted by scientists who invest less resources in more studies (Smaldino & McElreath, 2016). Any attempt to increase the per-study resource investment of the entire ecosystem must adopt one of two solutions: (1) directly change the institutional incentives that prioritize quantity of publication and/or (2) devise new institutions, which we call *big team science organizations*, that allow blocs of scientists to increase the resource investments in concert.

Direct change in the institutions of science is difficult because science itself has a decentralized structure – changing its reward structures requires buy-in from many independent actors. This limits the effectiveness of initiatives to directly change the incentives that prioritize quantity of publication. A similar problem afflicts the establishment of new organizations that coordinate the efforts of scientists to jointly invest in single projects. Fortunately, new information and communication technologies, such as the Google suite of collaboration tools, the Open Science Framework, Slack, and Zoom, have made the establishment of such organizations more feasible because they permit rapid, low-cost communication across far-flung countries and circumstances (Teasley & Wolinsky, 2001). Such communication lays the groundwork for new institutions that can change how scientific actors spend resources on their science (Spellman, 2015).

The function of an organization is to coordinate the activities of many actors simultaneously. An organization solves the collective action problem faced by individual actors by allowing multiple actors to pre-commit to a coordinated, simultaneous course of action. In the context of science, a *big team science organization* allows many scientists to pre-commit to investing their limited scientific resources into one big project in exchange for an individual reward – usually, authorship on a publication, though other rewards are possible, such as money or networking opportunities. During the project itself, the organization coordinates the actions of

the scientists so that they do not work at cross-purposes. These functions allow the per-study investments from big team science projects to eclipse what is achievable in small team science.

We believe that the primary and most definitional function of a big team science organization is to allow larger investments of material and intellectual resources in a given project. However, once established, big team science organizations have the potential to provide benefits to both individual scientists and the broader scientific ecosystem that are *external* to this primary function. As an example of one such *positive externality*, big team science organizations centralize many aspects of project administration in one organization rather than forcing individual labs to complete these functions on their own. This not only produces an efficiency for the scientific ecosystem, but also allows researchers at organizations without strong research infrastructure, such as instructors at teaching-focused universities (Wagge, Baciu, et al., 2019; Wagge, Brandt, et al., 2019) to do more research than would otherwise be possible. Big teams also allow individual researchers to specialize into roles that match their skills rather than occupying all roles simultaneously, producing an efficiency for the scientific ecosystem and giving the individual researcher the opportunity to develop a specialized skill and thereby command a higher salary. Researchers who join a big team science organization also gain access to an expansive community. For the individual researcher, this reduces intellectual isolation and opens the door to professional opportunities; for the scientific ecosystem, this can generate new collaborations that would not have been possible absent the expansive community. Finally, if big team science organizations become robust enough, they can become political actors of their own right within the broader scientific ecosystem. Thus, they can nudge the incentives and values of the entire scientific ecosystem by, for example, prioritizing the qualities of scientists to improve team productivity rather than individual productivity (Tiokhin et al., 2021).

Thus, big team science organizations can invest more resources into individual projects than is possible in small teams. Big team science organizations also have the potential to provide a variety of benefits to individual researchers and the broader scientific ecosystem that are external to this primary function. Whether these benefits materialize depend on establishing the organizations in the first place and ensuring the organizations do not fall prey to some of the unique risks and challenges of this style of science.

Big team science outside of psychology

Some disciplines have already adopted a big team science model, so they can provide useful examples of how big team science organizations can arise. These examples can also illustrate whether big team science can address the limitation of under-resourced studies. In the 1990s, behavioral genetics generally featured small studies examining the relationship between variation in a single candidate gene and a complex behavior or trait. In one prominent example, only 52 patients provided genetic material for an analysis of the relationship between the 5-HTTLPR polymorphism and major depression (Heils et al., 1996), a finding that spurred enormous interest in the biological mechanisms through which these genes might cause depression (Dannlowski et al., 2008; Gotlib et al., 2008). Unfortunately, and similar to the current situation in psychology, these early results were contradicted by failed replication studies (Gillespie et al., 2005). As more and more replication studies contradicted the earlier optimistic ones, researchers in behavioral genetics realized that changes were in order (Rieckmann, 2009).

The two decades that followed were a turning point for behavioral genetics because they gave rise to big team science in the form of large research consortia. Without pooling resources, researchers would have been unable to attain the scale of data necessary to advance our understanding of behavioral genetics. Before these consortia could attain such large-scale data, they needed to invest considerable research, time, and funds devising methodological workflows (Corvin et al., 2010), designing infrastructure to manage and harmonize datasets (Sullivan, 2010), and developing processes that increased the accuracy of the measurement of genotypes and phenotypes (23andMe Research Team et al., 2018; Corvin et al., 2010). With these solutions in hand, the research consortia served as a conduit through which multiple labs pooled resources to achieve large, heterogeneous samples that would otherwise have been out of reach for any single lab working independently. The UK Biobank is a prominent example, hosting a repository for data from over 500,000 participants, a sample that dwarfs the sample of 52 used for the first 5-HTTLPR study. The resulting findings have revolutionized our understanding of genes and behavior: we now know that, instead of single genetic variants exerting large influences, a large number of variants have small influences (23andMe Research Team et al., 2018). These impacts were subtle enough that they could not be observed and studied using the small, homogeneous samples typically employed in the early 1990s. Innovations in methods, workflows, and measurement allowed for a sea change in our understanding of the relationship between genes and behavior.

Particle physics provides another, still larger-scale, example. The 1950s saw a growing recognition among particle physicists that further advances would require a scale of research and equipment that would strain the budgets of entire nations, let alone single laboratories (Krige, 2004). For their part, Western governments saw cooperative large-scale investments in physics as a potential diplomatic tool; the projects served as a highly visible counterweight growing Soviet power, demonstrating the benefits that could accrue from an internationalist, democratic, and cooperative world order (Krige, 2004). These twin recognitions spurred the establishment of "mega-collaborations": large, international collaborations spanning multiple countries. Such mega-collaborations could only realize their aims through both technical and social achievements to establish the infrastructure, documentation, and workflows necessary to coordinate the efforts of huge numbers of scientists (Bakker, 1955; Brumfiel, 2011; for a modern workflow, see Espinosa et al., 2020). One of the pre-eminent examples of these mega-collaborations is the European Organization for Nuclear Research (CERN), which, together with Brookhaven National Laboratory, helped usher in a new era of particle physics: high-energy physics, in which large teams of scientists harness vast resources to probe the most fundamental constituents of matter (Bryant, 1994).

Name	Type	Description
Reproducibility Project: Psychology (Open Science Collaboration, 2015)	Ad-hoc	A large team-science project in psychology, which involved 270 contributing authors, who replicated 100 effects.
ManyLabs (Ebersole et al., 2016; Klein et al., 2014, 2018, 2019)	Ad-hoc	These projects involve dozens of researchers, each of whom collects data using their own local resources. The individual datasets are pooled to create a large common dataset.
Many Smiles Collaboration (Coles et al., 2019)	Ad-hoc	A collaborative project designed to find the best way to test the hypothesis that facial expressions influence emotions. The test was developed through consensus design. The resulting design uses multiple operationalizations and will be tested in a multi-site study.
Research contests to reduce implicit race preferences (Lai et al., 2014, 2016)	Ad-hoc	A series of contests to develop the most effective interventions to reduce implicit race preferences. Separate teams submitted interventions; all interventions were run together on a large online platform.
Registered Replication Reports (Simons et al., 2014)	Standing	An initiative to conduct multi-lab, preregistered close replications of previous studies. The initiative supervised by a hosting journal. Originally initiated at _Perspectives at Psychological Science_, other journals have adopted similar initiatives. Exact policies differ across journals.
Collaborative Replication and Education Project (Wagge, Baciu, et al., 2019; Wagge, Brandt, et al., 2019)	Standing	A framework for engaging undergraduates in replication research. On consultation with student advisors, the CREP team selects target effects for replication. CREP develops a set of templates that guide the replication process and uses a team of reviewers to ensure that the methods for each individual lab are true to the template.
Psychological Science Accelerator (Jones et al., 2021; Moshontz et al., 2018)	Standing	A standing, democratically structured network of labs focused on improving the national diversity of psychology samples.
ManyBabies (Byers-Heinlein et al., 2020)	Standing	A standing network that conducts multi-site infancy research.

Table 2. Eight big team science initiatives in psychology.

Big team science in psychology

Psychology has also started to witness an increase in big team science projects. Although most of these projects are recent, they have already had an outsized impact on the field (see Table 2 for details). However, these efforts have revealed three categories of obstacles that must be overcome if big team science is to maintain a sustained presence in the research landscape. These obstacles are: _incentivizing labor_ within the collaboration; _developing and maintaining infrastructure_ to coordinate team science activities; and _dealing with institutions_ designed around research conducted by smaller teams.

Barrier	Description	Solutions
Incentives	Academic prestige goes disproportionately to the first-listed author on publications	Consortium authorship Contributorship systems, like CRediT
Infrastructure	Coordinating team science requires dedicated infrastructure	Google Suite of tools for general online collaboration Open Science Framework for sharing materials & data Zoom and other videoconferencing for online meetings *formr* for online deployment of big team science projects *experimentum* for building experiments and managing projects *ScienceVerse* for documenting big team science projects
Institutions	Big team science must navigate frictions created by institutional policies developed around solo PIs	None at present

Table 3. Barriers to big team science and some potential solutions.

Incentives. Due to the central importance of prestige in obtaining academic rewards, academics are incentivized to obtain publications that can be used to enhance their prestige. This means that any collective research enterprise that wishes to direct the labor of academic scientists must either rely on sporadic volunteerism, find some other reward, like money, that substitute for prestige, or find a way to dole out prestige. We describe the risks of relying on volunteerism and the barriers to providing money as a reward in the sections on institutional barriers and risks of financial unsustainability. Here we deal more completely with prestige as a form of compensation.

A central problem with prestige as an incentive for participation in big team science is that, under the current system for awarding credit for publications, the bulk of credit goes to the first-listed author. This reward structure does not effectively incentivize the labor of the numerous other people who are necessary to produce big team science publications. Moreover, big team science organizations must do a tremendous amount of administration and coordination that, while critical to the success of the collaboration as a whole, is not easily creditable on publications.

We see at least two innovations that may help resolve this problem. The first is *consortium authorship*, in which a publication is credited to a collective entity rather than a group of individuals, or in which all members of the collaboration are listed alphabetically on all publications to render the individual subservient to the collective (Birnholtz, 2008). This was the approach to credit taken by the Open Science Collaboration when they published the

Reproducibility Project: Psychology (Open Science Collaboration, 2015). This is also the approach taken by some large physics collaborations, especially for papers that document details of infrastructure or instrumentation (e.g., ALICE collaboration, 2010; Harry & the LIGO Scientific Collaboration, 2010; The ATLAS collaboration et al., 2018). This innovation incentivizes collaboration by attempting to flatten the credit reward structure. Consortium authorship systems come at some risk, however: they can under-reward highly active but less visible consortium members, or lead consortium members to seek informal routes of gaining recognition, thereby effectively substituting the flat, formal system of credit with a more arbitrary informal system (Birnholtz, 2008).

The second innovation is to disclose contributions in a more fine-grained way through contributorship systems, such as CRediT (McNutt et al., 2018). These systems allow contributors to list within a manuscript the specific roles contributors play in projects. In principle, contributorship systems should allow people to develop reputations for filling certain types of project roles, such as data analysis or project management, and they should allow more fine-grained accountability when errors are detected in the project. Finally, contributorship systems enable contributors to provide evidence of excellence in a particular project role when applying for grants, jobs, or other professional rewards.

However, contributorship systems will only serve their purpose of giving recognition to excellent team scientists if people who control professional rewards, such as members of hiring committees, promotion committees, and funding committees, actually attend to and reward evidence of excellent team science contributions. Although we see promise in contributorship systems as a method to incentivize big team science labor (Holcombe, 2019), the existence of contributorship information is by itself not sufficient to ensure that incentives are aligned to reward big team science labor. The people who control professional rewards must also attend to these systems if they are to serve their intended functions.

Infrastructure. A second obstacle is the need for infrastructure to help facilitate and coordinate big team science projects. Some of this infrastructure already exists, such as the Slack, Google suite of collaboration tools, and Zoom. This existing infrastructure has been instrumental in propelling big team science to where it is today. For example, projects run through the Psychological Science Accelerator (Moshontz et al., 2018) use a combination of Slack and email for project coordination, Google Docs, Sheets, and Forms for collaborative project workflows (for a writing worklow, see Moshontz et al., 2019), a shared Google Drive for collaborative files management, and Zoom for conference calls. However, this existing infrastructure is general-purpose and therefore does not fully meet the specialized needs of behavioral science. For example, most psychology data collection platforms are designed for use by one or two users. This is insufficient in a team of researchers numbering in the hundreds. Although not every researcher needs to be part of the development of, for example, a survey, multiple users need to access survey instruments when helping with translation or other parts of the survey development process.

Another example of where research infrastructure is insufficient is when recruiting samples of participants that vary linguistically – a common occurrence when working with worldwide collaborators. Often, translated versions of the target measures are simply unavailable. Creating and validating a greater array of translated measure versions will go a long

way toward facilitating big team science projects. A lack of applicability across languages and countries apply to infrastructure across all stages of the research lifecycle. Data collection and management platforms in particular must cope with both issues involving language and translation and the varying ethical and legal standards that govern private data.

Fortunately, some promising initiatives already exist that, if they are properly supported, may help resolve these more specialized problems. For example, the *formr* experiment platform (Arslan et al., 2019) relies on some of the infrastructure that is already in place, namely Google Sheets, to permit easy, flexible, and collaborative construction of online experiments. Another promising project in this vein is the *experimentum* framework (https://psa.psy.gla.ac.uk/), which aims to develop an experiment builder and project management framework that is specialized for large, cross-linguistic, big team science settings. Finally, the *ScienceVerse* project (Lakens & DeBruine, 2021) aims to create a fully functional "grammar of science" that can be used for naming, describing, and identifying relationships between components of any scientific project. Such a grammar would be especially helpful for highly collaborative, cross-linguistic, big team science projects that must document a large number of components while navigating specialized ethical and legal standards.

These initiatives will only succeed in meeting the specialized infrastructure needs of big team science in the social sciences if they are supported financially and on a sustainable basis. Funders should recognize that these and other initiatives that support big team science also provide benefits for small team science. For example, although *formr* is especially useful in team settings, it provides a useful platform for teams of any size. Smart and ongoing investment in these projects will go a long way toward facilitating both small team and big team science.

Institutions. Another set of obstacles relates to the institutions that have developed assuming that projects are led by a solo (or small number of) Principal Investigators. These institutions cause friction in almost every aspect of large, collaborative science. The frictions start with funding mechanisms, most of which assume that projects are led by a sole PI and their staff. The European Research Council's Consolidator Grant, for example, supports a single PI and their host institution, and the European Research Council enforces highly detailed accounting rules to ensure all money is used in support of a solo PI's project. As another example, some grant funders, such as the US National Science Foundation, require applicants to list all their collaborators from a specified time period. This requirement places burdensome restrictions on prospective big team scientists whose collaborators can easily number in the hundreds. Even once money is awarded, administrative and legal barriers can make it difficult to send the money to the institutions and countries where it is needed.

Similar frictions plague almost every part of the big team science research lifecycle. Academic psychology departments do not typically train scientists to operate in large, distributed team settings and do not have specialized training tracks for scientists who wish to specialize in the many specific roles that big team science requires. Ethical Review Boards are often not prepared to evaluate a project that will be executed at a hundred sites worldwide. A big team science project with 100 collaborating labs may need to submit 100 variations of the same application to 100 Ethical Review Boards – a process that leads to immense waste and duplication of effort (Ervin et al., 2016; Schneider, 2015, pg. 44). Even the content of ethical regulations themselves differs across locations due to varying laws and norms. Some

organizations are not subject to clear ethical requirements at all – or, as is the case in East Africa, for example – they are only able to complete an ethical review in exchange for a costly fee that competes with other research costs (Kaplay, 2016; Osborne & Luoma, 2018).

As another example, journal submission portals are built assuming that articles are written by a relatively small number of authors – not 200 – and the process of entering author information is cumbersome and wasteful. A similar barrier obstructs the entry of authorship information in preprint servers like PsyArXiv. Once research is complete, tenure and evaluation committees do not know how to evaluate publications with hundreds of authors. This decision is high stakes, because if researchers that prioritize big team science are penalized by these committees, these "big team scientists" will be effectively selected out of the ecology of science (Smaldino & McElreath, 2016).

Here we are less certain what the future holds. If the institutions of psychological science adapt to the emergence of big team science, they will make this sort of science less costly and therefore more common. Alternatively, psychological science could develop an entirely *new* set of institutions, such as funders, ethical review boards, and journals, that are more accommodating of big team science. Whatever the future may hold, we hope that the institutions of science recognize the potential of big team science and act accordingly.

Risks of big team science

Although we believe that big team science approaches have great promise, they also bring a unique set of risks that stem from the very feature that provides their main benefits: the fact that they coordinate and centralize resources to be deployed in a single project (Stokols et al., 2008). These risks are *unaccountable leadership*, *management failures*, *conservatism*, *sustainability failures*, and *mega-mistakes*.

Unaccountable leadership. The primary virtue of big team science is its ability to deploy resources at a large scale. Often, this requires turning over resources to a single scientific organization, which carries the same risk as that entailed in turning over resources to a single governmental organization: the leadership of that organization could use the resources to centralize power within themselves and make themselves unaccountable to influence and criticism.

The negative consequences of unaccountable leadership power can impact many features of the scientific process. For example, in the idea generation phase, unaccountable leadership can select their own pet topics for investigation without heeding the ideas of people with less power, leading to ideas that are less creative and impactful (Wu et al., 2019). Unaccountable leadership can also stifle diversity in the strategies that are selected to develop and test the ideas that are selected (Devezer et al., 2019), slowing the pace of discovery. Finally, unaccountable leadership can have negative consequences for individual scholars who do not conform to leadership's perhaps narrow expectations of who is a proper "big team scientist" – especially if the leadership has influence over important mechanisms of career advancement (Azoulay et al., 2019).

Risk	Description	Mitigation strategies
Unaccountable leadership	Leadership of big team science organizations make themselves unaccountable to influence and criticism	Articulate organizational values Create structured bylaws Introduce democratic accountability Empower a board of directors
Management failures	Big team science organizations involve large numbers of relationships, and navigating the interpersonal issues that arise requires effective management of people	Articulate roles and responsibilities Create structured methods of making decisions Create formal and informal ways to solicit team member feedback Obtain formal management training
Conservatism	Big team science organizations may make decisions that cater to the median desire of the people in their group	Proactively partner with organizations that serve underrepresented groups Create mechanisms to facilitate mobility and advancement within the team Separate idea generation from project implementation
Sustainability failures	Big team science organizations require planning to be financially sustainable over the long term	Pay people the people who are responsible for the organization's maintenance Create and follow a sustainability plan
Mega-mistakes	When big team science organizations make mistakes, they tend to be big ones	Mitigate the other four types of risk Institute pro-active quality control processes

Table 4. Risks of team science and some risk mitigation strategies.

Mitigating this risk. At a baseline, big team science organizations should have a set of bylaws that lay out who is empowered to do what and a set of values. The bylaws need not specify a completely de-centralized power structure. Centralizing power within a small number of leaders can have important benefits, such as enabling those leaders to develop a coherent organizational strategy and allowing those leaders to make fast, agile decisions to adjust that strategy (Baum & Wally, 2003). However, to mitigate the risk of a lack of accountability, centralization of leadership must be accompanied with mechanisms to hold leadership accountable for its decisions if those leaders pursue directions with which important stakeholders disagree.

An important tool for enabling such accountability is a clearly articulated statement of the values that guide the big team science organization's mission. Once articulated, these values can

allow the organization's stakeholders to identify leadership mistakes that might require accountability by allowing those stakeholders to evaluate whether leadership actions do indeed fit with the organization's values. Insofar as leaders do issue regular statements about their organizational strategy, a set of well-articulated values also allows stakeholders to evaluate whether that strategy is indeed pursuing those values.

However, values are empty unless a group of organization stakeholders are empowered to take action if and when leaders violate those values. One example of an empowered stakeholder is a board of directors that is empowered by the organization's bylaws to remove leaders with whom they are unsatisfied. This board of directors should receive regular updates on the status of the big team science organization so that they can ask leaders to make corrections before a crisis occurs. Another possibility is to introduce explicit democratic mechanisms into the organization, such as regular elections for leadership positions. Whatever the form, accountability must be built into the organization's structure if it is to effectively check the risks of unaccountable leadership.

Management failures. The number of relationships within a team increases combinatorially with the number of people within the team. Larger teams also create more opportunities for people with widely different backgrounds to collide – or for single trouble-makers to create a toxic environment. Finally, larger teams are more likely to be characterized by specialized roles and communication channels. These features vastly increase the complexity of relationships in very large teams. In failure cases, this complexity can lead to *management failures*.

The first of these failures is role ambiguity. Role ambiguity occurs when the information a person has about the expectations that go along with the role, the methods for fulfilling those expectations, and the consequences for violating those expectations are unclear (Van Sell et al., 1981). Role ambiguity can occur in small groups, but the number of roles in a large scientific team, combined with the specialization of roles and complex decision-making apparatus, increase the risk of role ambiguity as teams increase in size. Role ambiguity increases the risk of errors because team members are uncertain which responsibilities apply to them. This uncertainly can lead to duplicated work on the one hand and unfulfilled tasks on the other. Role ambiguity also breeds feelings of dissatisfaction, leading people to grow disillusioned with the big team science organization and, in extreme failure cases, to leave it (Tubre & Collins, 2000).

The second failure is ineffective management of interpersonal issues. Some interpersonal issues, such as conflict between team members, are inevitable due simply to the number of relationships involved in big team science organizations. Some level of conflict can even be constructive when it is managed properly and kept task-focused (Forsyth, 2014; Loughry & Amason, 2014; Rahim, 2003). However, conflict can also create organizational dysfunction, potentially polarizing the team and derailing entire projects (Rahim, 2003). In the presence of sharp power differences between team members and unclear mechanisms for accountability, conflict can also take on a darker guise in the form of harassment and abuse (Berdahl & Raver, 2011). The risks of conflict can be especially high when team members have sharp cultural or epistemic differences, as might be expected in big team science settings that bring together people of varying personal, cultural, and disciplinary backgrounds (Bender et al., 2015).

Mitigating this risk. The most fundamental step a big team science organization can take to mitigate the risk of organizational failures is to establish a well-defined leadership structure that clearly communicates and enforces the organization's roles and decisions (Stokols, 2006). The leadership structure should also create both formal and informal ways of soliciting feedback from members of the big team science organization. These feedback mechanisms should include a means to report problematic behavior by individual team members, along with a means to hold problematic team members accountable, such as reporting the behavior to the team member's workplace. The feedback mechanisms should establish a feedback loop between decision-makers and team members so that the organization responds effectively to issues as they arise. The goal of the feedback loop will not eliminate interpersonal issues, as some number of issues are simply part of the human condition, but rather to instill a sense that procedures are fair when issues do occur (Konovsky, 2000). Effective communication also fosters an environment of psychological safety in which team members feel empowered to express issues, concerns, and points of disagreement without creating outright fights (Frazier et al., 2017).

Clarifying roles and instilling psychological safety both take time. However, their importance to effective collaboration means that team science organizations should prioritize creating structures that allow transparent communication, decision-making, and enforcement, ideally through a formal *collaboration agreement* that is drafted before the collaboration begins. Structured methods of communication, decision-making, and enforcement lay the groundwork to allow team members to build the mutual trust that is necessary for collaboration (Astuti & Bloch, 2012).

Management of teams is a vital skill that is seldom taught in the academy. Universities should recognize the central role that management plays in the success of projects, especially the larger ones that typify big team science. If big team science is to take root in the social sciences, management training needs to become a central part of the behavioral science curriculum.

Conservatism. Large organizations require large amounts of people to keep themselves running. If these organizations attempt to satisfy the desires of these large numbers of people, they will tend to cater to the median desire. This means that organizations will usually be *conservative* – or at least conservative with respect to the people within the group. This conservatism can manifest in two ways: in the selection of personnel and the selection of projects.

People tend to form social relationships with others who have similar characteristics and backgrounds (Kossinets & Watts, 2009; McPherson et al., 2001). Given that contemporary psychology is dominated by North America and Europe (Rad et al., 2018; Thalmayer et al., 2020), this raises a risk that big team science organizations will be similarly dominated by people from those continents. For example, not a single first author of the papers in Table 2 comes from a nation outside North America and Europe. Once big team science organizations dominated by North America and Europe are established, they may inadvertently crowd out organizations from elsewhere with different goals and personnel. Systematically excluding large subsets of humanity from psychology perpetuates unfair systems of inequality and can lead to a science that focuses unduly on the preoccupations of a small subset of humanity (Medin et al., 2017).

Conservatism can also manifest in project selection. The primary purpose of big team science is to scale up the resources that can be invested in a given project. However, this very virtue may induce big team science collaborations to select projects that they perceive as "safe" – both in terms of whether they can be feasibly executed and in terms of the degree to which they deviate from the scientific mainstream. Given that some degree of theoretical risk is necessary to spur scientific progress (Devezer et al., 2019; Meehl, 1978), this theoretical conservatism could feasibly slow scientific progress.

Mitigating this risk. To mitigate the risk of conservatism in personnel selection, we can draw on a general slate of strategies to improve the inclusion of underrepresented groups in science (Enriquez, 1979; Henrich et al., 2010; Syed & Kathawalla, 2020). These solutions must begin at the start of big team science organizations and must be continually re-evaluated throughout the organization's lifecycle. At the organization's founding, the founders must proactively and systematically partner with researchers in a broad array of non-Western and non-elite institutions – and especially those who are not part of the "virtual academic commons" formed by Twitter, Facebook, and science blogs. Once the organization has been founded, the organization's leadership must create mechanisms that allow for mobility within the team so that members of underrepresented groups can rise to positions of leadership. However, even these measures will not be sustainable without the direct investment of material resources into research infrastructure in under-resourced world regions.

Diversity in big team science personnel will not ensure diversity in big team science projects. Mitigating the risk of this type of conservatism requires maintaining a separation between the *idea generation* and the *project implementation* phases of big team science projects. During the idea generation phase, much smaller groups of scientists can identify problems and approaches unconstrained from a broader team consensus. The smaller teams then develop proposals based on their ideas and submit them for consideration by the larger consortium. The larger team can even explicitly build in mechanisms to solicit proposals from teams whose perspectives may differ from the scientific mainstream – such as those from outside North America and Europe.

Sustainability failures. The history of science is littered with promising initiatives and organizations that, once established, could not be sustained (Borgman et al., 2016). This problem recurs so frequently because scientific organizations are typically public-minded: they wish to provide a public good at little or no cost. Yet, this very public-mindedness invites free-ridership – people who will use the organization's service but who are either unwilling or unable to support the organization financially (Neylon, 2017).

This dynamic also threatens big team science organizations. All the initiatives listed in Table 2 were formed because the founders thought that large-scale collaboration could lead to better science rather than out of a desire to make money. If the initiative is intended to last for a single project, the project could survive using an ad hoc organizational structure run on the back of volunteer labor. However, continually creating one-shot organizations is wasteful and inefficient because such a model prevents the accumulation of organizational knowledge. If an initiative is to last beyond a single project, its leadership must at some point figure out how to generate the money required to retain long-term staff while not compromising the vision that inspired the project in the first place.

Mitigating this risk. The first step to mitigating the risk of a sustainability failure is to recognize that organizations cannot support themselves for free. If we want to have organizations that are capable of coordinating very large groups of scientists, those organizations must receive financial support so that they can generate and maintain the knowledge, staff, and other scientific resources necessary to fulfill their function. A corollary of this dictum is that, over the long haul, at least some skilled positions in a big team science organization must receive monetary compensation. Running entirely on volunteerism risks exploiting the idealistic people willing to donate their time (Kim et al., 2020), leading to burnout and interpersonal conflict.

The second step to mitigating this risk involves creating and following a *sustainability plan* (for an example, see Forscher & IJzerman, 2021) that maps out how the organization will generate the funds necessary to maintain itself. The funding models could entail what is most familiar to academic scientists: designating some team members as grantwriters who spend part of their time identifying and applying for large grants to support the big team science organization. However, this grantwriting model puts big team science organizations at the whims of large funders, who are often only willing to spend money on projects that generate discoveries, not the infrastructure that makes the discoveries possible (Zakaria et al., 2021). Most science grantmakers also rely on competitive peer review to select the most "worthy" projects, a funding model that risks inefficiency because it encourages grantwriters to spend more time on grants than the activities the grants would fund (Gross & Bergstrom, 2019). For these reasons, we believe that a grantwriting-based funding model carries inherent risks of both research waste and instability in personnel due to feast-or-famine funding cycles.

We believe that a more promising approach is to use one of the funding models that have led to sustainable funding for other large public-minded scientific organizations (Neylon, 2017). These funding models will likely involve either creating a system that "taxes" all beneficiaries of the big team science infrastructure by, say, imposing membership dues rather as a scientific society does, or, alternatively providing the infrastructure as a byproduct of selling another service. The other service could involve many things, but it should leverage what big team science organizations already do well – running multi-site studies. Thus, the organization might run some multi-site studies for a fee, or it might sell a service that it has needed to perfect in order to run multi-site studies, such as translation or project management. We believe that all these options are viable, though each has their own tradeoffs; membership dues could make the big team science organization inaccessible to lower-resourced members, whereas selling services could introduce conflicts of interest that threaten the organization's mission. Choosing and following a sustainability plan is necessary to maintain any public-minded organization, but the tradeoffs of the available options must also be evaluated with care.

Mega-mistakes. The final risk of big team science is that of what we term *mega-mistakes*. Big team science's primary virtue is its ability to "scale up" small projects into big ones by pooling resources across labs. However, this very virtue makes errors all the more costly: these errors risk wasting much more resources than would be wasted on a smaller project. These errors can even occur at the time of project selection if the topic of the project is not one that deserves the high investment of resources that big team science brings.

Errors in big team projects often have no one simple cause. Consider the example of the Human Brain Project. This project aimed to unlock massive advances for neuroscience by

developing and building large-scale computer simulations of brain regions, and eventually of entire brains (Markram, 2012). The leader of the project, Henry Markram, set an ambitious agenda; in 2009, he claimed that, after only ten more years of development the technology that powered the project would allow the simulation of an entire rodent brain, an achievement that he believed would revolutionize neuroscience (Abbott, 2020). The project's vision and ambition attracted major investment from both universities and funders; in 2015, the project had 112 university partners and a project budget of over €1 billion (Abbott, 2015).

Yet the project was plagued with problems. These problems spanned multiple domains, ranging from concerns that the project was monopolizing resources that could be allocated to other worthy projects, to disagreement over the project's core aims, to dissatisfaction with the project's autocratic leadership structure (Abbott, 2014). Regardless of the specific causes, by 2015 most scientists agreed that the project was not living up to Henry Markram's original vision (Abbott, 2015). Markram was removed from his leadership position the following year (Abbott, 2020). The project is still ongoing, but whether it will produce the revolution in neuroscience that Markram promised is anything but clear.

Mitigating this risk. As the example of the Human Brain Project illustrates, mega-mistakes can emerge due to a failure to manage other risks of big team science, such as the risk of mis-management. Thus, mitigating other risks of big team science may be one useful way of mitigating the risks of mega-mistakes.

However, big team science organizations can also proactively implement quality control processes that minimize the risks of mega-mistakes more directly. These quality control processes entail everything from formally reviewing project proposals, to instituting code review, pilot tests, and project "soft launches", to formal methods of optionally stopping data collection to avoid overinvestment in bad ideas (Lakens, 2014; Schönbrodt et al., 2017). Quality control can also involve partnering with outside scientific organizations that implement their own methods of quality control, such as journals that administer Registered Reports.

An additional promising strategy involves instituting audits of scientific processes by people who are formally independent of the big team science collaboration. These audits can take the form of research design review by outside experts, systems of back-translation to check the quality of forward-translation, formal code review, and even "red teams" who receive bounties for spotting bugs and other project flaws (Lakens, 2020).

All these quality control methods have high value in solo PI projects, but take on renewed importance in projects that command resources on the scale of big team science.

Conclusion

Big team science will never be a wholesale replacement to solo science – nor should it be. Absent the coordination constraints of a large team, solo scientists have the freedom to flexibly explore ideas that might be infeasible in a larger group setting without creating a risk of squandering a scale of resources that rises to the level of a "mega-mistake".

What big team science can do is vastly "scale up" the amount of resources – in other words, the money, person-hours, and specialized expertise – that can be deployed in a single,

well-chosen and well-vetted project. This increased resource investment can result in projects that are larger, more rigorous, and more representative of humanity. We believe this virtue is not to be underestimated, as low resource investment is likely a key ultimate cause of why psychology's progress has been slow. Moreover, other disciplines may find themselves facing similar problems. Ecologists identified similar problems in their own discipline, and team science is emerging as a potential solution (Fraser et al., 2013). Pre-clinical cancer biology faces problems even conducting the studies needed to assess whether a replicability problem exists (Friedl, 2019). A greater focus on big team science may help pre-clinical cancer biology "scale up" the resources devoted to the typical study in the discipline, which may help lay the groundwork to allow these replication studies to be conducted.

We believe that big team science has the potential to simultaneously tackle the many challenges that psychology faces. However, to fulfill this potential, psychological scientists must recognize and manage the many barriers and risks that this approach entails. If properly managed to leverage its virtues while mitigating its risks, we believe that big team science can be instrumental in the movement to build more reliable, informative, and rigorous science.

References

23andMe Research Team, COGENT (Cognitive Genomics Consortium), Social Science Genetic Association Consortium, Lee, J. J., Wedow, R., Okbay, A., Kong, E., Maghzian, O., Zacher, M., Nguyen-Viet, T. A., Bowers, P., Sidorenko, J., Karlsson Linnér, R., Fontana, M. A., Kundu, T., Lee, C., Li, H., Li, R., Royer, R., … Cesarini, D. (2018). Gene discovery and polygenic prediction from a genome-wide association study of educational attainment in 1.1 million individuals. *Nature Genetics*, *50*(8), 1112–1121. https://doi.org/10.1038/s41588-018-0147-3

Abbott, A. (2014). Row hits flagship brain plan. *Nature*, *511*(7508), 133–134. https://doi.org/10.1038/511133a

Abbott, A. (2015). Human Brain Project votes for leadership change. *Nature*, nature.2015.17060. https://doi.org/10.1038/nature.2015.17060

Abbott, A. (2020). Documentary follows implosion of billion-euro brain project. *Nature*, *588*(7837), 215–216. https://doi.org/10.1038/d41586-020-03462-3

ALICE collaboration. (2010). Alignment of the ALICE Inner Tracking System with cosmic-ray tracks. *Journal of Instrumentation*, *5*(03), P03003–P03003. https://doi.org/10.1088/1748-0221/5/03/P03003

Arslan, R. C., Walther, M. P., & Tata, C. S. (2019). formr: A study framework allowing for automated feedback generation and complex longitudinal experience-sampling studies using R. *Behavior Research Methods*. https://doi.org/10.3758/s13428-019-01236-y

Astuti, R., & Bloch, M. (2012). Anthropologists as Cognitive Scientists. *Topics in Cognitive Science*, *4*(3), 453–461. https://doi.org/10.1111/j.1756-8765.2012.01191.x

Azoulay, P., Fons-Rosen, C., & Zivin, J. S. G. (2019). Does Science Advance One Funeral at a

   Time? *American Economic Review*, *109*(8), 2889–2920.

   https://doi.org/10.1257/aer.20161574

Bakker, C. J. (1955). CERN: European Organization for Nuclear Research. *Physics Today*, *8*, 8–

   13.

Baum, J. R., & Wally, S. (2003). Strategic decision speed and firm performance. *Strategic

   Management Journal*, *24*(11), 1107–1129. https://doi.org/10.1002/smj.343

Bender, A., Beller, S., & Nersessian, N. J. (2015). Diversity as Asset. *Topics in Cognitive

   Science*, *7*(4), 677–688. https://doi.org/10.1111/tops.12161

Berdahl, J. L., & Raver, J. L. (2011). Sexual harassment. In S. Zedeck (Ed.), *APA handbook of

   industrial and organizational psychology, Vol 3: Maintaining, expanding, and

   contracting the organization.* (pp. 641–669). American Psychological Association.

   https://doi.org/10.1037/12171-018

Birnholtz, J. (2008). When Authorship Isn't Enough: Lessons from CERN on the Implications of

   Formal and Informal Credit Attribution Mechanisms in Collaborative Research. *The

   Journal of Electronic Publishing*, *11*(1). https://doi.org/10.3998/3336451.0011.105

Borgman, C. L., Darch, P. T., Sands, A. E., & Golshan, M. S. (2016). The Durability and

   Fragility of Knowledge Infrastructures: Lessons Learned from Astronomy.

   ArXiv:1611.00055 [Astro-Ph]. http://arxiv.org/abs/1611.00055

Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M.,

   Kirchler, M., Iwanir, R., Mumford, J. A., Adcock, A., Avesani, P., Baczkowski, B.,

   Bajracharya, A., Bakst, L., Ball, S., Barilari, M., Bault, N., Beaton, D., Beitner, J., …

Schonberg, T. (2019). *Variability in the analysis of a single neuroimaging dataset by many teams* [Preprint]. Neuroscience. https://doi.org/10.1101/843193

Brumfiel, G. (2011). High-energy physics: Down the petabyte highway. *Nature*, *469*(7330), 282–283. https://doi.org/10.1038/469282a

Bryant, P. J. (1994). *A Brief History and Review of Accelerators* (No. 261062). CERN.

Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, *14*(5), 365–376. https://doi.org/10.1038/nrn3475

Byers-Heinlein, K., Bergmann, C., Davies, C., Frank, M. C., Hamlin, J. K., Kline, M., Kominsky, J. F., Kosie, J. E., Lew-Williams, C., Liu, L., Mastroberardino, M., Singh, L., Waddell, C. P. G., Zettersten, M., & Soderstrom, M. (2020). Building a collaborative psychological science: Lessons learned from ManyBabies 1. *Canadian Psychology/Psychologie Canadienne*, *61*(4), 349–363. https://doi.org/10.1037/cap0000216

Coles, N. A., March, D. S., Marmolejo-Ramos, F., Arinze, N. C., Ndukaihe, I. L. G., Özdoğru, A. A., Aczel, B., Hajdu, N., Nagy, T., Basnight-Brown, D., Zambrano, D., Foroni, F., Willis, M., Pfuhl, G., Kaminski, G., IJzerman, H., Vezirian, K., Banaruee, H., Suarez, I., … Liuzza, M. T. (2019). *A Multi-Lab Test of the Facial Feedback Hypothesis by The Many Smiles Collaboration* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/cvpuw

Corvin, A., Craddock, N., & Sullivan, P. F. (2010). Genome-wide association studies: A primer. *Psychological Medicine*, *40*(7), 1063–1077. https://doi.org/10.1017/S0033291709991723

Cuccolo, K., Irgens, M. S., Zlokovich, M. S., Grahe, J., & Edlund, J. E. (2021). What Crowdsourcing Can Offer to Cross-Cultural Psychological Science. *Cross-Cultural Research*, *55*(1), 3–28. https://doi.org/10.1177/1069397120950628

Dannlowski, U., Ohrmann, P., Bauer, J., Deckert, J., Hohoff, C., Kugel, H., Arolt, V., Heindel, W., Kersting, A., Baune, B. T., & Suslow, T. (2008). 5-HTTLPR Biases Amygdala Activity in Response to Masked Facial Expressions in Major Depression. *Neuropsychopharmacology*, *33*(2), 418–424. https://doi.org/10.1038/sj.npp.1301411

Devezer, B., Nardin, L. G., Baumgaertner, B., & Buzbas, E. O. (2019). Scientific discovery in a model-centric framework: Reproducibility, innovation, and epistemic diversity. *PLOS ONE*, *14*(5), e0216125. https://doi.org/10.1371/journal.pone.0216125

Donoho, D., Arian, M., Imam, R., Morteza, S., & Stodden, V. (2008). *15 Years of Reproducible Research in Computational Harmonic Analysis* (pp. 1–26) [Technical Report]. Department of Statistics, Stanford University. http://statweb.stanford.edu/~wavelab/DonohoEtAlCISESubmission.pdf

Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., Baranski, E., Bernstein, M. J., Bonfiglio, D. B. V., Boucher, L., Brown, E. R., Budiman, N. I., Cairo, A. H., Capaldi, C. A., Chartier, C. R., Chung, J. M., Cicero, D. C., Coleman, J. A., Conway, J. G., … Nosek, B. A. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, *67*, 68–82. https://doi.org/10.1016/j.jesp.2015.10.012

Enriquez, V. G. (1979). Towards cross-cultural knowledge through cross-indigenous methods and perspective. *Philippine Journal of Psychology*, *12*(1), 9–15.

Ervin, A.-M., Taylor, H. A., & Ehrhardt, S. (2016). NIH Policy on Single-IRB Review—A New Era in Multicenter Studies. *New England Journal of Medicine*, *375*(24), 2315–2317. https://doi.org/10.1056/NEJMp1608766

Espinosa, J. P. A., Levcovitz, G. B., Bianchi, R.-M., Brock, I., Carli, T., Castro, N. F., Ciocio, A., Colautti, M., Menezes, A. C. D. S., da Fonseca, G. D. O., Alves, L. D. M., Hoecker, A., Ramos, B. L., Pinhão, G. L. L., Maidantchik, C., Malek, F., McPherson, R., Picco, G., & Santos, M. T. D. (2020). The ATLAS Publication Process Supported by Continuous Integration and Web Framework. *ArXiv:2005.06989 [Hep-Ex]*. http://arxiv.org/abs/2005.06989

Fanelli, D. (2010). "Positive" results increase down the hierarchy of the sciences. *PLoS ONE*, *5*(4), e10068. https://doi.org/10.1371/journal.pone.0010068

Finkel, E. J., Eastwick, P. W., & Reis, H. T. (2017). Replicability and other features of a high-quality science: Toward a balanced and empirical approach. *Journal of Personality and Social Psychology*, *113*(2), 244–253. https://doi.org/10.1037/pspi0000075

Forscher, P. S., & IJzerman, H. (2021, January 11). How should we fund the PSA? *Psychological Science Accelerator Blog*. https://psysciacc.org/2021/01/11/how-should-we-fund-the-psa/

Forsyth, D. R. (2014). *Group dynamics* (6th ed). Wadsworth Cengage Learning.

Fraser, L. H., Henry, H. A., Carlyle, C. N., White, S. R., Beierkuhnlein, C., Cahill, J. F., Casper, B. B., Cleland, E., Collins, S. L., Dukes, J. S., Knapp, A. K., Lind, E., Long, R., Luo, Y., Reich, P. B., Smith, M. D., Sternberg, M., & Turkington, R. (2013). Coordinated distributed experiments: An emerging tool for testing global hypotheses in ecology and

environmental science. *Frontiers in Ecology and the Environment*, *11*(3), 147–155.

https://doi.org/10.1890/110279

Frazier, M. L., Fainshmidt, S., Klinger, R. L., Pezeshkan, A., & Vracheva, V. (2017).

Psychological Safety: A Meta-Analytic Review and Extension. *Personnel Psychology*,

70(1), 113–165. https://doi.org/10.1111/peps.12183

Friedl, P. (2019). Rethinking research into metastasis. *ELife*, *8*, e53511.

https://doi.org/10.7554/eLife.53511

Gillespie, N. A., Whitfield, J. B., Williams, B., Heath, A. C., & Martin, N. G. (2005). The

relationship between stressful life events, the serotonin transporter (5-HTTLPR) genotype

and major depression. *Psychological Medicine*, *35*(1), 101–111.

https://doi.org/10.1017/S0033291704002727

Gotlib, I. H., Joormann, J., Minor, K. L., & Hallmayer, J. (2008). HPA Axis Reactivity: A

Mechanism Underlying the Associations Among 5-HTTLPR, Stress, and Depression.

Biological Psychiatry, *63*(9), 847–851. https://doi.org/10.1016/j.biopsych.2007.10.008

Gross, K., & Bergstrom, C. T. (2019). Contest models highlight inherent inefficiencies of

scientific funding competitions. *PLOS Biology*, *17*(1), e3000065.

https://doi.org/10.1371/journal.pbio.3000065

Hardwicke, T. E., Mathur, M. B., MacDonald, K., Nilsonne, G., Banks, G. C., Kidwell, M. C.,

Hofelich Mohr, A., Clayton, E., Yoon, E. J., Henry Tessler, M., Lenne, R. L., Altman, S.,

Long, B., & Frank, M. C. (2018). Data availability, reusability, and analytic

reproducibility: Evaluating the impact of a mandatory open data policy at the journal

Cognition. *Royal Society Open Science*, *5*(8), 180448.

https://doi.org/10.1098/rsos.180448

Harry, G. M. & the LIGO Scientific Collaboration. (2010). Advanced LIGO: The next generation of gravitational wave detectors. *Classical and Quantum Gravity*, *27*(8), 084006. https://doi.org/10.1088/0264-9381/27/8/084006

Heils, A., Teufel, A., Petri, S., Stöber, G., Riederer, P., Bengel, D., & Lesch, K. P. (1996). Allelic Variation of Human Serotonin Transporter Gene Expression. *Journal of Neurochemistry*, *66*(6), 2621–2624. https://doi.org/10.1046/j.1471-4159.1996.66062621.x

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*(2–3), 61–83. https://doi.org/10.1017/S0140525X0999152X

Holcombe, A. O. (2019). Contributorship, Not Authorship: Use CRediT to Indicate Who Did What. *Publications*, *7*(3), 48. https://doi.org/10.3390/publications7030048

Jones, B. C., DeBruine, L. M., Flake, J. K., Liuzza, M. T., Antfolk, J., Arinze, N. C., Ndukaihe, I. L. G., Bloxsom, N. G., Lewis, S. C., Foroni, F., Willis, M. L., Cubillas, C. P., Vadillo, M. A., Turiegano, E., Gilead, M., Simchon, A., Saribay, S. A., Owsley, N. C., Jang, C., … Coles, N. A. (2021). To which world regions does the valence-dominance model of social perception apply? *Nature Human Behaviour*. https://doi.org/10.1038/s41562-020-01007-2

Kaplay, S. (2016, July 6). In clinical trials, for-profit review boards are taking over for hospitals. Should they? *STAT*. https://www.statnews.com/2016/07/06/institutional-review-boards-commercial-irbs/

Kim, J. Y., Campbell, T. H., Shepherd, S., & Kay, A. C. (2020). Understanding contemporary forms of exploitation: Attributions of passion serve to legitimize the poor treatment of

workers. *Journal of Personality and Social Psychology*, *118*(1), 121–148.

    https://doi.org/10.1037/pspi0000190

Klein, R. A., Cook, C. L., Ebersole, C. R., Vitiello, C. A., Nosek, B. A., Chartier, C. R.,

    Christopherson, C. D., Clay, S., Collisson, B., Crawford, J., Cromar, R., Dudley, D.,

    Gardiner, G., Gosnell, C., Grahe, J. E., Hall, C., Joy-Gaba, J. A., Legg, A. M., Levitan,

    C., … Ratliff, K. A. (2019). *Many Labs 4: Failure to Replicate Mortality Salience Effect*

    With and Without Original Author Involvement [Preprint]. PsyArXiv.

    https://doi.org/10.31234/osf.io/vef2c

Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, Š., Bernstein, M. J., Bocian, K.,

    Brandt, M. J., Brooks, B., Brumbaugh, C. C., Cemalcilar, Z., Chandler, J., Cheong, W.,

    Davis, W. E., Devos, T., Eisner, M., Frankowska, N., Furrow, D., Galliani, E. M., …

    Nosek, B. A. (2014). Investigating Variation in Replicability: A "Many Labs"

    Replication Project. *Social Psychology*, *45*(3), 142–152. https://doi.org/10.1027/1864-

    9335/a000178

Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M.,

    Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D.

    R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., … Nosek, B.

    A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and

    Settings. *Advances in Methods and Practices in Psychological Science*, *1*(4), 443–490.

    https://doi.org/10.1177/2515245918810225

Konovsky, M. A. (2000). Understanding Procedural Justice and Its Impact on Business

    Organizations. *Journal of Management*, *26*(3), 489–511.

    https://doi.org/10.1177/014920630002600306

Kossinets, G., & Watts, D. J. (2009). Origins of Homophily in an Evolving Social Network. *American Journal of Sociology*, *115*(2), 405–450. https://doi.org/10.1086/599247

Krige, J. (2004). I. I. Rabi and the Birth of CERN. *Physics Today*, *57*(9), 44–48. https://doi.org/10.1063/1.1809091

Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J.-E. L., Joy-Gaba, J. A., Ho, A. K., Teachman, B. A., Wojcik, S. P., Koleva, S. P., Frazier, R. S., Heiphetz, L., Chen, E. E., Turner, R. N., Haidt, J., Kesebir, S., Hawkins, C. B., Schaefer, H. S., Rubichi, S., … Nosek, B. A. (2014). Reducing implicit racial preferences: I. A comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, *143*(4), 1765–1785. https://doi.org/10.1037/a0036260

Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., Calanchini, J., Xiao, Y. J., Pedram, C., Marshburn, C. K., Simon, S., Blanchar, J. C., Joy-Gaba, J. A., Conway, J., Redford, L., Klein, R. A., Roussos, G., Schellhaas, F. M. H., Burns, M., … Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, *145*(8), 1001–1016. https://doi.org/10.1037/xge0000179

Lakens, D. (2014). Performing high-powered studies efficiently with sequential analyses: Sequential analyses. *European Journal of Social Psychology*, *44*(7), 701–710. https://doi.org/10.1002/ejsp.2023

Lakens, D. (2020). Pandemic researchers—Recruit your own best critics. *Nature*, *581*(7807), 121–121. https://doi.org/10.1038/d41586-020-01392-8

Lakens, D., & DeBruine, L. M. (2021). Improving Transparency, Falsifiability, and Rigor by Making Hypothesis Tests Machine-Readable. *Advances in Methods and Practices in*

Psychological Science, *4*(2), 251524592097094.

https://doi.org/10.1177/2515245920970949

Landy, J. F., Jia, M. (Liam), Ding, I. L., Viganola, D., Tierney, W., Dreber, A., Johannesson, M.,

Pfeiffer, T., Ebersole, C. R., Gronau, Q. F., Ly, A., van den Bergh, D., Marsman, M.,

Derks, K., Wagenmakers, E.-J., Proctor, A., Bartels, D. M., Bauman, C. W., Brady, W. J.,

… Uhlmann, E. L. (2020). Crowdsourcing hypothesis tests: Making transparent how

design choices shape research results. *Psychological Bulletin*, *146*(5), 451–479.

https://doi.org/10.1037/bul0000220

LeBel, E. P., Berger, D., Campbell, L., & Loving, T. J. (2017). Falsifiability is not optional.

Journal of Personality and Social Psychology, *113*(2), 254–261.

https://doi.org/10.1037/pspi0000106

Loughry, M., & Amason, A. (2014). Why won't task conflict cooperate? Deciphering stubborn

results. *International Journal of Conflict Management*, *25*(4), 333–358.

https://doi.org/10.1108/IJCMA-01-2014-0005

Markram, H. (2012). The Human Brain Project. *Scientific American*, *306*(6), 50–55. JSTOR.

McNutt, M. K., Bradford, M., Drazen, J. M., Hanson, B., Howard, B., Jamieson, K. H., Kiermer,

V., Marcus, E., Pope, B. K., Schekman, R., Swaminathan, S., Stang, P. J., & Verma, I. M.

(2018). Transparency in authors' contributions and responsibilities to promote integrity in

scientific publication. *Proceedings of the National Academy of Sciences*, *115*(11), 2557–

2560. https://doi.org/10.1073/pnas.1715374115

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social

Networks. *Annual Review of Sociology*, *27*(1), 415–444.

https://doi.org/10.1146/annurev.soc.27.1.415

Medin, D., Ojalehto, B., Marin, A., & Bang, M. (2017). Systems of (non-)diversity. *Nature Human Behaviour*, *1*(5), 0088. https://doi.org/10.1038/s41562-017-0088

Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, *46*, 806–834. https://doi.org/10.1037/0022-006X.46.4.806

Moshontz, H., Campbell, L., Ebersole, C. R., IJzerman, H., Urry, H. L., Forscher, P. S., Grahe, J. E., McCarthy, R. J., Musser, E. D., Antfolk, J., Castille, C. M., Evans, T. R., Fiedler, S., Flake, J. K., Forero, D. A., Janssen, S. M. J., Keene, J. R., Protzko, J., Aczel, B., … Chartier, C. R. (2018). The Psychological Science Accelerator: Advancing Psychology Through a Distributed Collaborative Network. *Advances in Methods and Practices in Psychological Science*, *1*(4), 501–515. https://doi.org/10.1177/2515245918797607

Moshontz, H., Ebersole, C. R., Weston, S. J., & Klein, R. A. (2019). *A Guide for Many Authors: Writing Manuscripts in Large Collaborations* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/92xhd

Newell, A. (1973). *You can't play 20 questions with nature and win: Projective comments on the papers of this symposium*. Visual Information Processing, Carnegie-Mellon University.

Neylon, C. (2017). Sustaining Scholarly Infrastructures through Collective Action: The Lessons that Olson can Teach us. *KULA: Knowledge Creation, Dissemination, and Preservation Studies*, *1*(1), 3. https://doi.org/10.5334/kula.7

Obels, P., Lakens, D., Coles, N. A., Gottfried, J., & Green, S. A. (2019). *Analysis of Open Data and Computational Reproducibility in Registered Reports in Psychology* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/fk8vh

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), aac4716–aac4716. https://doi.org/10.1126/science.aac4716

Osborne, T. L., & Luoma, J. B. (2018). Overcoming a primary barrier to practice-based research: Access to an institutional review board (IRB) for independent ethics review. *Psychotherapy*, *55*(3), 255–262. https://doi.org/10.1037/pst0000166

Rad, M. S., Martingano, A. J., & Ginges, J. (2018). Toward a psychology of Homo sapiens: Making psychological science more representative of the human population. *Proceedings of the National Academy of Sciences*, *115*(45), 11401–11405. https://doi.org/10.1073/pnas.1721165115

Rahim, M. A. (2003). Toward a Theory of Managing Organizational Conflict. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.437684

Rieckmann, N. (2009). Gene-Environment Interactions and Depression. *JAMA*, *302*(17), 1859. https://doi.org/10.1001/jama.2009.1578

Schneider, C. (2015). *The censor's hand the misregulation of human-subject research*. The MIT Press.

Schönbrodt, F. D., Wagenmakers, E.-J., Zehetleitner, M., & Perugini, M. (2017). Sequential hypothesis testing with Bayes factors: Efficiently testing mean differences. *Psychological Methods*, *22*(2), 322–339. https://doi.org/10.1037/met0000061

Sears, D. O. (1986). College sophomores in the laboratory: Influences of a narrow data base on social psychology's view of human nature. *Journal of Personality and Social Psychology*, *51*(3), 515–530. https://doi.org/10.1037/0022-3514.51.3.515

Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., Bahník, Š., Bai, F., Bannard, C., Bonnier, E., Carlsson, R., Cheung, F., Christensen, G., Clay, R., Craig,

M. A., Dalla Rosa, A., Dam, L., Evans, M. H., Flores Cervantes, I., … Nosek, B. A. (2018). Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science*, *1*(3), 337–356. https://doi.org/10.1177/2515245917747646

Simons, D. J., Holcombe, A. O., & Spellman, B. A. (2014). An Introduction to Registered Replication Reports at *Perspectives on Psychological Science*. *Perspectives on Psychological Science*, *9*(5), 552–555. https://doi.org/10.1177/1745691614543974

Smaldino, P. E., & McElreath, R. (2016). The natural selection of bad science. *Royal Society Open Science*, *3*(9), 160384. https://doi.org/10.1098/rsos.160384

Spellman, B. A. (2015). A Short (Personal) Future History of Revolution 2.0. *Perspectives on Psychological Science*, *10*(6), 886–899. https://doi.org/10.1177/1745691615609918

Stokols, D. (2006). Toward a Science of Transdisciplinary Action Research. *American Journal of Community Psychology*, *38*(1–2), 79–93. https://doi.org/10.1007/s10464-006-9060-5

Stokols, D., Misra, S., Moser, R. P., Hall, K. L., & Taylor, B. K. (2008). The Ecology of Team Science. *American Journal of Preventive Medicine*, *35*(2), S96–S115. https://doi.org/10.1016/j.amepre.2008.05.003

Sullivan, P. F. (2010). The Psychiatric GWAS Consortium: Big Science Comes to Psychiatry. *Neuron*, *68*(2), 182–186. https://doi.org/10.1016/j.neuron.2010.10.003

Syed, M., & Kathawalla, U.-K. (2020). *Cultural Psychology, Diversity, and Representation in Open Science* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/t7hp2

Teasley, S., & Wolinsky, S. (2001). Scientific Collaborations at a Distance. *Science*, *292*(5525), 2254–2255. https://doi.org/10.1126/science.1061619

Thalmayer, A. G., Toscanelli, C., & Arnett, J. J. (2020). The neglected 95% revisited: Is American psychology becoming less American? *American Psychologist*. https://doi.org/10.1037/amp0000622

The ATLAS collaboration, Aaboud, M., Aad, G., Abbott, B., Abdinov, O., Abeloos, B., Abidi, S. H., AbouZeid, O. S., Abraham, N. L., Abramowicz, H., Abreu, H., Abreu, R., Abulaiti, Y., Acharya, B. S., Adachi, S., Adamczyk, L., Adelman, J., Adersberger, M., Adye, T., … Zwalinski, L. (2018). Searches for heavy ZZ and ZW resonances in the ℓℓqq and ννqq final states in pp collisions at s = 13 $$ \sqrt{s}=13 $$ TeV with the ATLAS detector. *Journal of High Energy Physics*, *2018*(3), 9. https://doi.org/10.1007/JHEP03(2018)009

Tiokhin, L., Hackman, J., Munira, S., Jesmin, K., & Hruschka, D. (2019). Generalizability is not optional: Insights from a cross-cultural study of social discounting. *Royal Society Open Science*, *6*(2), 181386. https://doi.org/10.1098/rsos.181386

Tiokhin, L., Panchanathan, K., Smaldino, P. E., & Lakens, D. (2021). *Shifting the level of selection in science* [Preprint]. MetaArXiv. https://doi.org/10.31222/osf.io/juwck

Tubre, T. C., & Collins, J. M. (2000). Jackson and Schuler (1985) Revisited: A Meta-Analysis of the Relationships Between Role Ambiguity, Role Conflict, and Job Performance. *Journal of Management*, *26*(1), 155–169. https://doi.org/10.1177/014920630002600104

Uhlmann, E. L., Ebersole, C. R., Chartier, C. R., Errington, T. M., Kidwell, M. C., Lai, C. K., McCarthy, R. J., Riegelman, A., Silberzahn, R., & Nosek, B. A. (2019). Scientific Utopia III: Crowdsourcing Science. *Perspectives on Psychological Science*, *14*(5), 711–733. https://doi.org/10.1177/1745691619850561

Van Sell, M., Brief, A. P., & Schuler, R. S. (1981). Role Conflict and Role Ambiguity: Integration of the Literature and Directions for Future Research. *Human Relations*, *34*(1), 43–71. https://doi.org/10.1177/001872678103400104

Vuorre, M., & Curley, J. P. (2018). Curating Research Assets: A Tutorial on the Git Version Control System. *Advances in Methods and Practices in Psychological Science*, *1*(2), 219–236. https://doi.org/10.1177/2515245918754826

Wagge, J. R., Baciu, C., Banas, K., Nadler, J. T., Schwarz, S., Weisberg, Y., IJzerman, H., Legate, N., & Grahe, J. (2019). A Demonstration of the Collaborative Replication and Education Project: Replication Attempts of the Red-Romance Effect. *Collabra: Psychology*, *5*(1), 5. https://doi.org/10.1525/collabra.177

Wagge, J. R., Brandt, M. J., Lazarevic, L. B., Legate, N., Christopherson, C., Wiggins, B., & Grahe, J. E. (2019). Publishing Research With Undergraduate Students via Replication Work: The Collaborative Replications and Education Project. *Frontiers in Psychology*, *10*, 247. https://doi.org/10.3389/fpsyg.2019.00247

Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, *566*(7744), 378–382. https://doi.org/10.1038/s41586-019-0941-9

Yarkoni, T. (2019). *The Generalizability Crisis* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/jqw35

Zakaria, S., Grant, J., & Luff, J. (2021). Fundamental challenges in assessing the impact of research infrastructure. *Health Research Policy and Systems*, *19*(1), 119. https://doi.org/10.1186/s12961-021-00769-z