The Oxford Martin Working Paper Series on
Technological and Economic Change

Disrupting Science

Chinchih Chen, Carl Benedikt Frey, and Giorgio Presidente

Working Paper No. 2022-4

Disclaimer: This is a working paper and represents research in progress. This paper represents the opinions of the authors and does not represent the position of the Oxford Martin School or other institutions or individuals.

For more information on the Oxford Martin Programme on Technological and Economic Change, please visit: https://www.oxfordmartin.ox.ac.uk/technological-economic-change/

For more information on the Oxford Martin Programme on the Future of Work, please visit: https://www.oxfordmartin.ox.ac.uk/future-of-work/

Disrupting Science

Chinchih Chen[1], Carl Benedikt Frey[1], and Giorgio Presidente[1]

[1]Oxford Martin School, University of Oxford

April 26, 2022

Abstract

In this paper, we explore how the rise of remote collaboration has shaped the pursuit of new ideas in scientific discovery. To systematically distinguish between disruptive breakthroughs and incremental discoveries, we use citations data for over ten million research teams—publishing in eleven fields of research from 1961 to 2020. On average, we document a robust and significant negative impact of remote collaboration on breakthrough discovery. However, beginning in the 2010s, the negative impact tapers off and even becomes positive. We provide evidence suggesting that the reversal is driven by improvements in key technologies needed for effective remote collaboration.

Keywords: science; remote work; innovation; teams; colocation

JEL classification: M1; O31; O43; O18;

1 Introduction

A large body of literature emphasizes the importance of face-to-face interactions for the cross-fertilisation of ideas (Marshall 1890, Jacobs 1969, Saxenian 1996, Moretti 2012, Jaffe et al. 1993, Glaeser & Gottlieb 2009, Helsley & Strange 2004, Duranton & Puga 2001). Despite this, however, the share of geographically-distributed teams in scientific research has steadily risen since the 1970s and accelerated further following the ICT revolution of the 1990s (Forman & Zeebroeck 2012). Over the same period, as shown in Figure 1, the percentage of fundamental or disruptive scientific discoveries has declined in tandem, implying that an ever-larger share of researchers focus on incrementally developing existing ideas, which without further breakthroughs eventually runs into diminishing returns (e.g. Ettlie et al. 1984, Dewar & Dutton 1986).

In this paper, we explore the link between the rise of remote collaboration and the increasingly incremental nature of scientific discovery. We do so against the background of an emerging literature suggesting that new ideas have become harder to find (e.g. Bloom et al. 2020, Park et al. 2021, Gordon 2017, Cowen 2011).[1] For example, applying the growth accounting of Solow (1957) to the production function for new ideas, Bloom et al. (2020) argue that sustained economic growth stems from the effective number of researchers and their research productivity. While the former has increased, they document that the latter has fallen in recent decades. Their findings, in other words, are consistent with a growing emphasis on the development of existing ideas at the expense of the discovery of new ones (Figure 1). And yet, the empirical literature has largely remained silent on the relationship between the striking organizational changes in scientific discovery in recent years and the

[1] In addition, Jones (2009, 2010) shows that the age at which inventors first patent as well as the size of research teams have increased, suggesting that longer stints of learning are required to gain the expertise necessary to push the technological frontier.

types of ideas and projects being pursued.[2]

To bridge this empirical void, we begin by identifying colocated and distributed teams based on authors' affiliation. Then, we classify ideas according to their novelty or disruptiveness. Thus, our analysis builds on the simple intuition that some ideas and scientific discoveries are more important than others (Gordon 2017, Wu et al. 2019, Mokyr 1992, Frey 2019, Acemoglu et al. 2020, Akcigit & Kerr 2018). For example, early breakthroughs in mRNA research opened entirely new vistas for scientific progress, making possible the most effective vaccines against Covid-19. The case of mRNA also illustrates the importance of face-to-face interactions for serendipitous discovery: Katlin Karikó and Drew Weissman, whose seminal 2005 paper demonstrated the promise of the technology, met randomly while taking turns on a Xerox machine. In contrast, BionTech's mRNA vaccine update in response to the Omicron variant is an example of an incremental adjustment of an existing idea.

To systematically distinguish between disruptive and incremental discoveries, we employ the methodology of Funk & Owen-Smith (2017) and Wu et al. (2019), who measure "disruptiveness" based on citation patterns. The underlying intuition goes as follows: if a paper replaces the corpus of research it cites in subsequent citations, it can be deemed to have disrupted the existing body of knowledge. To implement this approach, we use citations data for over ten million research teams—publishing in eleven fields of research from 1961 to 2020—from Microsoft Academic Graph (MAG), which also provides the affiliations of the authors on each paper. Crucially, this allows us to infer whether a team is colocated or distributed.

Using a difference-in-difference design, we exploit instances in which colocated teams

[2] Ever since Romer (1990), the production of ideas has become central to theories of modern growth, but the idea generation process has largely remained a black box for economists, though Weitzman (1998) and Jones (2009) are noteworthy exceptions.

become distributed and compare the disruptiveness of their publications before and after the split, controlling for team characteristics and those of their research field. Doing so, we find a robust and persistent negative impact of switching to a distributed model on disruption. Specifically, our baseline estimate suggests that a team is around 2 percentage points less likely to produce disruptive research after becoming distributed. On average, across the sixty years we study, the estimated probability of a team producing disruptive research is roughly 10 percent, implying that remote collaboration significantly reduces the chance of a team generating scientific breakthroughs.

However, the relationship between face-to-face collaboration and disruption is unlikely to be time invariant as coordination costs have fallen with recent advances in ICT. Indeed, exploring the relationship between geographical dispersion and disruption decade by decade, we find evidence that the above described pattern reversed in recent years, suggesting that the benefits of colocation may have declined over time (Hellmanzik & Kuld 2021, Head et al. 2019, Kim et al. 2009).[3] For example, in the 2010s, the key technologies needed for effective remote collaboration had begun to proliferate, including video conferencing technology and the cloud. While online communication tools, like Skype and Dropbox, began to spread in the 2000s, it is only after 2010 that many key remote work technologies emerged, including Trello/Office 356 (2011); Zoom/Google Drive (2012); Slack (2013); Overleaf (2014); Microsoft Teams (2017). Consistent with the hypothesis that improvements in remote work technology are responsible for the reversal in the relationship between distribution and disruption, we provide evidence that colocation matters less when the split involves team members in countries with better broadband infrastructure.

[3] Some studies attribute this to the decline in travel costs (Catalini et al. 2020, Dong et al. 2018, Agrawal et al. 2017). In this paper, we focus on the role of communication technologies, like Forman & Zeebroeck (2012) and Furman et al. (2021).

Our paper adds to several literatures. First, a growing body of work examines the economic effects of remote work. In a pioneering study of the Chinese travel agency Ctrip, Bloom et al. (2015) found that the switch to working from home increased performance by 13 percent through fewer breaks, sick days, and a more convenient working environment. These findings have also been vindicated in more recent experiments. For example, studies of a Fortune 500 retailer show that when on-site workers took up opportunities to go remote in 2018-2019, their productivity rose by 7 percent (Emanuel & Harrington 2020). But while remote work might increase productivity in routine and repetitive activities, for which knowledge spillovers are less important, it might inhibit the generation of new ideas, as distributed teams are more siloed (Yang et al. 2021). To the best of our knowledge, we are the first to study the impact of remote work on creativity in scientific discovery.

Second, we build on a series of papers on scientific collaboration (Agrawal et al. 2017, Dubois et al. 2014, Azoulay et al. 2010),[4] and a subset of this literature exploring the impact of technology on the organization of research (e.g. Agrawal & Goldfarb 2008, Forman & Zeebroeck 2012, Ding et al. 2010).[5] As argued by Becker & Murphy (1992), collaborative work has historically been hampered by coordination costs that increase with geographic dispersion. Yet while the dramatic reduction in the cost of coordinating production at distance—brought by the proliferation of the Internet in the 1990s—led to predictions about "the death of distance" (Cairncross 1997, Friedman 2005), virtual interactions have complemented co-location rather than substituting for it (Agrawal & Goldfarb 2008, Glaeser 1998).[6] Co-location remains particularly important in serendipitous discov-

[4] Scientific research is becoming increasingly collaborative (Catalini et al. 2020, Wuchty et al. 2007), allowing for greater specialization and to access complementary talent (Haeussler & Sauermann 2020) as well as more combinations of skills and knowledge (Freeman et al. 2015). Increasing complexity has also led to the rise of interdisciplinary teams (Falk-Krzesinski et al. 2011, Milojević 2014).

[5] For instance, Freeman et al. (2015) show that remote collaboration has a heterogeneous impact on research productivity in physics, nanoscience, biotechnology and applied microbiology.

[6] In fact, most evidence suggests that agglomeration matters more since the digital age in a variety of settings (Forman & Zeebroeck 2012, Furman et al. 2021, Cairncross 1997, Agrawal & Goldfarb 2008, Berger

ery (Catalini 2018), while communication that involves tacit knowledge (Polanyi 1958, Von Hippel 1994) still benefits from face-to-face interactions (Gaspar & Glaeser 1998, Rosenthal & Strange 2001, Storper & Venables 2004).[7] Thus, Catalini et al. (2020) find that a reduction in travel costs between locations led to an increase in collaboration. In contrast, we show that the decline in communication costs has shaped the *type* of ideas being pursued, leading to a relative increase in incremental discovery.[8]

Third, our paper relates to an ongoing debate surrounding our future growth prospects following the productivity slowdown of the mid-2000s (Byrne et al. 2016, Syverson 2017, Goldin et al. 2020, Vollrath 2020). While Gordon (2017) argues that the great inventions of the late nineteenth century cannot be repeated, Brynjolfsson et al. (2021) point out that productivity growth tends to follow a J-curve as complementary investments are required to realize the benefits from new technologies. While we find that the rise of remote collaboration—driven by the ICT revolution (Forman & Zeebroeck 2012)—led to a secular decline in the disruptiveness of science over the past half century, we also find that this relationship was reversed after 2010 as complementary investments were made in remote work technologies.[9] This rhymes with evidence from the historical record, showing that harnessing the benefits of electricity and steam power also required complementary investments and reorganization, leading to delayed productivity gains (David 1990, Crafts 2004).

The remainder of this paper is structured as follows. Section 2 describes our approach for identifying distributed teams, as well as measuring the disruptiveness of scientific papers. In section 3, we outline our empirical strategy, which we take to the data in section 4.

& Frey 2016, Leamer & Levinsohn 1995, Berger & Frey 2017, Blum & Goldfarb 2006, Agrawal et al. 2015).

[7] In addition, Glaeser & Ponzetto (2010) argue that improvements in ICT have increased the returns to new ideas.

[8] Though we note a reversal in recent years.

[9] That the disruptiveness of research has declined is corroborated by the evidence in Park et al. (2021).

Section 5 discusses the potential role of ICT in shaping the impact of remote collaboration on the disruptiveness of research output over time. Finally, in section 6, we outline our conclusions.

Figure 1: The negative correlation between the production of disruptive ideas and geographical distribution.

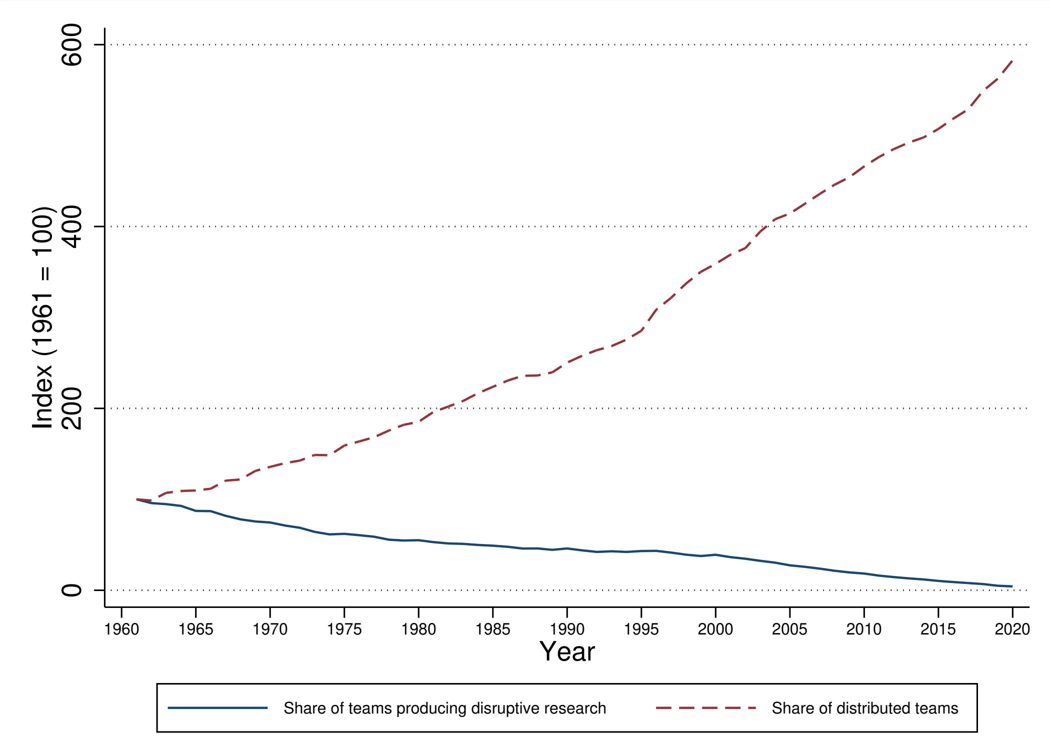

The figure presents the evolution of the shares of disruptive papers (solid line) and spatially distributed teams (dashed line), both normalized to 100 in 1961. Disruptive papers are defined as those with a value of the disruption index larger than the 90th percentile of the index distribution. Spatial distributed are teams in which at least one of the team members is affiliated with a different institution.

2 Data and Measurement

For our analysis, we use Microsoft Academic Graph (MAG), which contains information on 87 million scientific articles, including authors names, affiliations, research fields, and citation relationships between articles. In the below, we describe how this data is used to define research teams (section 2.1), measure the disruptiveness of publications (section 2.2), and create our final panel dataset (section 2.3). We also discuss potential measurement errors and strategies to address them (section 2.4).

Since our extensive cleaning procedure leaves us with relatively few observations before the 1960s, we start our analysis in 1961—the first year for which we have more than 10,000 teams. Our final sample consists of an unbalanced panel of 12,876,720 research teams publishing papers in eleven scientific fields over sixty years—from 1961 to 2020. Further key summary statistics are provided in Appendix Table I.

2.1 Colocated *vs* Spatially Distributed Teams

We extract the authors' name and affiliation from each publication in MAG, and assign a unique identifier to each research team, defined as a group of more than one author publishing at least one paper together. By definition, teams are always composed by the same set of authors. If in a given year, one team adds or loses an author, we consider it to be a different team. Since we focus on teams, we drop 34% of single-author publications from MAG, which leaves us with roughly sixty million papers to analyse.

To infer whether a team is colocated or distributed, we rely on the affiliation of each coauthor. We deem a team to be "colocated" if all authors are affiliated with the same research institution, making it more likely that the collaboration involves a higher degree of face-to-face interaction relative to electronic or digital communication. Conversely, if at least one coauthor is affiliated with a different institution, we deem the team to be "distributed".

A challenge with this approach is that 12% of teams have authors with multiple affiliations. For example, if the first team member has affiliation A, while the second member is affiliated with A and B, how can we infer whether a team is distributed or not? To maximise sample size, in case of multiple institutions, we use the first listed institutions. However, our results remain the same if we restrict our analysis to teams composed by authors with a single affiliation. After excluding articles with missing authors' affiliation or typos, we are left with some 22 million publications.

2.2 Measuring Disruption

While citation counts are often used to evaluate the importance of an invention or scientific discovery, they are at best a crude measure of the importance of an idea. For example, even though studies show that articles by large research teams receive more citations (Wuchty et al. 2007, Klug & Bagrow 2016), smaller teams are more likely to make breakthrough discoveries that disrupt the status quo (Wu et al. 2019). Consider two famous articles. One is the Bak–Tang–Wiesenfeld (BTW) model, which was the first discovered example of a dynamical system displaying self-organized criticality (Bak et al. 1988). The other examines Bose–Einstein condensation in a gas of sodium atoms (Davis et al. 1995). Both articles have a similar number of citations, but while the BTW-model is typically cited without mentions of references from that article, the Bose–Einstein condensation article is almost always co-cited with the earlier work of Bose (1924) and Einstein (1924). In other words, the BTW model launched new streams of research, whereas the Bose–Einstein article built on previous work in a cumulative manner, like most science.

As noted, science that opens up new vistas for progress is critical for long-run growth since incremental improvements eventually run into diminishing returns. Yet this crucial difference is rarely adequately reflected in the empirical literature. A recent notable exception is a paper by Funk & Owen-Smith (2017) which, to the best of our knowledge, developed the first index to systematically distinguish between incremental and disruptive ideas, focusing on patented inventions.[10] Another is Wu et al. (2019), which applied the same approach to scientific papers and software projects. In addition, they provided a comprehensive validation of the measure, showing that Nobel-prize-winning papers, on average, score among the 2% most disruptive articles, while review articles tend to register at the bottom end of the distribution. They also validate the disruptiveness index based on expert surveys across a host of disciplines.[11]

Following Wu et al. (2019), we construct a directed network in which papers are nodes and citations are links between them using the MAG. The simple idea underlying our approach is that papers which are co-cited with the earlier work it cites incrementally improve an existing body of knowledge, whereas papers that are not co-cited with the papers they cite create new paradigms. Specifically, based on our network, we calculate three quantities: i) the number of papers citing a "focal paper", but not the papers cited by it (n_a); ii) the number of papers citing both the focal paper and at least one of its references (n_b); and iii) the number of papers citing at least one of the references of the focal paper, but not the focal paper itself (n_c). Formally, the disruption index D is defined as:

$$\tilde{D} \equiv \frac{n_a - n_b}{n_a + n_b + n_c}$$

Thus, our index varies between -1 and 1, with values close to 1 reflecting greater novelty or disruptiveness. Conversely, values close to -1 imply incremental or cumulative

[10] However, Funk & Owen-Smith (2017) refer to disruptive ideas as "destabilizing", and incremental ideas as "consolidating." We adopt language that is more widely used in the literature, including by Wu et al. (2019), who uses the same measure.

[11] For details, see Wu et al. (2019).

improvements of the existing body of knowledge. Reassuringly, the Bose-Einstein condensation paper registers at the lower end of the $\tilde{D}$ distribution, despite its high number of citations. More broadly, as shown in Figure I of the online appendix, there is a U-shaped relationship between $\tilde{D}$ and the number of citations received by a paper. This implies that both articles with radical ideas as well as with marginal contributions (in terms of disruptiveness) may receive a high number of citations, making citation counts a noisy measure for our purposes.

Using data from the MAG, we apply the methodology above to calculate the disruptiveness of each publication p by team i in year t, $\tilde{D}_{p,i,t}$. By construction, however, our index can only be calculated for 72% of articles with at least one citation, leaving us with approximately 16 million articles.

2.3 Panel Construction

To aggregate $D_{p,i,t}$ at the team- and year-level, which is the variation we exploit in our empirical analysis (see section 3), we experiment with different approaches. Our baseline approach is taking the simple average of $\tilde{D}_{p,i,t}$ for each team-year pair. Specifically, letting $N_{i,t}$ be the number of publications by team i in year t, we compute the following quantity:

$$\overline{\tilde{D}}_{i,t} = \frac{1}{N_{i,t}} \sum_{N_{i,t}} \tilde{D}_{p,i,t}$$

We then construct a dummy variable $D_{i,t} = 1$ if in a given year, $\overline{\tilde{D}}_{i,t} > 90$th percentile of sample distribution and zero otherwise. This approach has the advantage of making the interpretation of our results more transparent in terms of probability of disruption. One potential issue with this approach, however, is that average disruptiveness might be low because of a high number of publications ($N_{i,t}$ is high). Thus, we assign $D_{i,t} = 1$ if the team i publishes at least one disruptive paper in year t, where disruption is again evaluated in terms of the overall underlying distribution of the index.

We note that both approaches rely on a general definition of disruptiveness, which is based on the full distribution of the index across years. However, in evaluating the importance of ideas, we might need to take into account the fact that some publication might be disruptive *relative* to a particular scientific discipline or time period. To do so, we exploit the fact that articles in MAG are classified into research fields. One potential issue, for example, is that as disciplines have evolved and knowledge has become increasingly specialised over the years, meaning that papers on similar topics might be classified in different fields depending on the time period. To mitigate related concerns, we use eleven major fields, which allows us to consistently evaluate whether a publication is disruptive within its respective discipline, as well as within a discipline and year. Appendix Table II presents the relative incidence of research areas in our sample. Section 4 shows that our results are consistent regardless of the approach we use to quantify disruptiveness.

2.4 Measurement Errors

A crucial assumption underpinning our empirical approach is that the affiliation authors provide on published papers accurately reflects their geographical location during the research project. This is likely to be true in the vast majority of cases. Consider a colocated team starting a project at time p. Suppose that one of the members changes institution at $s > p$. Let t be the time of publication. If $p \leq t < s$, we have no measurement issues: affiliations accurately reflects the location of all team members while conducting their research.

It must be noted, however, that the longer the time from the generation of an idea to publication (ITP), $t - p$, the higher the possibility that a team switches status between the beginning of a project and its publication ($t > s$). This could generate measurement errors, because we would infer from the affiliations that the team is distributed, even though members might have collaborated onsite during the early stages of the project (or vice versa). Nonetheless, our approach is reasonable if affiliations reflect the institution where the bulk of the research was conducted, even when the ITP is long.[12] Yet there are challenging cases in which authors use their current affiliation, regardless of where they worked on the project, and this will be more problematic in fields where publication cycles are lengthy.[13] To deal with this issue, at least in part, we drop social sciences and humanities, where the first response time as well as the review duration is particularly long (Huisman & Smits 2017), making measurement errors less likely.[14]

A related worry is that the ITP is affected by the time required to obtain results. For instance, some research requires clinical trials, which can result in lengthy time spans between the beginning of the project and publication. While we cannot observe this aspect directly in our data, we perform robustness checks in which we exclude disciplines like medicine, where clinical trials and the collection of experimental data is common.

Another concern is that in some cases, scholars might be visiting other departments while remaining affiliated with their institution. This is more likely to happen when research requires special equipment (i.e., a particle accelerator) onsite, leading scholars to work in a different location from their affiliated institution. To that end, we perform robustness tests in which we restrict our analysis to research fields that are less likely to require special equipment, using data from the National Science Foundation Survey of Research and Development Expenditures at Universities and Colleges.[15] Crucially, for our purposes, the questionnaire asks all U.S. research institutions to report R&D expenditures

[12] This would mean reporting the old affiliation whenever $s - p > t - s$.
[13] For instance, they might do so to improve departmental evaluations based on published papers.
[14] However, this reduces our sample by 26%.
[15] The repositories can be found at https://www.nsf.gov/statistics/herd/pub_data.cfm.

for the purchase of capitalized equipment by research field. Appendix Figure II reports the share of equipment expenditures by research field for all decades spanning our dataset.[16] We note that in the early years of our sample, life sciences dominate in terms of spending on capital equipment. In more recent decades, for which we have more granular data, medicine and biology stand out together with engineering.

3 Empirical Strategy

For our empirical analysis, we relate the dummy for disruptive research, $D_{i,t}$, to whether a team is colocated or distributed using the following linear model:

$$D_{i,t} = \beta_0 + \beta_1 DISTR_{i,t} + \nu_{i,t} \tag{1}$$

We define a dummy variable $DISTR_{i,t} = 1$ if in a given year, a team is distributed and zero otherwise. Model (1) can be estimated with OLS, provided that the $DISTR_{i,t}$ is uncorrelated to the error term $\nu_{i,t}$, i.e. $E[DISTR_{i,t}\nu_{i,t}] = 0$. However, this condition might not be satisfied in general, because $DISTR_{i,t}$ is not randomly assigned across teams. For instance, research institutions with the largest resources might host expensive lab equipment (e.g., a radiation or magnetic field laboratory). In this case, researchers are more likely to belong to that institution (i.e. $DISTR_{i,t} = 0$), while at the same time having better chances of developing disruptive ideas. A related concern is that $DISTR_{i,t}$ depends on team ability. For instance, teams composed by particularly talented scholars might be more likely to produce disruptive research while at the same time being more likely to be distributed across different institutions due to more plentiful job opportunities among other things.

[16] In particular, we use survey waves for the years 2020, 2010, 2000, 1990, 1980, and 1972, which is the first wave available.

To mitigate such concerns, we exploit the fact that 13% of teams publish in multiple years, which allows us to estimate the within-team impact of a switch between onsite and remote collaboration. Hence, we impose the following error structure:

$$\nu_{i,t} \equiv u_{f,t} + u_i + \varepsilon_{i,t} \tag{2}$$

where $u_{f,t}$ is the research field-year fixed effect; u_i is the team fixed effect, and $\varepsilon_{i,t}$ an idiosyncratic error term. Crucially, team fixed effects absorb team-specific, time-invariant unobserved characteristics and allows us to estimate the within-team change in disruptiveness before and after a team splits.[17] We notice that over 70% of teams changing their mode of collaboration, move from colocated to distributed.

Research field-year fixed effects purge our estimates from a number of potential additional sources of bias. First, $\tilde{D}$ (and so D) might be biased towards zero since the most recent publications have had less time to accumulate citations (see section 2.2).[18] Second, as noted in section 2.1, there is an unobserved component of TIP—the time from the generation of an idea to the first journal submission—which might affect our measurement of whether a team is colocated or distributed, introducing measurement errors. By including field-year fixed effects, we are able to account for such systematic differences in the unobserved component of TIP. Third, researchers in specific fields might be more likely to generate disruptive research, while at the same time being more or less prone to work remotely. The inclusion of field-year fixed effects purge our estimates of β_1 from such potential confounding factors.

Once we account for persistent differences in ability across teams and control for field-specific factors and with the error structure in (2), the variables underlying the realisation

[17] For example, team fixed effect absorbs the impact of different team sizes, which previous research finds to be associated with the disruptiveness of ideas (Wu et al. 2019).

[18] Few citations are reflected in a large n_c, which increases the denominator of $\tilde{D}$.

of $DISTR_{i,t}$ are unlikely to be correlated with the error term. Variation in $DISTR_{i,t}$ might be driven, for instance, by authors' personal circumstances, such as whether their partners get a job in a different city.

4 Can Distributed Teams Disrupt?

We begin by estimating model (1) with and without fixed effect, as to exploit the whole sample. Unless differently stated, we cluster standard errors at the team-level. Our main results are presented in Table 1. In column 1, we compare the probability of disruption between colocated and distributed teams. The coefficient is negative and significant at the 99 percent confidence level. The estimated magnitude implies that colocated teams are 2 percentage points more likely to produce disruptive research. The estimated difference is substantial compared to the average probability of conducting disruptive research, which is captured by the constant term and is equal to less than 11 percent.

In column 2, we include team fixed effect and exploit within-team changes in the mode of collaboration. The coefficient of interest is again negative and significant at the 99 percent confidence level. The estimated magnitude implies that a team is 0.6 percentage points less likely to produce disruptive research when it switches from being colocated to distributed. While the average team-level impact of becoming a geographically dispersed team is small in absolute terms, it is still economically meaningful compared to the average probability of conducting disruptive research.[19]

However, the negative impact might only be transitory, as team members adapt to the organizational change. To examine this possibility, we run an event-study regression in the spirit of Goodman-Bacon (2021).[20] The event of interest is the switch from being colocated

[19] Appendix Table III shows that the results very similar if we drop teams composed by authors with multiple affiliations (see section 2.1).

[20] As in (1), we include team and research-year fixed effects. We cluster the errors at the team-level. The

to spatially distributed. We implement the event study with four leads and four lags.[21] Figure 2 presents our results. After teams split geographically, the probability of producing disruptive research appears to decline persistently by more than 2 percentage points in the fourth year after the switch. This leads us to conclude that spatial distribution entails a substantial and long-lasting negative impact on the probability of producing disruptive research across our full sample.

Table 1: Spatial distribution and research disruptiveness.

	(1) Disruptive research	(2) Disruptive research
Distributed team	-0.020***	-0.006***
	(0.0001)	(0.001)
Constant	0.108***	0.102***
	(0.0001)	(0.0003)
Observations	13,878,717	1,772,910
R-squared	0.094	0.582
Team FE	no	yes
Field-year FE	yes	yes

The table presents OLS estimates of the relationship between spatial distribution and disruptiveness. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

4.1 Robustness

Naturally, a key concern with our approach is reverse causality. In particular, we are concerned that $DISTR_{i,t}$ might be affected by past values of $D_{i,t}$—i.e., teams producing disruptive research could be more or less likely to become distributed. To that end, in column

complete estimation results can be found in appendix Table IV.

[21] Additional lags and leads are not statistically significant due to the limited number of observations per team.

Figure 2: Distributed teams and disruption: results from an event-study regression.

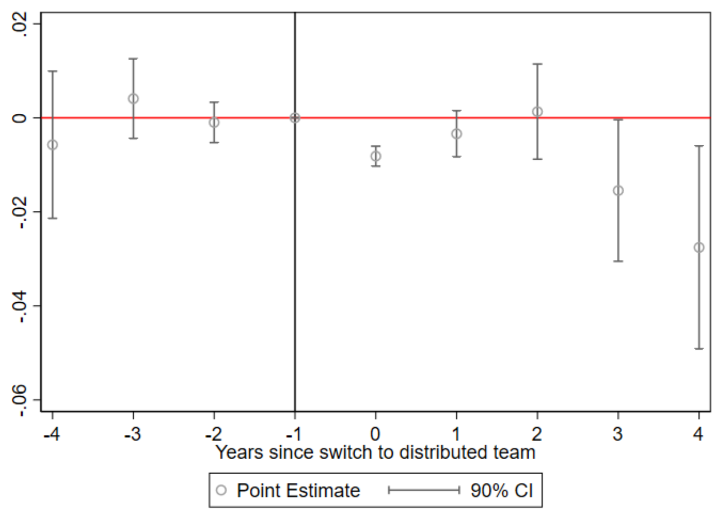

The figure presents the estimated coefficient and 90 percent confidence intervals from an event-study regression with four year leads and four year lags. The dependent variable is the a dummy variable equal to 1 if the disruption index is above the 90th percentile of the index distribution. The treatment variable is a dummy taking value 1 if in one year, at least one of the team members is affiliated with a different institution. The regression includes team and research-year fixed effects. Standard errors are clustered at the team-level.

1 of appendix Table V, we regress $DISTR_{i,t}$ on two years lags of $D_{i,t}$. Reassuringly, none of the coefficients are statistically significant.

Still, it is possible that disruptiveness drives the geographical dispersion of teams. For instance, the working paper version of an article could receive wide recognition and thus affect the job market prospects of the team members. In this case, future values of $D_{i,t}$ (when the paper is finally published) could affect whether the team has become colocated or distributed. While we partly address this by focusing on research fields where the publication process is relatively short, we further explore this possibility in column 2, where we regress $DISTR_{i,t}$ on two year leads of $D_{i,t}$. Also in this case, we do not find any statistically significant coefficient. In other words, expected disruptiveness seems unlikely to drive our results.

In addition, to mitigate concerns about the impact of measurement errors, Table VI presents results based on three sub-samples. In column 1, we assess the robustness of our results by dropping medicine, which has potentially long ITPs due to factors such as clinical trials. In column 2, we only consider fields that should not be reliant on clinical data, such as computer science, mathematics, and engineering. In column 3, we exclude equipment-intensive disciplines (based on appendix Figure II), including medicine, biology, and engineering. In all four samples, the coefficient of spatial distribution is at least as large as in the baseline and statistically significant at the 99 percent statistical level.

Another concern is that geographic dispersion as well as the disruptiveness of a team's output is a function of the length of collaboration. In column 1 of appendix Table VII, we regress $DISTR_{i,t}$ on the cumulative number of publishing years for a team. The positive and significant coefficient suggests that the longer a team is active, the higher the probability of some of its members changing institution. The estimated coefficient implies that every additional year a team appears in our data increases the probability that the team is distributed by 2 percentage points.

We note that the impact of the length of collaboration on disruption might be either negative or positive. For instance, it might take time for newly-formed teams to collaborate effectively, which would have a negative impact on disruptiveness. At the same time, new teams—that is, teams appearing for the first time in our data—might develop more novel ideas and then turn to exploiting a breakthrough discovery through marginal improvements in follow-on publications. If this is the case, we would expect a negative correlation between the length of collaboration and disruptiveness. We find support for the latter: column 2 of Table VII shows that the cumulative number of publishing years has a negative and significant impact on disruptiveness, suggesting that established teams primarily exploit existing ideas. Since the length of collaboration might drive our results, we include it as a control in column 3. Reassuringly, the coefficient of distributed teams remains similar in size and it is still significant at the 99 percent confidence level.

A remaining concern is whether our findings are robust to different definitions of what constitutes a distributed team. So far, we have defined distributed teams as those where one member is at a different institution to the rest of the group. But should we consider a team to be distributed if only one of its five members collaborates at distance? To explore this, for all distributed teams, we compute the share of different institutions hosting each member. We then use this share to weight estimates from (1), assigning full weight to colocated teams. The results are presented in column 1 of Table VIII, which shows that the coefficient and the standard errors are very similar to the baseline. In column 2, we present an alternative test in which we replace $DISTR_{i,t}$ with a dummy equal to 1 only if more than one-third of a team's members belong to different institutions and zero otherwise. We label these teams "highly distributed". The coefficient is again very similar to the baseline and significant at the 99 percent confidence level.

In a similar vein, we check that our main results are robust to alternative definitions of disruption. For example, in column 1 we obtain qualitatively identical results when using a dummy equal to 1 for publications above the 99th percentile of the index distribution. Next, we turn to experiment with a dummy equal to 1 if a team publishes *at least one* disruptive paper. The results are presented in column 2, in which the coefficient is still negative and statistically significant. Finally, in column 3, we assign $D_{i,t} = 1$ if the average disruptiveness of a team is larger than the 90th percentile of the index distribution within the respective research field in each year. Again, the coefficient remains negative and significant at the 99 percent confidence level.

In addition, results are similar when we use other indicators for valuable ideas—though these indicators are silent on whether a publication constitutes a breakthrough. Appendix

Table X shows the estimated impact of spatial distribution on total citations (column 1) and the number of papers published (column 2). Both variables are expressed in logs. The negative and significant coefficient in column 1 implies that distributed teams receive 6 percent fewer citations after the switch. Column 2 shows that there is not sufficient variation in the number of papers per team to identify the coefficient. This is due to the fact that most teams publish only one paper per year together. Therefore, in column 3 we remove team the fixed-effect and estimate that distributed teams publish 7 percent less papers than colocated ones.

Finally, we replace our baseline linear probability model with an alternative estimator. Since approximately 90 percent of total observations take the value 0—which includes teams that never produce disruptive research, as well as teams that consistently produce disruptive publications over all years of observation—we implement a fixed effect Logit model.[22] The results are presented in appendix Table XI. The coefficient is still negative and statistically significant at the 99 percent confidence level, but it is now much larger in absolute value. The estimated magnitude implies that a team is 15 percentage points less likely to produce disruptive research when it switches from being colocated to distributed. While we still prefer our baseline specification with the full set of fixed effects and clustered errors, the results from the FE Logit estimator suggest that the OLS coefficients might be biased toward zero, and so constitutes a lower bound.

[22] With this specification, we experience convergence problems related to the concavity of the log pseudo-likelihood function. We solve this issue by including year fixed effects and research field-year trends instead. The statistical software we use does not support clustering of standard errors at the team-level with the fixed effect Logit. As an alternative, we implement a wild bootstrapping procedure with 50 repetitions.

5 Mechanisms

We next turn to exploring the mechanisms underlying Table 1, which suggests that on average, spatial distribution is detrimental to breakthrough discovery. One possible explanation is that collaborative work is hampered by coordination costs that increase with geographic dispersion, as argued by Becker & Murphy (1992). This would imply that improvements in ICT infrastructure, which allows for more frictionless communication, should mitigate the negative effect of spatial distribution. To test this, we interact $DISTR_{i,t}$ with the average number of broadband subscriptions per capita in the countries hosting the coauthors' institutions. In our sample, authors are affiliated with institutions located in 186 countries.[23]

One concern is that the technological trends driving the change in broadband subscriptions might be common to all countries. Thus, we account for the potential serial correlation of errors by imposing a more restrictive two-way clustering structure, by team and year. The results from this exercise are presented in Table 2.[24] We note that the coefficient of the main term is negative and significant as before. However, the interaction coefficient is positive and significant, suggesting that the negative impact of spacial distribution is mitigated when people in the countries where academics locate have better access to digital infrastructure. Unsurprisingly, the coefficient on broadband per capita is also statistically significant, implying that the Internet matters also for colocated teams as they benefit from better access to information and more frictionless communication with researchers outside their respective teams and departments.

Clearly, information and communications technology has improved enormously in recent years and decades, suggesting that the relationship between distribution and disrup-

[23] Roughly 30% of articles are published by teams with at least one member belonging to a research institution in the United States. Around 15% of teams have authors based in Germany, Spain, France, and Italy; 5% from the United Kingdom; 10% from Canada, Australia, and Japan and 10% China.

[24] The number of observations is lower due to missing broadband data for some countries.

Table 2: The role of ICT infrastructure.

	(1) Disruptive research
Distributed team	-0.010***
	(0.003)
Distributed team x broadband per capita	0.016*
	(0.010)
Broadband per capita	0.067**
	(0.027)
Observations	948,807
R-squared	0.580
Team FE	yes
Field-year FE	yes

The table presents OLS estimates of the relationship between spatial distribution, disruptiveness and ICT infrastructure. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Broadband per capital is the team-level average number of internet subscriptions per capita, obtained by taking the average country-specific number of internet subscriptions per capita in the countries hosting each team member. Standard errors are clustered at the team- and year-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

tion is unlikely to be time invariant. For example, while email and the Web arrived in the 1990s, Skype and Dropbox were invented in the 2000s. In addition, the 2010s saw the proliferation Slack, Google Drive, Zoom, and MS Teams. Thus, the negative correlation between spatial distribution and disruptiveness might be more muted in recent years.

To explore this intuition in greater detail, we augment model (1) with interaction terms between $DISTR_{i,t}$ and dummy variables flagging the years before 1990, between 1990 and 1999, 2000 and 2009, 2010 and 2020, and finally between 2015 and 2020.[25] Figure 3 plots the related point estimates and 90 percent confidence intervals of the relationship between disruptiveness and spatial distribution by time period. The impact of spatial distribution and disruption is negative and significant up to 2010. However, the coefficient

[25] These results are similar with alternative temporal splits of the sample.

turns positive and not statistically significant for the years after 2010, and positive and significant after 2015. In other words, the negative relationship between spatial distribution and disruptiveness appears more muted over the final decade of the sample, and has even reversed in recent years.

Figure 3: The impact of spatial distribution on disruption in different time periods.

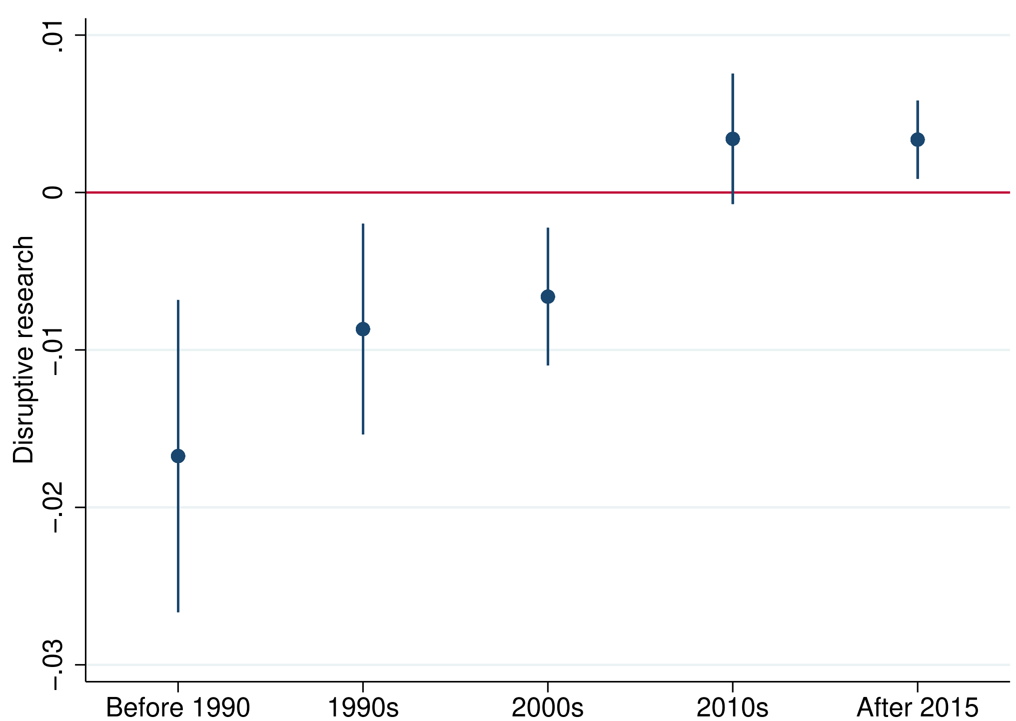

The figure presents OLS estimates and 90% confidence intervals of the relationship between spatial distribution and disruptiveness over different time periods. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level.

The reversal after 2010 is surprising since even the latest remote work technologies are imperfect substitutes for face-to-face interactions. For example, as Mills (2017) and Glaeser (1998) have pointed out, digital technologies only allow for interactions that are planned on at least one end, making sporadic encounters, which are highly important for innovation (Jacobs 1969, Andrews 2019), unlikely to happen. On the other hand, there are potentially countervailing forces that might outweigh such costs, including complementary knowledge, skills, and resources between different institutions. Thus, improved remote work technologies, which reduce communication frictions, allow team members to better take advantage of such complementarities. For example, local networks of individual members in their respective institutions might suddenly benefit the wider distributed team to a larger extent.

To that end, we explore how the effect of switching from onsite to distributed work differs depending on the quality of the institution a researcher moves to. If a team member moves to a more highly ranked institution, for example, it is more likely to benefit from local knowledge spillovers or complementary resources. Specifically, we rely on the Times Higher Education (THE) university ranking to assign a score to each university in our sample.[26] For each team and year, we compute the average team ranking score based on the affiliated institution of each team member.[27] We then create a dummy variable equal to 1 if the average team ranking score improves with respect to the previous year observed in the data. We note that the number of observations is lower in this specification since only 1,200 institutions are ranked by THE—while our sample includes a total of 19,928 institutions—and we lose one observation per team, since we need one period lag to define the ranking dummy.

We next compute the average ranking of a team in each year and interact $DISTR_{i,t}$ with a dummy equal to 1 if the average ranking of a team increases and 0 otherwise.

[26] The ranking is available at https://www.timeshighereducation.com/world-university-rankings.

[27] We create a categorical variable $Rank = 1$ if the institution is a top-50, $Rank = 2$ if ranked between 51 and 100, and so on up to $Rank = 4$. Then, constrained by the structure of the THE ranking, we assign $Rank = 5$ if the institution is ranked between 200 and 300, and so on up to $Rank = 11$, which corresponds to institutions ranked below the 1200th position, which are no longer ranked. Alternative specifications, such as finer thresholds between positions 1 and 200, deliver very similar results.

We also include a triple interaction by multiplying the latter with a dummy taking the value 1 from 2010. Table 3 presents our results from this exercise. As expected, we find that spatial distribution harms disruptiveness irrespective of whether team members joined better-ranked universities before 2010, pointing towards communication frictions.[28] Conversely, the triple interaction, which considers team members moving to more highly ranked institutions after 2010, is positive and significant. We take this to imply that before the last generation of remote work technologies, distributed teams were unable to harness the benefits of local knowledge spillovers and complementary assets due to communication frictions within the team. For example, if one team member moves to a better research institution, they will most likely benefit from feedback and sporadic interactions with their local faculty. But their wider team has only benefited from such input in recent years as communication frictions within distributed teams have been diminished.

6 Conclusions

The past half century has seen a striking increase in remote collaboration in scientific discovery as the cost of travel and electronic communication has fallen (Catalini et al. 2020, Dong et al. 2018, Agrawal et al. 2017, Forman & Zeebroeck 2012, Furman et al. 2021). In this paper, we explore how the rise of remote collaboration has shaped the trajectory of science between 1961 and 2020. Building on the intuition that some ideas are more fundamental than others (Mokyr 1992, Christensen 2013, Gordon 2017), we implement a novel approach to distinguish between disruptive breakthroughs and incremental improvements of existing ideas (Wu et al. 2019, Funk & Owen-Smith 2017), using citations data from Microsoft Academic Graph (MAG). Exploiting instances in which colocated teams be-

[28] Clearly, this does not imply that individual authors do not benefit from joining a better institution as we are only considering the disruptiveness of the ongoing collaborations.

Table 3: The role of complementary knowledge, skills, and resources across different institutions.

	(1) Disruptive research
Distributed team	-0.00971***
	(0.00229)
Better university ranking	-0.00815
	(0.00967)
Distributed team x better university ranking	-0.00484
	(0.0130)
Distributed team x better university ranking x years since 2010	0.0279***
	(0.0105)
Observations	812,205
R-squared	0.582
Team FE	yes
Field-year FE	yes

The table presents OLS estimates of the relationship between spatial distribution, disruptiveness and the average university ranking of the team. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Better university ranking is a dummy variable equal to 1 if the average team ranking score improves with respect to the previous year it is observed. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

come distributed, we use a difference-in-difference design to compare the disruptiveness of research outputs before and after teams split. Doing so, we document a significant negative impact of geographic dispersion on disruptive discovery. Our findings speak to those of Yang et al. (2021), showing that remote teams are more siloed, which they argue is detrimental to innovation. The relative increase in incremental science also speaks to the findings of Bloom et al. (2015), suggesting that new ideas have become harder to find. To the best of our knowledge, we provide the first evidence linking the decline in research productivity to the reorganization of scientific discovery.

At the same time, the relationship between face-to-face collaboration and disruption has not been monotonic. When we turn to exploring the relationship between geographical dispersion and disruption decade by decade, we find evidence that the negative effect of remote collaboration has become more muted over time. It even reversed in recent years, suggesting that the benefits of colocation have declined as information and communications technology has progressed (Hellmanzik & Kuld 2021, Head et al. 2019, Kim et al. 2009). In particular, the 2010s witnessed some striking improvements in the technologies needed for effective remote collaboration, including video conferencing technology and the cloud. Consistent with the idea that improvements in remote work technology explain the reversal in the relationship between distribution and disruption, we find colocation to be less important when the split involves team members in countries with better broadband infrastructure.

Our results have implications for the ongoing debate surrounding the future of economic growth (Gordon 2017, Mokyr 2014). As Brynjolfsson et al. (2021) has point out, productivity growth tends to follow a J-curve as complementary investments and organizational changes are required to realize the benefits of new technologies. In a similar fashion, our findings suggest that harnessing the benefits of the ICT revolution for remote collaboration required complementary investments in technologies that support remote work. As the Covid-19 pandemic has sparked a sharp increase in patenting related to remote work technologies (Bloom et al. 2021), this could mark the beginning of a revival of disruptive science and faster productivity growth.

That said, our findings should not be taken to suggest that face-to-face interactions no longer matter. In particular, digital technologies still only allow for interactions that are planned on at least one end (Glaeser 1998), meaning that the kind of sporadic encounters that have driven innovation historically are still relevant today (Frey 2020). Indeed, we find a beneficial effect of of switching to remote but joining a more highly ranked institution after 2010. We interpret the finding as suggesting that the distributed team is more likely to benefit from local knowledge spillovers, presumably due to fewer communication frictions within the team. This implies that local knowledge networks and digital networks are complements rather than substitutes. Indeed, by connecting local knowledge networks, remote collaboration increases the size of what Muthukrishna & Henrich (2016) have called the "collective brain".

References

Acemoglu, D., Akcigit, U. & Celik, M. A. (2020), 'Radical and incremental innovation: The roles of firms, managers and innovators', *AEJ Macroecon (forthcoming)* .

Agrawal, A., Catalini, C. & Goldfarb, A. (2015), 'Crowdfunding: Geography, social networks, and the timing of investment decisions', *Journal of Economics & Management Strategy* **24**(2), 253–274.

Agrawal, A. & Goldfarb, A. (2008), 'Restructuring research: Communication costs and the democratization of university innovation', *American Economic Review* **98**(4), 1578–90.

Agrawal, A., McHale, J. & Oettl, A. (2017), 'How stars matter: Recruiting and peer effects in evolutionary biology', *Research Policy* **46**(4), 853–867.

Akcigit, U. & Kerr, W. R. (2018), 'Growth through heterogeneous innovations', *Journal of Political Economy* **126**(4), 1374–1443.

Andrews, M. (2019), 'Bar talk: Informal social interactions, alcohol prohibition, and invention', *Unpublished Working Paper* .

Azoulay, P., Graff Zivin, J. S. & Wang, J. (2010), 'Superstar extinction', *The Quarterly Journal of Economics* **125**(2), 549–589.

Bak, P., Tang, C. & Wiesenfeld, K. (1988), 'Self-organized criticality', *Physical review A* **38**(1), 364.

Becker, G. S. & Murphy, K. M. (1992), 'The division of labor, coordination costs, and knowledge', *The Quarterly Journal of Economics* **107**(4), 1137–1160.

Berger, T. & Frey, C. B. (2016), 'Did the computer revolution shift the fortunes of us cities?

technology shocks and the geography of new jobs', *Regional Science and Urban Economics* **57**, 38–45.

Berger, T. & Frey, C. B. (2017), 'Industrial renewal in the 21st century: evidence from us cities', *Regional Studies* **51**(3), 404–413.

Bloom, N., Davis, S. J. & Zhestkova, Y. (2021), Covid-19 shifted patent applications toward technologies that support working from home, *in* 'AEA Papers and Proceedings', Vol. 111, pp. 263–66.

Bloom, N., Jones, C. I., Van Reenen, J. & Webb, M. (2020), 'Are ideas getting harder to find?', *American Economic Review* **110**(4), 1104–44.

Bloom, N., Liang, J., Roberts, J. & Ying, Z. J. (2015), 'Does working from home work? evidence from a chinese experiment', *The Quarterly Journal of Economics* **130**(1), 165–218.

Blum, B. S. & Goldfarb, A. (2006), 'Does the internet defy the law of gravity?', *Journal of international economics* **70**(2), 384–405.

Bose, S. N. (1924), *Plancks gesetz und lichtquantenhypothese*, Springer.

Brynjolfsson, E., Rock, D. & Syverson, C. (2021), 'The productivity j-curve: How intangibles complement general purpose technologies', *American Economic Journal: Macroeconomics* **13**(1), 333–72.

Byrne, D. M., Fernald, J. G. & Reinsdorf, M. B. (2016), 'Does the united states have a productivity slowdown or a measurement problem?', *Brookings Papers on Economic Activity* **2016**(1), 109–182.

Cairncross, F. (1997), *The death of distance: how the communications revolution will change our lives*, Harvard Business School.

Catalini, C. (2018), 'Microgeography and the direction of inventive activity', *Management Science* **64**(9), 4348–4364.

Catalini, C., Fons-Rosen, C. & Gaulé, P. (2020), 'How do travel costs shape collaboration?', *Management Science* **66**(8), 3340–3360.

Christensen, C. M. (2013), *The innovator's dilemma: when new technologies cause great firms to fail*, Harvard Business Review Press.

Cowen, T. (2011), *The great stagnation: How America ate all the low-hanging fruit of modern history, got sick, and will (eventually) feel better: A Penguin eSpecial from Dutton*, Penguin.

Crafts, N. (2004), 'Steam as a general purpose technology: a growth accounting perspective', *The Economic Journal* **114**(495), 338–351.

David, P. A. (1990), 'The dynamo and the computer: an historical perspective on the modern productivity paradox', *The American Economic Review* **80**(2), 355–361.

Davis, K. B., Mewes, M.-O., Andrews, M. R., van Druten, N. J., Durfee, D. S., Kurn, D. & Ketterle, W. (1995), 'Bose-einstein condensation in a gas of sodium atoms', *Physical review letters* **75**(22), 3969.

Dewar, R. D. & Dutton, J. E. (1986), 'The adoption of radical and incremental innovations: An empirical analysis', *Management science* **32**(11), 1422–1433.

Ding, W. W., Levin, S. G., Stephan, P. E. & Winkler, A. E. (2010), 'The impact of information technology on academic scientists' productivity and collaboration patterns', *Management Science* **56**(9), 1439–1461.

Dong, X., Zheng, S. & Kahn, M. E. (2018), The role of transportation speed in facilitating high skilled teamwork, Technical report, National Bureau of Economic Research.

Dubois, P., Rochet, J.-C. & Schlenker, J.-M. (2014), 'Productivity and mobility in academic research: evidence from mathematicians', *Scientometrics* **98**(3), 1669–1701.

Duranton, G. & Puga, D. (2001), 'Nursery cities: urban diversity, process innovation, and the life cycle of products', *American Economic Review* **91**(5), 1454–1477.

Einstein, A. (1924), 'Quantentheorie des einatomigen idealen gases', *SB Preuss. Akad. Wiss. phys.-math. Klasse* .

Emanuel, N. & Harrington, E. (2020), Working remotely? selection, treatment and the market provision remote work", Technical report, Harvard mimeo.

Ettlie, J. E., Bridges, W. P. & O'keefe, R. D. (1984), 'Organization strategy and structural differences for radical versus incremental innovation', *Management science* **30**(6), 682–695.

Falk-Krzesinski, H. J., Contractor, N., Fiore, S. M., Hall, K. L., Kane, C., Keyton, J., Klein, J. T., Spring, B., Stokols, D. & Trochim, W. (2011), 'Mapping a research agenda for the science of team science', *Research Evaluation* **20**(2), 145–158.

Forman, C. & Zeebroeck, N. v. (2012), 'From wires to partners: How the internet has fostered R&D collaborations within firms', *Management science* **58**(8), 1549–1568.

Freeman, R. B., Ganguli, I. & Murciano-Goroff, R. (2015), *Why and Wherefore of Increased Scientific Collaboration*, University of Chicago Press.

Frey, C. B. (2019), *The technology trap*, Princeton University Press.

Frey, C. B. (2020), 'The great innovation deceleration', *MIT Sloan Management Review* **61**(4), 1–3.

Friedman, T. L. (2005), *The world is flat: A brief history of the twenty-first century*, Macmillan.

Funk, R. J. & Owen-Smith, J. (2017), 'A dynamic network measure of technological change', *Management Science* **63**(3), 791–817.

Furman, J. L., Nagler, M. & Watzinger, M. (2021), 'Disclosure and subsequent innovation: Evidence from the patent depository library program', *American Economic Journal: Economic Policy* **13**(4), 239–70.

Gaspar, J. & Glaeser, E. L. (1998), 'Information technology and the future of cities', *Journal of urban economics* **43**(1), 136–156.

Glaeser, E. L. (1998), 'Are cities dying?', *Journal of economic perspectives* **12**(2), 139–160.

Glaeser, E. L. & Gottlieb, J. D. (2009), 'The wealth of cities: Agglomeration economies and spatial equilibrium in the united states', *Journal of economic literature* **47**(4), 983–1028.

Glaeser, E. L. & Ponzetto, G. A. (2010), *Did the Death of Distance Hurt Detroit and Help New York?*, University of Chicago Press.

Goldin, I., Koutroumpis, P., Lafond, F. & Winkler, J. (2020), 'Why is productivity slowing down?', *Oxford Martin Working Paper Series on Economic and Technological Change Working Paper No. 2021-6* .

Goodman-Bacon, A. (2021), 'Difference-in-differences with variation in treatment timing', *Journal of Econometrics* .

Gordon, R. J. (2017), *The rise and fall of American growth*, Princeton University Press.

Haeussler, C. & Sauermann, H. (2020), 'Division of labor in collaborative knowledge production: The role of team size and interdisciplinarity', *Research Policy* **49**(6), 103987.

Head, K., Li, Y. A. & Minondo, A. (2019), 'Geography, ties, and knowledge flows: Evidence from citations in mathematics', *Review of Economics and Statistics* **101**(4), 713–727.

Hellmanzik, C. & Kuld, L. (2021), 'No place like home: geography and culture in the dissemination of economic research articles', *Empirical Economics* **61**(1), 201–229.

Helsley, R. W. & Strange, W. C. (2004), 'Knowledge barter in cities', *Journal of Urban Economics* **56**(2), 327–345.

Huisman, J. & Smits, J. (2017), 'Duration and quality of the peer review process: the author's perspective', *Scientometrics* **113**(1), 633–650.

Jacobs, J. (1969), 'Strategies for helping cities', *The American Economic Review* pp. 652–656.

Jaffe, A. B., Trajtenberg, M. & Henderson, R. (1993), 'Geographic localization of knowledge spillovers as evidenced by patent citations', *the Quarterly journal of Economics* **108**(3), 577–598.

Jones, B. F. (2009), 'The burden of knowledge and the "death of the renaissance man": Is innovation getting harder?', *The Review of Economic Studies* **76**(1), 283–317.

Jones, B. F. (2010), 'Age and great invention', *The Review of Economics and Statistics* **92**(1), 1–14.

Kim, E. H., Morse, A. & Zingales, L. (2009), 'Are elite universities losing their competitive edge?', *Journal of Financial Economics* **93**(3), 353–381.

Klug, M. & Bagrow, J. P. (2016), ‘Understanding the group dynamics and success of teams’, *Royal Society open science* **3**(4), 160007.

Leamer, E. E. & Levinsohn, J. (1995), ‘International trade theory: the evidence’, *Handbook of international economics* **3**, 1339–1394.

Marshall, A. (1890), ‘Principles of economics’, *London: Macmillan* pp. 1–627.

Mills, E. S. (2017), Sectoral clustering and metropolitan development, *in* ‘Sources of metropolitan growth’, Routledge, pp. 3–18.

Milojević, S. (2014), ‘Principles of scientific research team formation and evolution’, *Proceedings of the National Academy of Sciences* **111**(11), 3984–3989.

Mokyr, J. (1992), *The lever of riches: Technological creativity and economic progress*, Oxford University Press.

Mokyr, J. (2014), ‘Secular stagnation? not in your life’, *Secular Stagnation: Facts, Causes and Cures* **83**.

Moretti, E. (2012), *The new geography of jobs*, Houghton Mifflin Harcourt.

Muthukrishna, M. & Henrich, J. (2016), ‘Innovation in the collective brain’, *Philosophical Transactions of the Royal Society B: Biological Sciences* **371**(1690), 20150192.

Park, M., Leahey, E. & Funk, R. (2021), ‘The decline of disruptive science and technology’, *arXiv preprint arXiv:2106.11184* .

Polanyi, M. (1958), *Personal knowledge*, Routledge.

Romer, P. M. (1990), ‘Endogenous technological change’, *Journal of political Economy* **98**(5, Part 2), S71–S102.

Rosenthal, S. S. & Strange, W. C. (2001), 'The determinants of agglomeration', *Journal of urban economics* **50**(2), 191–229.

Saxenian, A. (1996), 'Inside-out: regional networks and industrial adaptation in silicon valley and route 128', *Cityscape* pp. 41–60.

Solow, R. M. (1957), 'Technical change and the aggregate production function', *The review of Economics and Statistics* pp. 312–320.

Storper, M. & Venables, A. J. (2004), Buzz: the economic force of the city, *in* 'Journal of Economic Geography', Citeseer.

Syverson, C. (2017), 'Challenges to mismeasurement explanations for the us productivity slowdown', *Journal of Economic Perspectives* **31**(2), 165–86.

Vollrath, D. (2020), *Fully Grown*, University of Chicago Press.

Von Hippel, E. (1994), '"sticky information" and the locus of problem solving: implications for innovation', *Management science* **40**(4), 429–439.

Weitzman, M. L. (1998), 'Recombinant growth', *The Quarterly Journal of Economics* **113**(2), 331–360.

Wu, L., Wang, D. & Evans, J. A. (2019), 'Large teams develop and small teams disrupt science and technology', *Nature* **566**(7744), 378–382.

Wuchty, S., Jones, B. F. & Uzzi, B. (2007), 'The increasing dominance of teams in production of knowledge', *Science* **316**(5827), 1036–1039.

Yang, L., Holtz, D., Jaffe, S., Suri, S., Sinha, S., Weston, J., Joyce, C., Shah, N., Sherman, K., Hecht, B. et al. (2021), 'The effects of remote work on collaboration among information workers', *Nature human behaviour* pp. 1–12.

Online Appendix (not for publication)

Table I: Summary statistics.

	(1) N	(2) mean	(3) sd	(4) min	(5) max
Disruptiveness index	1.388e+07	0.00264	0.0467	-0.904	1.000
Dummy disruptiveness index > 90th percentile	1.388e+07	0.100	0.300	0	1
Dummy disruptiveness index > 99th percentile	1.388e+07	0.0102	0.100	0	1
Citations per team	1.388e+07	24.91	88.65	1	81,953
Number of papers per team	1.388e+07	1.284	1.088	1	127
Share of distributed teams	1.388e+07	0.423	0.494	0	1
Dummy disruptiveness (distributed teams)	5.863e+06	0.0694	0	0.0694	0.0694
Dummy disruptiveness (colocated teams)	8.001e+06	0.122	0	0.122	0.122
Broadband subscriptions per capita	7.178e+06	0.167	0.149	0	0.617

Notes: This table presents summary statistics of the main variables used in our analysis.

Figure I: Correlation between citation counts and disruptiveness.

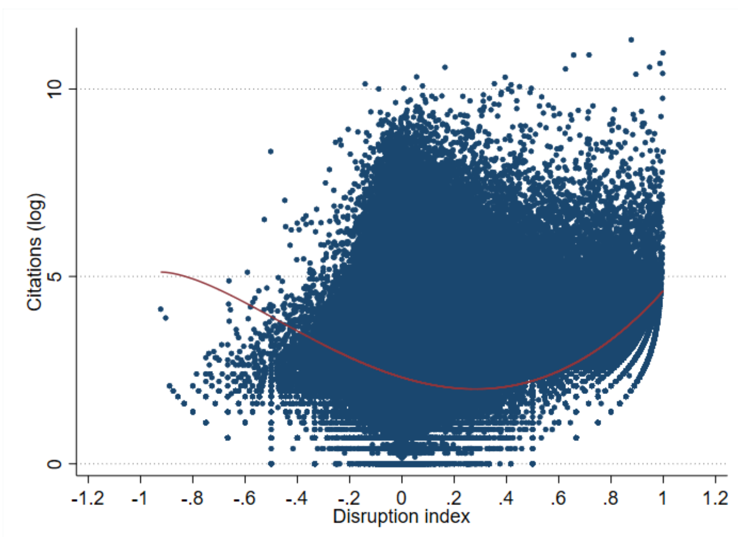

Notes: This figure presents the team-year-level correlation between the disruption index and log-citations.

Table II: Incidence of research fields.

Research field	Sample incidence
Art	0.1%
Biology	22.2%
Business	0.31%
Chemistry	19.9%
Computer science	3.33%
Economics	2.67%
Engineering	3.03%
Environmental science	0.18%
Geography	0.18%
Geology	2.67%
History	0.03%
Materials science	6.85%
Mathematics	5.81%
Medicine	18.89%
Philosophy	0.03%
Physics	7.2%
Political science	0.10%
Psychology	5.78%
Sociology	0.87%

Source: Authors' own calculations based on MAG data.

Figure II: Reliance of research fields on physical capital equipment, by decade.

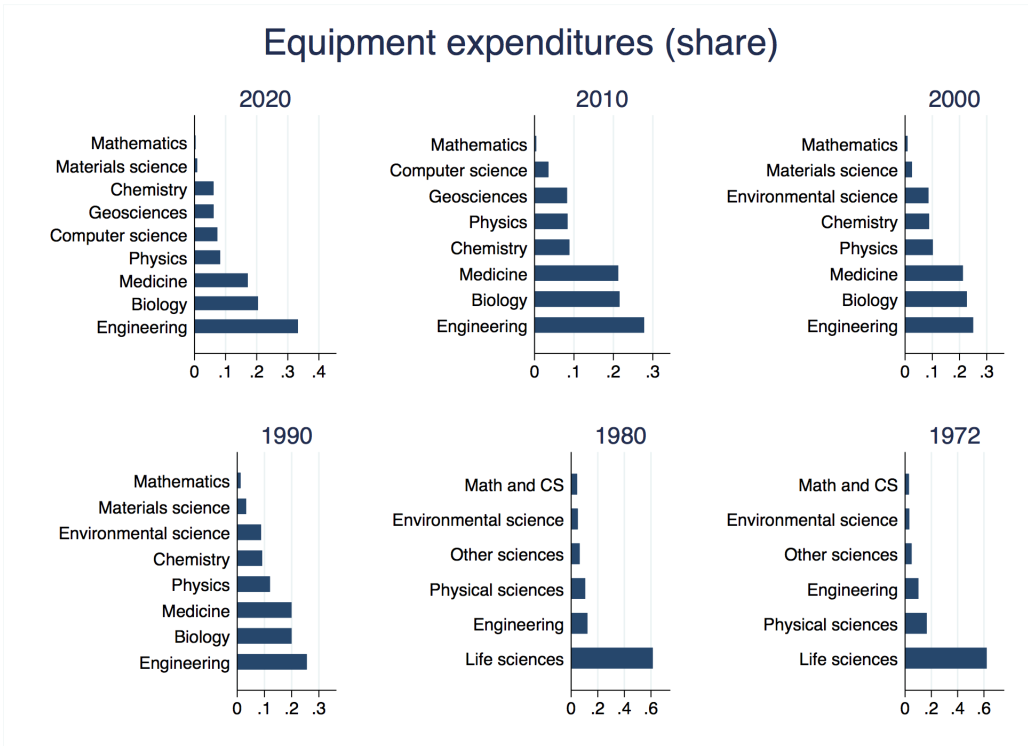

Notes: This figure presents the share of expenditures on capitalised capital equipment by research field and decade. The data refers to research institutions in the United States.
Source: Authors' calculation based on the National Science Foundation Survey of Research and Development Expenditures at Universities and Colleges.

Table III: Spatial distribution and research disruptiveness: dropping teams with multi-affiliation authors.

	(1) Disruptive research	(2) Disruptive research
Distributed team	-0.016***	-0.007***
	(0.000)	(0.002)
Constant	0.114***	0.109***
	(0.000)	(0.000)
Observations	12,159,097	1,550,198
R-squared	0.091	0.584
Team FE	no	yes
Field-year FE	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution and disruptiveness. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table IV: Event study regression.

	(1) Disruptive research
lead4	-0.006
	(0.009)
lead3	0.003
	(0.005)
lead2	0.000
	(0.003)
lag0	-0.009***
	(0.001)
lag1	-0.004
	(0.003)
lag2	0.001
	(0.006)
lag3	-0.021**
	(0.009)
lag4	-0.027**
	(0.013)
Constant	0.100***
	(0.000)
Observations	1,772,910
R-squared	0.582
Team FE	yes
Field-year FE	yes

Notes: This table presents the estimated coefficient and standard errors from an event-study regression with four year leads and four year lags. The dependent variable is a dummy variable equal to 1 if the disruption index is above the 90th percentile of the index distribution. The treatment variable is a dummy taking value 1 if in one year, at least one of the team members is affiliated with a different institution. The regression includes team and research-year fixed effects. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table V: Disruptiveness impact on mode of collaboration: lags and leads.

	(1) Distributed team	(2) Distributed team
Disruptive research (t-1)	0.005	
	(0.006)	
Disruptive research (t-2)	0.003	
	(0.006)	
Disruptive research (t+1)		-0.003
		(0.006)
Disruptive research (t+2)		-0.001
		(0.007)
Observations	27,577	27,580
R-squared	0.905	0.883
Team FE	yes	yes
Field-year FE	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution and disruptiveness. The dependent variable is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Disruptive research is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table VI: Short ITP.

	(1) Excluding medicine	(2) Non-clinical trials research fields	(3) Non-equipment-intensive research fields
Distributed team	-0.006***	-0.011***	-0.007***
	(0.001)	(0.003)	(0.002)
Observations	1,542,457	329,524	1,043,636
R-squared	0.585	0.575	0.583
Team FE	yes	yes	yes
Field-year FE	yes	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution and disruptiveness. Non-clinical trials research fields are: mathematics, computer science, engineering; while non-equipment-intensive research fields include: excluding medicine, biology, engineering. The dependent variable is a dummy taking the value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table VII: Length of collaboration.

	(1) Distributed team	(2) Disruptive research	(3) Disruptive research
Length of collaboration	0.021***	-0.009***	-0.009***
	(0.0001)	(0.0001)	(0.0001)
Distributed team			-0.004***
			(0.001)
Observations	1,772,910	1,772,910	1,772,910
R-squared	0.875	0.582	0.582
Team FE	yes	yes	yes
Field-year FE	yes	yes	yes

The table presents OLS estimates of the relationship between spatial distribution, length of collaboration and disruptiveness. The dependent variable in column 1 is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. The dependent variable in columns 2 and 3 is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table VIII: Highly distributed teams.

	(1) Disruptive research	(2) Disruptive research
Distributed team	-0.006***	
	(0.001)	
Strongly distributed team		-0.006***
		(0.001)
Observations	1,772,897	1,772,910
R-squared	0.587	0.582
Weights	yes	no
Team FE	yes	yes
Field-year FE	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution and disruptiveness. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index from Wu et al. (2019). Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Weights are the shares of members in different institutions; highly distributed teams have at least 2/3 of their members in different institutions. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table IX: Alternative disruption measures.

	(1) Disruption (dummy p99)	(2) Disruption (at least 1 paper)	(3) Disruption (within field-year)
Distributed team	-0.001**	-0.007***	-0.006***
	(0.0004)	(0.001)	(0.001)
Observations	1,772,910	1,772,910	1,772,910
R-squared	0.494	0.583	0.533
Team FE	yes	yes	yes
Field-year FE	yes	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution and disruptiveness. The dependent variable in column (1) is a dummy taking value 1 if the average team disruption is above the 99th percentile of the distribution of the disruptiveness index; in column (2), a dummy equal to 1 if in a given year, a team produces at least one paper above the 99th percentile of the distribution of the disruptiveness index; and in column (3), a dummy equal to 1 if the average disruptiveness of a team is above the 90th percentile of the index distribution within a research field and year. Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table X: Alternative research productivity indicators.

	(1) Citations	(2) Number of papers	(3) Number of papers
Distributed team	-0.057***	-0.000	-0.071***
	(0.004)	(0.000)	(0.000)
Observations	1,772,910	1,772,910	13,878,717
R-squared	0.727	1.000	0.036
Team FE	yes	yes	no
Field-year FE	yes	yes	yes

Notes: This table presents OLS estimates of the relationship between spatial distribution, disruptiveness and alternative research productivity indicators. All dependent variables are in logs. Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.

Table XI: Alternative estimator.

	(1) Disruptive research FE Logit
Distributed team	-0.152***
	(0.025)
Observations	291,832
Number of teamid	115,967
Team FE	yes
Field-year FE	no
Field-year trend	yes
Year FE	yes

Notes: This table presents Fixed Effect Logit estimates of the relationship between spatial distribution and disruptiveness. The dependent variable is a dummy taking value 1 if the average team disruption is above the 90th percentile of the distribution of the disruptiveness index. Distributed team is a dummy variable equal to 1 if in a given year, a team is spatially distributed (based on authors' affiliation) and 0 otherwise. Standard errors are clustered at the team-level. The coefficients with *** are significant at the 1% level, with ** are significant at the 5% level, and with * are significant at the 10% level.