4-13-2020
Working Paper Number 20007

The Case for Remote Work

Matthew Clancy
Iowa State University, mclancy@iastate.edu

The Case for Remote Work

Abstract
The case for remote work goes well beyond its use during the covid-19 global pandemic. Over the last ten years, research from a variety of subdisciplines in economics and other social sciences collectively makes a strong case for the viability of remote work for the long-run. This paper brings this research together to argue remote work (also called telework) is likely to become far more common in the future for four reasons. First, the productivity of individual workers who switch to remote work is comparable or higher than their colocated peers, at least in some industries. Second, matching firms to geographically distant workers is becoming easier thanks to technological and social developments. Third, remote workers tend to be cheaper because workers value geographic flexibility and the ability to work remotely. Fourth, the benefits of knowledge spillovers from being physically close to other knowledge workers has been falling and may no longer exist in many domains of knowledge. While the prevalence of remote work (pre-covid-19) is small, I show it was already rising rapidly with plenty of room to continue growing. Finally, I argue remote work has positive externalities and should be promoted by policy-makers.

Keywords
remote work, telework, knowledge spillovers, covid-19

Disciplines
Behavioral Economics | Labor Economics | Work, Economy and Organizations

The Case for Remote Work

Matthew Clancy

This Draft: April 13, 2020

The case for remote work goes well beyond its use during the covid-19 global pandemic. Over the last ten years, research from a variety of subdisciplines in economics and other social sciences collectively makes a strong case for the viability of remote work for the long-run. This paper brings this research together to argue remote work (also called telework) is likely to become far more common in the future for four reasons. First, the productivity of individual workers who switch to remote work is comparable or higher than their colocated peers, at least in some industries. Second, matching firms to geographically distant workers is becoming easier thanks to technological and social developments. Third, remote workers tend to be cheaper because workers value geographic flexibility and the ability to work remotely. Fourth, the benefits of knowledge spillovers from being physically close to other knowledge workers has been falling and may no longer exist in many domains of knowledge. While the prevalence of remote work (pre-covid-19) is small, I show it was already rising rapidly with plenty of room to continue growing. Finally, I argue remote work has positive externalities and should be promoted by policy-makers.

Thanks to Stian Westlake and Paul Hill for helpful comments, and Lyman Stone for sharing data. All errors are my own. *Matthew Clancy is an assistant teaching professor in the economics department at Iowa State University. Please email mclancy@iastate.edu with any comments and suggestions.*

As I am writing today, in the midst of the Covid-19 global pandemic, the case for short-term remote work is obvious. What other choices do organizations have to keep the lights on when their workers are isolating themselves to slow the spread of disease? I'm one of these people: Iowa State University, where I teach, has rapidly switched to remote instruction for the rest of the semester.

But I began writing this essay several months before the first cases of Covid-19 were reported in the popular press, and the ideas in it have been coalescing for more than a year. Covid-19 makes an obvious case for short-term remote work, also called telework. But I am arguing that the case for *long-term* remote work is much stronger than has typically been appreciated. Remote work does not have to be merely an emergency response, to be discarded when the pandemic subsides. It can be the new normal for many industries. Moreover, I think it *will* be the new normal and that this is desirable and should be encouraged. These reasons go far beyond maintaining business continuity during pandemics.

My arguments mostly rely on academic research published within the last ten years. Individually, any one of these articles is suggestive. Taken together I think they make a forceful case that remote work is going to become much more common in the years ahead. This work is split across multiple academic subdisciplines, and I have not seen it brought together in one place. That is the goal of this essay.

Before beginning, let's agree on some terms. By remote work, I mean individuals working physically distant from collaborating coworkers. In this essay, remote work can certainly mean working from home, but it can also mean working in a coffee shop, or a coworking space, or even in a satellite office. As long as the work is done mostly physically separated from collaborating coworkers, I call the work remote. In this essay, when collaborating coworkers are physically clustered together in the same physical workplace, I say the firm is colocated.

Let me also make clear what I am not claiming. I am not claiming that all industries will go remote - only more. I am not claiming that we will all work from home - only that this will become more common, alongside working from coworking spaces and satellite offices. I am not claiming that big cities will "disappear" - only that the choice to live in them will not be driven as much by the need to colocate near firms.

Most importantly, I am not claiming that remote collaboration is "better" than face-to-face communication. All else equal, collaborating face-to-face probably *is* better than collaborating remotely. But collaborating remotely offers its own advantages: firms that use a remote workforce can hire more productive workers for less.

What I am arguing in this essay is that technological and social changes have reached a tipping point and remote work will become more and more common going forward.

1. Theoretical Framework

In this paper, the framework for thinking about the decision to hire a remote worker is based on a simple model of a firm deciding whether to make a position remote or colocated. Firms prefer whichever offers a higher return, defined as the value the worker brings in divided by their wage. For economist readers, the value of either option is the marginal product of labor multiplied by the price of the firm's output *p*. For non-economist readers, that's simply the extra revenue the position brings in: marginal product of labor means the additional goods or services that the position lets the firm sell for price *p*.

In the model, a remote worker's marginal product is denoted by Q_r, which represents the expected "quality" of a new remote hire. The expected quality of a new colocated hire is denoted Q_c but their marginal product is *augmented* by two advantages of colocation. I use *A* to denote the increased productivity of an individual colocated worker, relative to a remote counterpart, and *S* to denote the increased productivity that stems from local knowledge spillovers. I'll discuss each of these in more detail below, but one reason I separate *A* and *S* because evidence related to each belongs to distinct academic literatures. Taken together, the marginal product of a colocated worker is the product ASQ_c, whereas the marginal product of a remote worker is simply Q_r.

We can think of *A* as measuring the importance of colocated traits of work (relative to remote analogues) such as: the relative ease of communicating ideas face-to-face or in a conference with a shared whiteboard; relatively better access to physical machines, artifacts, and customers; relatively stronger incentives to work hard when under the supervision of colocated managers and surrounded by colocated peers. If workers are more productive when colocated, then *A* is greater than 1. "*A*" falls when remote technology for collaboration improves.

The term *S* is a measure of the increase in productivity of colocated workers, relative to remote ones, due to "knowledge spillovers" from other colocated people. Think of this as measuring the productivity-enhancing exposure to new ideas and fresh perspectives from other geographically proximate workers (both in the firm and the city), as well as the increased ease of transferring this knowledge. If a worker's productivity is increased by being physically near other workers and their knowledge, then *S* is greater than 1. As it gets easier to access their ideas remotely, *S* falls.

Figure 1 illustrates how *A* and *S* might vary for a set of five example positions.

Figure 1. Plotting Positions in Terms of Relative Productivity (*A*) and Local Knowledge Spillovers (*S*)

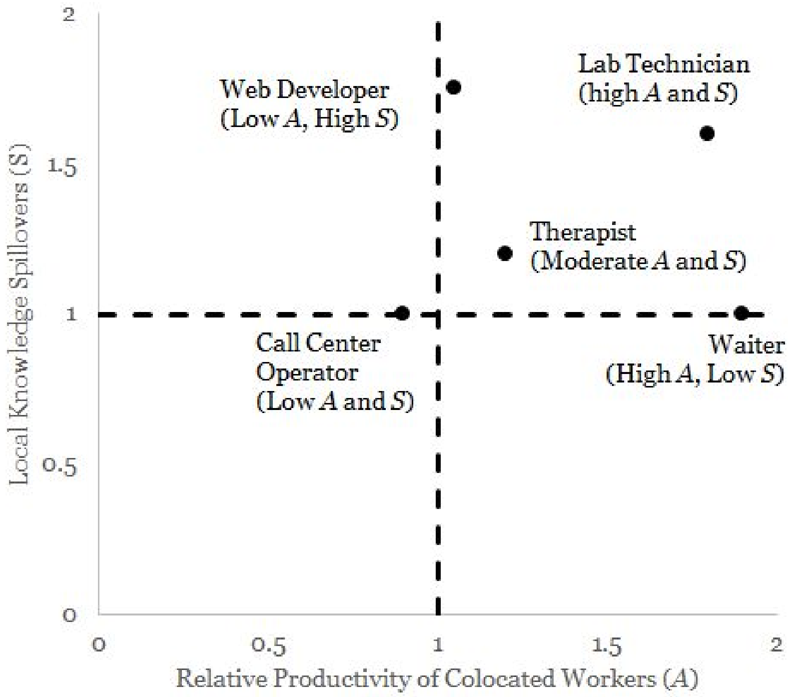

Note: This figure illustrates how the relative productivity advantage of colocation varies across five hypothetical positions.

In figure 1, the dashed lines indicate $A = 1$ and $S = 1$, which corresponds to the case where there are no advantages to a worker of a given quality being colocated. In the upper right quadrant, colocated workers of a given quality are more productive than remote ones. We would expect this to be the case for many positions, and I have given a few examples.

In the upper right corner of figure 1 is the position of lab technician, who needs to be physically present in the lab to perform their duties and also benefits significantly from being around other technicians who can share knowledge and ideas. The position cannot easily be done remotely, and this is represented by high values of *A* and *S*. Alternatively, the web developer position can perform their work tasks nearly as well at home as in the office (low *A*). However, the position benefits a lot from being near other developers, with whom they can share ideas and discuss problems. The presence of these strong knowledge spillovers is represented by a high *S*. In contrast, the waiter position does not benefit at all from knowledge spillovers ($S = 1$), but a worker needs to be physically present to perform their work tasks (high *A*).

In this example, a therapist has moderate values of *A* and *S*, indicating there is some real advantage to being physically present and colocated with other workers, but this advantage is not large. Lastly, a call center operator does not benefit from knowledge spillovers ($S = 1$), and is slightly more productive when working remotely than when working in an office ($A < 1$).

Looking at figure 1, clearly firms will prefer to make a position remote if it falls in the lower left quadrant, where A and S are less than 1 and remote workers more productive than colocated ones. But this is not the only case where remote workers are preferred.

In this paper, I focus on two additional advantages of remote work: (1) remote positions may attract better candidates and (2) remote positions may cost less. As noted above, expected worker quality is denoted by Q_i, where $i = r$ or c for remote and colocated respectively. Since remote positions can draw on a larger pool of candidates, it may be possible to recruit better workers in a remote position. This would be represented by the ratio $Q_r/Q_c > 1$.

I assume the cost of a worker is their wage w_i where $i = r$ or c for remote and colocated respectively. This may vary with the cost of living in different places, and if workers think of the option to work remotely as a benefit. If colocated workers are more expensive than remote ones, this can be represented by the ratio $w_c/w_r > 1$.

In this simple framework, firms will make a position remote (instead of colocated) if the return on a remote worker (the revenue they bring in divided by their cost) exceeds the return on a colocated worker. The return on a remote worker is pQ_r divided by the wage of a remote worker w_r. The return on a colocated worker is $pASQ_c$ divided by the wage of a remote worker w_c.. The condition that the return on a remote worker exceeds the return on a colocated one can be expressed with the following inequality:

$$pASQ_c/w_c \leq pQ_r/w_r$$

This equation can be rearranged as follows:

(1) $$A \times S \leq Q_r/Q_c \times w_c/w_r$$

This equation organizes the rest of this paper. It says firms will prefer to make a position remote if the productivity advantage of colocation (AS) is more than offset by access to higher quality remote workers (Q_r/Q_c) and/or higher costs of colocated labor (w_c/w_r).

For example, suppose that the higher cost of living in urban centers and a preference to work at home means the colocated wage is 30% higher than the comparable remote wage $(w_c/w_r = 1.3)$ and that access to a larger labor market means the expected quality of a remote worker is 15% greater than the local labor market $(Q_r/Q_c = 1.15)$. Since $1.3 \times 1.15 = 1.5$, this implies a colocated worker must be at least 1.5 times as productive as a remote one for a firm to prefer to make a position colocated $(A \times S > 1.5)$.

This example is illustrated in figure 2 below. All combinations of $AS < 1.5$ are shaded in blue. Firms prefer to make any position with A and S in this region a remote one. As we would expect, firms prefer to make all positions in the lower-left quadrant remote; this corresponds to positions where remote workers have higher productivity than colocated ones. However, firms also prefer to make some of the positions in the upper right quadrant remote. These are positions where colocated workers of the same quality are more productive than remote ones, but not so much more productive that it is worth forgoing the cost savings and higher quality of matches that a remote position enables.

Figure 2: Types of Positions where Remote Workers are Preferred

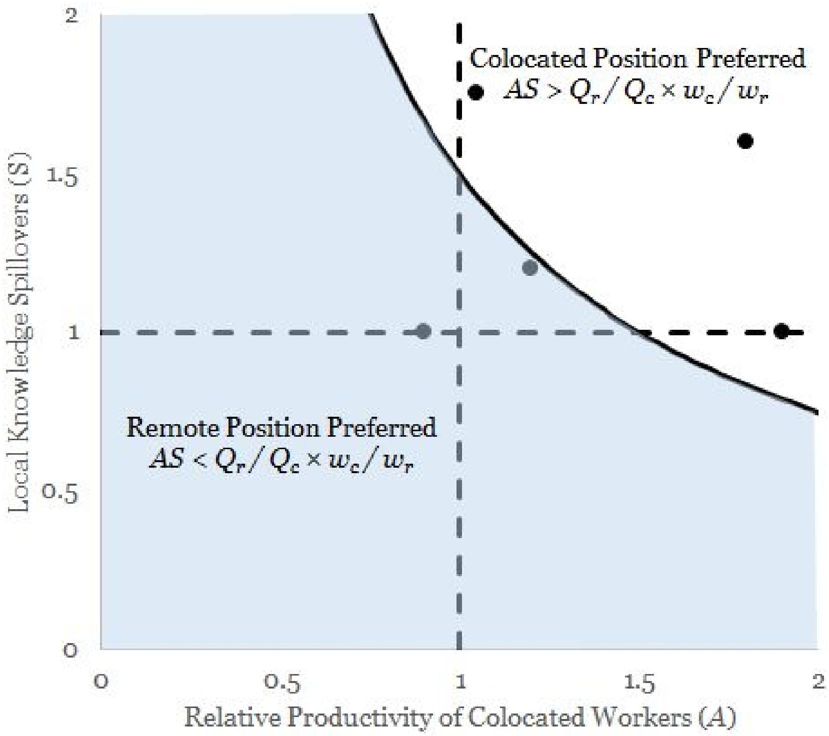

Notes: This figure illustrates the set of positions with values of A and S where remote workers are preferred, highlighted in blue, if $Q_r/Q_c \times w_c/w_r = 1.5$.

Equation 1 and figure 2 makes it clear that it is *not* necessary for the technology of remote work to be as good as colocated work ($A \leq 1$), nor for there to be no spillovers from colocation ($S \leq 1$), in order for firms to prefer to hire remotely. Note, for example, in figure 2 that a firm would prefer to make the therapist position remote, even though there are real advantages to colocation in the position.

The rest of this paper will present evidence that the left-hand side of equation (1) is not as large as commonly believed and shrinking; and the right-hand side of equation (1) is large and growing. Expressed in terms of figure 2, I will be arguing that positions are drifting down and to the left, so that positions that used to be better colocated (such as the web developer in Figure 1) are now better remote (i.e., they have or will move into the blue region in Figure 2). At the same time, I will argue that the quality and wage differentials are also shifting in a manner favorable to remote work. In terms of the figure, the blue region is expanding up and to the right.

Specifically, section 2 of this paper will review evidence on the productivity of remote versus colocated workers and argue A is close to or less than 1 in many cases. Section 3 will argue the expected quality of remote workers relative to colocated ones (Q_r/Q_c), is rising. Section 4 will present evidence that cost-of-living differences and a preference for remote work make the cost of remote positions less than the cost of colocated ones, and give reasons why the colocated wage premium may be rising. Section 5 will argue that the spillover term S is declining. Section 6 closes the main argument by presenting empirical evidence that even before the outbreak of covid-19, the prevalence of remote work was rising and had substantial scope to continue rising.

Explaining Recent Trends

Before moving on, in the remainder of this section I will show how this simple model can also explain why we have observed an increase in geographic concentration, even as the technology for remote work improves. Since at least 2005, the growth of employment in innovative industries has risen fastest in cities that had the largest innovative sector employment in 2005 (Atkinson, Muro, Whiton 2019). A simple way to model this is to assume the price p of the output of innovative industries has risen. Since the benefit of all workers (remote and colocated) is their marginal product multiplied by the price p, this raised the return on all workers, and firms in innovative industries created a large number of new positions. These could be remote or colocated.

Let us suppose that over the last several decades, improvements in communication technology have steadily reduced A (from very high starting levels). As equation (1) and figure 2 make clear, improvements in A do not necessarily lead to an increase in the prevalence of remote work. So long as $AS \geq Q_r/Q_c \times w_c/w_r$, firms will still choose to meet the increased demand for workers with colocated workers. And once they have decided to use colocated workers, they will obtain higher productivity from their workers if they locate in cities where spillovers are likely to be large, due to the presence of other innovative workers.

While changes in A may be mostly driven by improvements in communication and technology, the value of S are driven in part by the collective hiring decisions of all firms. The increased demand for knowledge workers, and the clustering of them into cities, will tend to increase S over time, since there will be more workers colocated in technology hubs to exchange ideas with. The rise in S can partially or totally offset declines in A over the same period. The result is increasing concentration, even as the technology for remote work steadily improves. This is consistent with what has happened over the last few decades.

But the model makes clear that this is just one possible equilibrium. Technology is likely to continue reducing A for the foreseeable future. Technological progress can also reduce S over time by enabling remote communication of ideas and knowledge. Meanwhile, the increase in S due to clustering tends to run into diminishing returns, since rising congestion costs eventually constrain the expansion of a city's population and density. At some point, parameters change enough so that equation (1) begins to hold. At that point, firms will switch to hiring remote, rather than colocated workers. Once this dynamic begins to hold, it has the long-run effect of lowering S, since workers no longer cluster together geographically. This further undermines the case for colocation, resulting in a new equilibrium where remote work is widespread.

2. Relative Productivity of Remote Work (A)

I first examine evidence related to the parameter A, which measures the relative increase in productivity for a worker of given quality, if they are colocated with other workers. Since this parameter changes with the state of information technology, I will emphasize more recent studies, when they are available.

Gajendran and Harrison (2007) is a good summary of the literature up through the mid-2000s. The paper highlights some of the main theoretical factors that may affect the productivity of remote and colocated workers. While it is typically assumed remote work is inferior to colocation ($A > 1$), there are a variety of reasons this may not be true. Most importantly, by (potentially) granting workers more flexibility in their time of work or their physical work environment, remote work can allow workers to better optimize for their own idiosyncratic preferences. Some workers may be more productive working early in the morning and others late at night; some are more productive in a silent enclosed space and others in a noisy coffee shop. By reducing stress about coordinating work schedules with family schedules, remote work can also allow greater focus while at work. Lastly, the reduction in travel time can extend the hours employees are able to work.

To assess the efficacy of remote work, Gajendran and Harrison (2007) collects 46 studies of remote work, all of which involved actual remote work rather than lab-based simulations. All of these studies, however, were observational, and therefore selection effects likely bias the results. Unsurprisingly, they find remote work is associated with greater perceptions of worker autonomy and job satisfaction. It's also associated with lower perceptions of work-family conflict, stress, and intent to seek new employment. For a smaller set of studies they also study the difference between occasional (less than 2.5 days per week) and intensive (more than 2.5 days per week) remote work. In general, the differences are not substantive, though they find evidence the quality of relationships with coworkers is worse for remote-workers working from home more often.

Gajendran and Harrison also find remote work has no statistically significant impact on work performance, as rated by the individual worker, and a positive effect when rated by the supervisor. While we will see these results are consistent with causally identified evidence, they should be interpreted with caution, since it may be that only high-performing workers are allowed to work remotely. This is also suggested by the positive association between perceived relationship quality with a supervisor and remote work, which Gajendran and Harrison also document.

Bercovitz and Feldman (2011) also studies the early era of remote work for the specific context of research collaborators who disclose inventions to their university technology transfer office. Over the period 1988-1999 they study 1,425 invention disclosures with multiple inventors. They find disclosures where some of the inventors are not local are no less likely to receive patents, and actually receive higher royalties for their inventions.

Similarly, Freeman, Ganguli, and Murciano-Goroff (2014) study the citations received by academic papers for a sample of 126,000 papers published between 1990 and 2010 in the fields of particle and field physics, nanoscience and nanotechnology, and biotechnology and applied microbiology. For each paper, they map the institution of each co-author to the associated city. They find papers with an international collaboration frequently receive fewer citations, conditional on the number of coauthors and references. However, papers by exclusively US coauthors, who are not colocated in the same city, receive more citations in the physics and biology fields studied, but less in the nanotechnology fields.

Turning from academic collaboration to firms, Monteiro, Straume, and Velento (2019) provide recent evidence about the efficacy of remote work for a broad sample of firms. Monteiro, Straume, and Velento (2019) use a 2011-2016 survey of a representative sample of 5,000-6,500 Portuguese firms to assess the impact of remote work on firm-level productivity. Lacking a precise measure of remote work, they use as a proxy a positive response to the question "Did your enterprise provide to the persons employed remote access to the enterprise's e-mail system, documents, and applications?" In some specifications, they also interact this binary response with the share of employees who report using a computer in their work, to derive a measure of the *potential* extent of remote work.

Simply comparing the productivity of firms that enable remote access to those that do not (and including some basic controls), remote-accessible firms are 15% more productive. But this appears to be driven by more productive firms offering remote work, rather than remote work making firms more productive. Although they have data on thousands of firms, their preferred specification includes firm-level fixed effects, meaning their identification is driven by the sub-sample of 394 firms whose remote work status changed over 2011-2016. Using this sample, they find remote work has a small but statistically significant *negative* effect on firm-level productivity. Going from 0 to 100% (potentially) remote workers decreases productivity by 7-10%, depending on how the sample is weighted.

However, these headline numbers conceal important variation among types of firms. It turns out the negative impact of remote work is concentrated in lower-performing firms: mid-size and large firms do not see reduction in productivity when they switch to remote work; neither do exporters; neither do firms with high-skilled employees. In fact, firms that do R&D increase their productivity by 9% when they go from 0 to 100% (potentially) remote.

By controlling for time-invariant firm characteristics, and industry-specific trends, Monteiro, Straume, and Velento (2019) provide more compelling evidence that the productivity of remote workers is just as high as colocated ones, at least in high-performing firms. However, a weakness of the paper is the noisy estimate of remote work, which is based more on the possibility of remote work, rather than actual uptake.

Better evidence for the causal effect of remote work comes from various natural and field experiments. Sherman (2019) conducts a field experiment where 187 employees of a life-sciences company in Cambridge, UK worked from home on even or odd weeks of a four-week period (allowing each participant to serve as their own control). During remote work weeks, employees worked, on average, 2.14 days from home, compared to 0.49 during in-office weeks. Self-rated performance was higher during remote work weeks.

Bloom et al. (2015) describes another remote work experiment made possible by the remarkable fact that one of the co-authors was a co-founder, first CEO, and Chairman of a large Chinese travel-booking company, CTrip. This allowed the authors to run a real experiment on a real company at large scale on employees who operated out of a Chinese call center.

In the experiment, CTrip conducted a lottery to assign half of 249 interested and qualified employees to work from home four days per week (to qualify, they needed six months tenure, broadband internet, and a private room to work in). Once chosen, workers were not allowed to switch back to a colocated office until the nine-month experiment concluded. The experiment found home worker performance increased 13% relative to the control, mostly due to an increase in time spent on calls during the workday, which workers attributed to quieter working conditions. Attrition also dropped 50%, relative to the control. Based on the results of the experiment, CTrip estimated participation in the program saved $2000 per worker per year, and expanded the remote work option to its entire company.

While Bloom et al. (2015) provides clean evidence that remote work resulted in higher productivity, it is not obvious how applicable the results are outside of their specific context: a call center where workers were still geographically constrained by the need to appear in the office one day a week. Choudhury, Foroughi, and Larson (2019) provides complementary evidence that full-time remote work policies can increase productivity.

Choudhury, Foroughi, and Larson (2019) study various remote work programs at the US patent and trademark office (USPTO). Analogous to the program described in Bloom et al. (2015), qualifying patent examiners have access to a program where they work from home four days a week and in the office once a week. As with Bloom et al. (2015), USPTO remote workers were more productive than their colocated peers, though this analysis is observational, not experimental.

The main focus of Choudhury, Foroughi, and Larson (2019), however, is on a more generous program for which it is possible to identify the causal impact. In 2011, the USPTO began implementing a "work-from-anywhere" program that allowed qualifying examiners to work remote all the time, and to relocate to anywhere (though they could be required to return to headquarters for up to 12 days or 5 trips per year). This program was oversubscribed, resulting in quasi-random allocation of the program to interested workers. Choudhury, Foroughi, and Larson (2019) compare the performance of workers selected for the work-from-anywhere program to unsuccessful applicants in the standard work-from-home program.

They find remote workers in the work-from-anywhere program were 4.4% more productive than those in the standard work-from-home program (who were already more productive than colocated workers). Suggesting the importance of modern technology for enabling productive remote work, they also find a requirement that workers use modern collaborative software raised productivity for less experienced examiners (but had no effect on experienced ones).

These studies provide strong identification that remote workers can actually have higher productivity than colocated ones. But again, the applicability of these findings to other forms of work could be questioned. Remote work may be less productive than colocation when extensive collaboration is required, since face-to-face communication becomes impossible with a distributed workforce. Sherman (2019) does include measures of "task interdependence" and finds that workers with above median task interdependence did experience *less* benefit from remote work (and statistically significant only at the 10% level), though not a negative effect of remote work. It may well be that the comparatively isolated work of patent examiners and call center operators is unusually suited to remote work.

For some suggestive evidence on the effects of remote work on an organization requiring extensive collaboration, Karis, Wildman, and Mane (2016) is a deep dive into how remote collaboration works (or doesn't) at google, based on surveys, observations of meetings, interviews, and data on equipment use. The paper concludes, "Technical tools make keeping in contact and sharing information much easier than in the past... Although there is no effective solution to the time differences involved in long distance East-West collaboration, the other challenges are slowly decreasing, and the advantages of having a distributed workforce can now be more fully realized."

Suggesting the ongoing challenges alluded to by Karis, Wildman, and Mane (2016) continue to shrink a 2019 internal study by google (Gilrane 2019) found no difference in the effectiveness, performance ratings, or promotions for teams and individuals whose work required remote collaboration. Stripe also recently made a large push into remote work, noting in its 2019 announcement (Singleton 2019) that "the technological substrate of collaboration has gotten *shockingly* good over the last decade" (emphasis in original) and that "while we did not initially plan to make hiring remotes a huge part of our engineering efforts, our remote employees have outperformed all expectations." All this evidence together suggests the term A in equation (1) is shrinking, small, and likely less than 1 in many industries.

3. Matching (Q_r/Q_c)

Section 2 argues that giving existing workers the option to work remotely does not significantly reduce their productivity; indeed, it may raise it. But firms that use remote workers also have access to different workers than colocated firms. This section argues the expected quality of remote workers, relative to the expected quality of colocated ones (Q_r/Q_c), is increasing.

The basic argument is simple. Remote workers can be hired from the local market but also from outside it. Therefore, remote positions can, in principle, access higher quality workers anytime there are higher quality workers outside the local market.

The problem for a firm is to identify these workers (and attract them). The next section will argue workers find remote positions attractive. This section discusses the steady improvement in technology for matching remote workers to interested firms. In the absence of direct evidence about trends in the hiring of geographically dispersed remote workers, I focus on three other trends that support the ability of distant firms and workers to match. These are: (1) increased use of the internet by job-seekers and employers; (2) development of effective algorithms and information on online job markets; and (3) greater potential to form and maintain distant relationships via social networks.

Trend 1: Increased use of the internet by job-seekers and employers
The main technology enabling better matching is, of course, the internet. A large number of studies have documented how the internet improves matching of buyers and sellers across many markets, such as airlines, books, rental and home vacancies, and food trucks (Goldfarb and Tucker 2019). Surprisingly, however, studies of the spread of the early internet found little evidence it facilitated better matching in labor markets. Studying the use of internet job search in 2000, Kuhn and Skuterand (2004) found it was no better than traditional job search. Kroft and Pope (2014) look at the city-by-city rollout of craigslist over 2005-2007 and also found no detectable impact on local labor market outcomes (though craigslist did reduce apartment vacancy rates and the use of print-based classified ads).

These results are a useful reminder that full utilization of new technologies is not instantaneous. Faberman and Kudlyak (2016) update Kuhn and Skuterand (2004) with data from 2011. In the 11 years between 2000 and 2011, the share of the unemployed with home internet access rose from 40% to 74%, and the share of unemployed job seekers using the internet to find work tripled from 26% to 76%. Using data from 2011, Farberman and Kudlyak (2016) find internet job searchers were significantly more likely to find work within one year. Indeed, as they point out, the use of internet job search and job postings had become so routine at the time of their writing that they are regularly used to study overall labor market dynamics by economists.

Online social networks - especially professional networks like LinkedIn - have also emerged as a secondary channel through which recruiters can learn about and contact potential hires. In 2019, 27% of US adults have LinkedIn profiles (Iqbal 2019). An emerging literature has documented that use of online social networks by recruiters is common across a variety of countries, though it presents its own challenges (Stopfer and Gosling 2013, Berkelaar 2017).

Trend 2: Development of effective algorithms and information on online job markets
A bigger pool of candidates and positions can potentially make sorting through job posts and applications more challenging, but a variety of tools and algorithms are being developed to facilitate online job matching. Early examples included referral links (which allowed users to send promising leads to other people) and online application portals that verify applicants meet minimum criteria. Brenčič and Norris (2012) study the use of these kinds of early job search tools by employers posting jobs on Monster.com over 2004-2006, finding they were indeed used by firms with the most experience with online job sites.

A number of papers have specifically studied the design of *oDesk* (now *upWork*), an online labor market for short-term remote work contracts. Horton (2017) studies the efficacy of algorithmic recommendation tools on match quality on *oDesk*. Horton's experiment relies on the fact that employers may invite specific *oDesk* workers to apply for positions. In 2011, the company performed an experiment where it provided recommendations of workers to invite to a random set of new employers. These recommendations were based on the availability of the worker (as inferred by *oDesk*), prior satisfactory performance, and skill match with the posted contract. A control group of employers were not shown recommended workers. Horton (2017) finds the algorithmically generated recommendations increased the probability of finding a suitable worker by 20% for technical jobs (but had no effect on nontechnical jobs).

Online labor markets also face issues related to trust. Since *oDesk* focuses on short-term contracts, these issues are particularly salient and the company has a number of features designed to reassure employers workers are not shirking. For example, employers can require jobs be completed while using a program that captures screenshots of the worker's computer. More importantly for the question of match quality, *oDesk* provides information about past performance and does not allow employees to delete all this information (Agrawal et al. 2013).

This kind of basic information on past performance appears to be quite important in these markets. Pallais (2014) performs an experiment where 952 randomly selected *oDesk* workers are hired for a 10-hour data entry task, and then given short or detailed publicly viewable performance reviews. Workers with these reviews are substantially more likely to be hired again on the market, compared to those not hired as part of the experiment. Agrawal, Lacetera, and Lyons (2013) also find validated information on experience significantly predicts hiring, even after controlling for other observables, and that this effect is strongest for workers from less developed countries.

Trend 3: Greater potential to form and maintain distant relationships via social networks
While the emergence and steady improvement of online labor markets facilitates matching, more than 50% of jobs have typically been found through social ties (Gee, Jones, and Burke. 2017). These matches also tend to be good ones, at least in terms of having lower turnover and higher productivity (Burks et al. 2015). To the extent that social networks are built by physical proximity, this is a powerful force for agglomeration and colocation. Our neighbors become friends, and our friends help us find jobs, and the jobs they help us find tend to be local.

That said, the emergence of email, videochat, and online social networks like Facebook, Twitter, and LinkedIn have the potential to expand our networks of geographically distant "weak ties" (relationships that are more acquaintance than friend). These weak ties have long been thought to play an important role in job search, since their connections to different communities makes them more likely to transmit novel information (Granovetter 1973).

Gee, Jones, and Burke (2017) test the weak ties hypothesis using data on 6 million Facebook users over 2007 to 2011. They find weak ties (as measured by the extent of overlapping friends or the extent of interactions) collectively help users find more jobs than close friends, simply because people don't have many close friends but have many weak ties. Meanwhile, Dunbar (2017) finds some evidence that social networks help (young) people acquire more weak ties. In two 2015 surveys of UK adults, Dunbar finds most users of social networks report 150-180 friends on social networks, which is not noticeably larger than the typical size of offline social networks of around 150 (the so-called "Dunbar number"). However, when Dunbar breaks his data down by age, it turns out 18-24 year-olds report many more friends - 282 on average. Most of the increased size of young people's social networks appears to reflect an expansion of weak ties, not close ties.

There's also some evidence social networks help users maintain geographically distant social ties. Laniado et al. (2018) study user interactions on a popular Spanish online social network in 2010. While they find users are more likely to *be* friends on the network if they are geographically close, once friendships are formed distance barely matters. The extent of online interaction between geographically close friends and geographically distant friends is very similar.

Online social networks may support the maintenance of a larger network of dispersed friends, but they do not necessarily help users form new distant relationships. There is some evidence though that this does occur in some online environments that function as "third places;" places where people "hang out" online and potentially meet new geographically distant people. McCulloch (2019) cites twitter as one example and it has also long been recognized that online multiplayer games can function as third places (Steinkuehler and Williams, 2006), since they are, essentially, a place for people to hang out and interact. Molyneux, Vasudevan, and Zúñiga (2015) surveyed a large nationally representative group of US adults and found players responded affirmatively to statements normally used to measure social capital among offline communities, such as "I feel close to the people I play games with" and "I feel like I am part of a community of gamers." Trepte, Reinecke, and Juechems (2012) survey 811 mostly European players of e-sports and find similar evidence that online gaming leads to the creation of social capital and offline social support. Online gaming is a very new way to form friendships, but one that might become more prevalent in the future. Forty-one percent of American adults play video games with other people online, averaging nearly five hours per week (Entertainment Software Association, 2019).

Finally, the internet and online social networks can also help people maintain relationships that are formed offline. A small literature on academic conferences and collaboration suggests short in-person meetings at infrequent conferences can be sufficient to establish lasting relationships between geographically dispersed individuals (Freeman, Ganguli, and Murciano-Goroff 2015, Chai and Freeman 2019, Campos, Leon, and McQuillen 2018, and Boudreau et al. 2017). Of particular relevance is Bakhshi, Davies, and Mateo-Garcia (2015), which explicitly measures the impact of attending the 2012 Le'Web tech conference on attendees' twitter social network. Following the conference, attendees "followed" other conference participants at a much higher rate than they followed non-attendees in the six weeks after it's conclusion.

To sum up, one major advantage of remote work over colocated work is that remote work enables firms to access a larger pool of potential hires, and job-seekers to access a larger pool of potential employers. This increases the relative quality of the best match for a remote worker, relative to a colocated one (Q_r/Q_c). However, realizing this advantage requires employers and employees to find each other. Fortunately, we have a variety of reasons to believe the ability of firms to match with remote workers is steadily increasing. Using the internet to search for employees and employers has become common, with better algorithms helping both parties sort through the larger pool of potential matches for the best fit. Social networks are another important part of job matching, but online social networks and communication technology are likely to increase the number of weak ties we can form and maintain across geographic distance.

A potential bottleneck to remote work is the tendency of people to most easily form relationships with people who are physically proximate. While the internet may make it easier to strengthen relationships formed online or via travel, physical proximity remains a major determinant in the formation of relationships.

But, framed another way, this can be viewed as a strength of remote work. Remote work offers the possibility of remaining physically close to our emotionally close relationships, while also pursuing job opportunities that are located far away. While the importance of geographic proximity is a problem for creating good matches, it is actually a potential strength for attracting geographically distant workers. The next section turns to the general question of how firms attract remote workers.

4. Wage Differentials (W_c/W_r)

Remote work capabilities allow a firm to access a larger labor market, and this will lead to higher quality matches if workers and firms can find each other, and if firms can entice distant workers to work for them remotely. The last section reviewed evidence on matching at a distance. This section reviews the costs of attracting and retaining remote workers.

Colocation is expensive. Autor (2019) shows that while the urban wage premium for high-school educated workers fell significantly between 1970 and 2015 (as middle-skill urban jobs disappeared), the urban wage premium for college-educated workers remains as strong as ever. In 2015, a college educated worker in the densest part of the country earned an average hourly wage of approximately $25/hr, compared to approximately $15/hr in the least dense parts of the country (Economist 2019). As will be discussed in section 6, college educated workers also tend to have the greatest scope for remote work. If the productivity of remote workers is sufficiently close to those for colocated workers, then this suggests significant savings are possible for companies that hire remotely.

Variation in the cost of living may also benefit remote workers. Moretti (2013) estimates approximately one-quarter of the college-educated wage premium is accounted for by higher cost-of-living in the places where college-educated workers live. In the USPTO's work-from-anywhere program, workers were paid the same salary, regardless of where they chose to reside. Choudhury, Foroughi, and Larson (2019) show that participants in the program enjoy substantial increases in their real wage by moving to places with a lower cost-of-living.

There are potential cost savings beyond salary as well. In Bloom et al. (2015)'s study of remote work at CTrip, the company estimated it saved $1250 per worker per year on reduced office space costs. The US Patent and Trademark office estimated that it's work-from-anywhere program saved the agency $38.2mn in office space costs in 2015 alone (Choudhury, Foroughi, and Larson 2019).

Bloom et al. (2015) also found remote work reduced turnover by 50%, which led to annual savings of $400 per worker. The fact that turnover fell so dramatically suggests workers like remote jobs. Mas and Pallais (2017) attempt to directly estimate the value workers place on the option to work from home by performing an experiment with job applicants for a position in a call center in 2016. As part of the application process, applicants were shown two versions of the position they were applying for, with different (randomly selected) wage differentials, and asked to choose which one they would prefer, if both were available. On average, workers were willing to accept an 8% pay cut to enjoy the option of working from home. Notably, this is far more than they are willing to accept in exchange for other perks, such as the option to set their own schedule.

That said, Bloom et al. (2015) find some evidence that work from home was not as valuable after workers tried it. After their nine-month experience with remote work, 50% of the remote workers opted to return to the colocation option. Only 35% of those who originally wanted to work from home took up the option, when it was extended to them after the experiment. Interviews suggested workers felt lonely and isolated at home.

On the other hand, in a supplemental experiment Mas and Pallais (2017) find employees already working from home were willing to take a 19% pay cut on a hypothetical job to keep the option to work from home (compared to a 9% pay cut for those without an existing work-from-home arrangement). These results suggest the population that self-selects in remote work values it more than the general population. In general, there is no reason to believe the value of remote work will not vary across people, firms, industries, and countries.

Taken together, variation in the cost of living, potential savings on office space and turnover, and the amenity value of remote work itself are all reasons why the cost of remote workers is likely to be lower than the cost of colocated ones.

Declining Geographic Mobility

There is also reason to believe this wage differential may be increasing over time, at least for remote workers with the option to work from anywhere. Between 1985 and 2010, the annual probability a working age US civilian adult moved to a new state declined from 3.5% to under 2%, while the share who moved specifically for a new job or transfer nearly halved between 2000 and 2010 (Kaplan and Schulhofer-Wohl 2017). The trend of declining geographic mobility presents another argument for remote work. In section 3, I argued that, compared to colocation, remote work has access to a larger labor market and can therefore obtain higher quality matches. This, however, ignored the possibility of colocated firms hiring distant workers willing to move. This appears to be becoming more difficult though.

Exactly why geographic mobility is on the decline remains an open question. Consistent with the evidence in section 3, Kaplan and Schulhofer-Wohl (2017) suggest better information about geographically distant jobs has led to reduced mobility (for example, because fewer people move and then move back after discovering they dislike the job or region). Hsieh and Moretti (2019) attribute part of the decline to land use restrictions in major cities, that strongly curtailed the expansion of the housing stock. Johnson and Kleiner (2017) argue rising occupational licensing requirements are partially to blame.

But another strand of research suggests the "problem" may be rising attachment to our homes: when we have the chance to move for a higher salary, more people prefer the friends and family to the money. Dahl and Sorenson (2010a) and Dahl and Sorenson (2010b) attempt to quantify the monetary value of physical proximity to close relationships using a 2003 and 2006 sample of Danish adults. For example, in one exercise, they look at where Danish blue collar workers move after the plant they work at closes. By correlating the typical salary of blue collar workers in the region where workers eventually choose to reside with the presence of friends, family, and other characteristics, they can estimate the dollar value of proximity to home. They find a doubling of the distance to the hometown is associated with a salary that is approximately $10,000 higher (compared to an average salary of $32,000). They find similar sized effects for movement away from friends. Effects are even larger in absolute and proportional terms for Danish scientists and engineers who work at establishments that are closed.

There are several reasons why home attachment may be rising. First; like any other normal good, as society grows richer, it will be willing to "purchase" more proximity to friend and family, in this case by forgoing higher salaries not to move. Second; the propensity to move declines with age, and our society has grown older. Karahan and Rhee (2017) argue that the aging of the workforce has also had the indirect effect of more local recruitment effort by firms, since they know older (distant) workers will be less willing to move; this also means younger workers are more likely to be locally recruited though. Third; Coate and Mangum (2019) show the propensity to move is declining in the depth of "roots" an individual has in their home region. As families put down deeper roots in the relatively "new" cities of the US South and West, their propensity to move from those cities has fallen substantially. If workers are less willing to move for the same package of wage and benefits as in the past, then this further strengthens the case for remote work.

5. Local Knowledge Spillovers (*S*)

We now turn to *S*, the strength of local knowledge spillovers. It has been amply documented that innovation and economic activity tends to cluster in cities. For example, Bettencourt et al. (2007) show patent production and economic activity cluster in cities at a faster than proportional rate - if city *X* has twice the population of city *Y*, it will tend to have *more* than twice the number of patents and economic activity. A variety of rationales for geographic clustering have been suggested, and these can be grouped under the broad headings of sharing, matching, and learning (Duranton and Puga 2003). In brief, cities allow better access to shared assets (infrastructure, but also labor pools and risk), they facilitate matching of (colocated) workers to firms, and they allow for the easier exchange of information (learning). Because these traits make firms more productive, firms will tend to locate in cities, despite the higher costs of doing so.

This section focuses on the last of these reasons, the access to knowledge that supposedly requires physical proximity. To the extent this is true, the potential advantages of remote work are attenuated. Suppose a firm allows workers to work remotely, but requires them to reside in the same city in order to enjoy spillovers *S*. Such a firm is restricted to the same local labor market as a colocated firm (or has to induce workers to move), with wages that reflect cost-of-living in the region. For such a firm, Q_r/Q_c and W_c/W_r are not much larger than 1. Without access to some of the strongest advantages of remote work, firms may decide to ask workers to work in a colocated office too.

This section presents evidence that *S* is declining and may be equal to 1 for some types of economic activity. Balland et al (2019) look at the extent to which patent activity clusters into more populous cities over 1850-2010. They calculate the scaling coefficient for various indicators of innovation and economic activity. The interpretation is, that if the scaling exponent is *x*, then a 2-fold increase in population leads to a 2^x-fold increase in whatever is being measured. Like Bettencourt et al. (2007), they find a superlinear relationship ($x > 1$) between patenting and population still holds. But they also find the scaling coefficient peaked in 1960. Since then, the extent to which patents cluster into cities has declined.

In this section, I'll argue the decline in the clustering of innovative activity reflects the decline in the importance of local knowledge spillovers - the term *S* is equation (1). To begin, we'll look more closely at measures of local knowledge spillovers and show that they too are quite likely declining. Then, I'll argue the decline in *S* is driven by better information and transportation technology that have made physical proximity less important.

The Strength of Local Knowledge Spillovers is Falling

The canonical paper establishing the importance of local knowledge spillovers is Jaffe, Trajtenberg, and Henderson (1993). Jaffe, Trajtenberg, and Henderson (1993) attempt to establish the importance of local knowledge spillovers through patent citations. The idea is that if (1) patent citations indicate knowledge transfers and (2) patents are more likely to cite the patents of local firms and inventors, then (3) this is evidence that knowledge transfers most readily when it's local.

But there is a complication. Suppose Jaffe, Trajtenberg, and Henderson (1993) establish patents are indeed more likely to cite local patents. This could be because local knowledge transfers more easily than distant knowledge. But it could also be because more of the relevant knowledge is nearby. As noted above, clustering can happen for several reasons, including reasons that have nothing to do with knowledge transfer (e.g., sharing or matching).

To establish citations reflect knowledge transfers, Jaffe, Trajtenberg, and Henderson (1993) asks us to consider three patents: *X*, *X**, and *Y*. *X* and *X** are technologically similar patents, and *Y* is a patent cited by *X* but not *X**. Jaffe, Trajtenberg, and Henderson (1993) show *X* and *Y* are more likely to belong to the same city than *X** and *Y*. In other words, patents are more likely to cite local patents, relative to non-citing controls that are technologically similar.

Kwon, Lee, Lee, and Oh (2019) bring Jaffe, Trajtenberg, and Henderson (1993) into the modern era. They replicate Jaffe, Trajtenberg, and Henderson (1993)'s methodology, as well as a critique from Fox-Kean and Thompson (2005), but bring data up through 2015 (and with much larger samples of patents). They find patents are even more likely to cite local patents today, relative to (various) non-citing controls.

This would seem to suggest local knowledge spillovers are actually strengthening. However, two other papers suggest citation-based measures of local knowledge spillovers have become increasingly unreliable.

Kuhn, Younge, and Marco (2020) show there has been a sharp increase since 2000 in the number of citations of dubious quality, driven by an increase in citing essentially every patent ever encountered by the applicant (because of a belief that this will better insulate their patent against legal challenge). To demonstrate the importance of changing patent standard, Kuhn, Younge, and Marco (2020) do a quick and simple update of Jaffe, Trajtenberg, and Henderson (1993) through 2010, using both the standard set of patent citations and a selected sample more likely to reflect genuine knowledge flows (the "improved selection"). Using a conventional set of citations, they show the probability a citing and cited patent belong to the same city increased from 3.9% to 6.4% between 1975 and 2010. Using only their sample of high-quality citations, it only increased from 4.3% to 4.7% though.

But as noted above, these numbers could be driven by non-learning factors that lead to geographic clustering. When they implement Jaffe, Trajtenberg, and Henderson (1993)'s control patent methodology with high-quality citations, they find the increase in the probability a citing and cited patent belong to the same city has actually declined. In 1975, a citing and cited patent were 3x as likely to belong to the same city as a control and the cited patent. By 2010, this was under 2x. In other words, patents are a lot less likely to cite other local patents than they used to be.

Arora, Belenzon, and Lee (2018) takes a second look at Jaffe, Trajtenberg, and Henderson (1993) and essentially show that a placebo analysis based on patent citations unlikely to reflect knowledge flow obtains the same result. In particular, they focus on "citation reversals", which occur when patent *X* cites patent *Y,* but patent *Y* is applied for *after* patent *X*. This occurs about 4% of the time. For citation reversals, it is impossible for the citing patent to have read the cited patent, and so it is impossible patent *X* learned about patent *Y* by reading it. While it may have learned about the patent through other channels, compared to a typical citation it is more likely the citation was added to define the scope of the patent's claims or establish patentability, and not as a receipt for borrowing knowledge. Therefore, (1) if "typical" citations are more likely to reflect genuine knowledge transfers than citation reversals, and (2) local knowledge transfers most easily, then (3) typical citations should have a stronger local bias than citation reversals.

Yet Arora, Belenzon, and Lee (2018) show, for a sample of 1.7mn citations over 2001-2014, there is no statistically significant difference in the propensity of citations and citation reversals to point to local patents (relative to controls). Since citation reversals *do* show a local bias but are more likely to reflect technological similarity than learning, one interpretation is that the Jaffe, Henderson, Trajtenberg (1993) methodology fails to adequately control for technological similarity.

I take two main lessons from Kuhn, Younge and Marco (2019) and Arora, Belenzon, and Lee (2018). First, citations are an increasingly poor measure of anything beginning after the year 2000, because of changes in citation practice. Second, the Jaffe, Henderson, and Trajtenberg (1993) methodology of using the patent classification system to identify "control" patents may be inadequate. Accordingly, I turn to some papers that use neither patent citations (after 2000) nor the Jaffe, Henderson, Trajtenberg (1993) approach to inferring measure the intensity of local knowledge spillovers.

Fitjar and Rodriguez-Pose (2017) do not rely on patents at all. To test whether cities facilitate chance and casual encounters that result in innovation, in 2013 they ask 542 Norewegian firms about the most important partner involved in the creation of a new product or process in the preceding three years, and then how this partnership formed. They find casual and chance encounters account for about one fifth of these partnerships (targeted search accounts for the remaining four fifths). However, there is no statistically significant difference in the probability a partnership is formed via chance in and outside cities. Indeed, if anything, casual and chance encounters are more likely to result in new partnerships for firms based outside of cities.

Mewes (2019) investigates a similar question using US patent data. Mewes (2019) assumes innovation is about discovering new ways to combine existing knowledge. Access to knowledge about different technologies locked away in other people's heads is quite useful for this kind of innovation. And indeed, a paper by Berkes and Gaetani (2019) finds city centers do generate more unusual combinations of technologies.

Mewes' work uses patents, but not patent citations so it shouldn't have the same data quality issues as citation based papers. Instead he uses the technology classifications assigned to patents; when a patent is classified as belonging to technology classes *X* and *Y*, Mewes interprets this as an indicator that the patent combined pre-existing technologies *X* and *Y*. He is particularly interested in novel and unusual combinations of technology, since we would expect cities to have disproportionate advantage for these novel combinations. Mewes looks at how the size advantage of big cities in producing "atypical" combinations has changed between 1850 and 2010.

Mewes' approach is to calculate the scaling exponent for county population over 1850-2010, as in Balland et al. (2019). The interpretation here is, that if the scaling exponent is x, then a 2-fold increase in population leads to a 2^x-fold increase in atypical technology combinations. The higher is x, the bigger the advantage of a big population in generating unusual combinations. But the scaling exponent peaked in the 1970s and has been on a steady decline since. Big cities no longer have as strong an advantage in generating unusual combinations of technology.

Packalen and Bhattacharya (2015) take a different tack again. They scan the text of all patents and pull out important one-, two-, and three-word sequences (e.g., "microprocessor" and "polymerase chain reaction"). They call these word(s) "concepts" and interpret them as technological ideas. Because they can observe the date each concept is first mentioned in a patent, they can measure the "age" of the idea. They then look at whether patents in cities use newer ideas than patents from outside cities. They find the patents of big cities used to have a much higher probability of mentioning a very young concept, but this faded over time. By the 2000s, they can no longer rule out zero effect.

Besides patents, academic papers represent another avenue for studying the importance of local knowledge spillovers. Unlike patents, there is little concern over the quality of citations in academic references. Focusing on a set of 5.5 million pairs of papers from 2012 and 2014, Wuestman, Hoekman, and Frenken (2019) study the importance of geographic proximity in academic research by estimating the probability a cited paper's authors work within 20km of the citing paper's authors. And they find papers are indeed more likely to cite local work.

However, just as with patents, this may not actually reflect the fact that knowledge transfers more easily at the local level. It may simply be that different regions specialize in different kinds of research (for reasons that have nothing to do with local knowledge spillovers) and local citations merely reflect the location of relevant knowledge. This seems to be partially the case. When Wuestman, Hoekman, and Frenken (2019) restrict attention to citations between papers drawing on similar knowledge, they find local papers are no longer any more likely to be cited. That said, citations between papers that are dissimilar does seem to retain a local bias. But the analysis does not let us ascertain whether this has fallen over time or not.

Lastly, Kim, Morse, and Zingales (2009) examine the research productivity of academics in economics and finance over 1970-2000. By keeping track of the moves of individual academics, they can see how the productivity of researchers changes when they move to top universities (where research productivity is measured in terms of how many articles they write per year, adjusted for the quality of journal). In this way they can infer the impact of colocation with other academics on the individual productivity of a researcher. They find the positive impact of being in a top university declined in each decade, such that there was no positive impact of being in a top university during the 1990s. Similarly, the research productivity of colleagues in a department was positively correlated with an individual's research output in the 1970s, but after controlling for individual productivity, this effect disappears in the 1990s.

Better Transport and Communication Technology Facilitates Access to Distant Knowledge

Thus, recent work finds the strength of local knowledge spillovers, measured in various ways is falling or quite low in the last decade. In this section, we present evidence that this is driven by improvements in transportation and communication technology.

Catalini, Fons-Rosen, and Gaule (2018) look at academic collaboration between chemists living in different cities after Southwest Airlines opens a new route connecting them. They find in the years after new (low-cost) airline routes connect them, chemists publish 50% more articles co-authored with chemists on the other end of the route. The effect is stronger for collaborations across different fields and when both chemists are more productive than the average for their department - both cases when being able to reach outside your local contacts is important. Similar effects exist for other disciplines.

It's not only planes. Dong, Zheng, and Kahn (2018) look at collaboration between academics in China when cities are connected by high speed rail. They find similar results (although their results are more fragile and can partially disappear depending on the econometric method): after a high-speed rail line is built between two cities, there is an increase in the number of papers coauthored by academics based in the cities. This effect is strongest when a "secondary" city is connected to a "mega" city, and when the cities are close enough so that high-speed rail becomes faster than air travel.

The evidence from airplanes and high-speed trains is consistent with roads as well. Agrawal, Galasso, and Oettl (2017) look at what happens to innovation when US regions build more highways. Agrawal, Galasso, and Oettle (2017) look at private sector innovation and focus on the local impact, rather than how interstates enable collaboration across regions. They find a 10% increase in regional highways is associated with 1.7% more regional patents over 5 years.

But peek beneath the surface and this is another story of how falling transportation costs erode the importance of local knowledge. The authors focus on citations patents make to other patents from the same region: the more roads, the greater the distance between these patents. They also show the impact of roads is strongest in low-density cities, where inventors are more geographically disperse. Intuitively, after my city builds a new interstate I'm more likely to cite patents from across the city instead of across the street, especially if there aren't many inventors nearby. Roads enable more local-but-not-that-local knowledge flows.

In general, anything that increases access to distant knowledge can erode the importance of being physically proximate to people. Furman, Nagler, and Watzinger (2018) study the impact of patent depository libraries on local innovation. Comparing recipients of these libraries to nearby eligible sites that did not receive libraries, they show increased access to information (i.e., the text of patents) increases local patenting by 17%. Moreover, patents from inventors living near patent libraries are more likely to cite patents belonging to more distant inventors. The library apparently reduces the need to be near other inventors to make use of their ideas.

Of course, you still need to be near the library. Or at least, you did, until the internet. Furman, Nagler, and Watzinger (2018) also show the local impact of patent depository libraries evaporates once the first internet searchable patent databases become available in 1995.

Improvements in digital communication technology have the same effect of facilitating remote collaboration and reducing the need to be geographically proximate. Forman and Zeebroeck (2012) and Forman and Zeebroeck (2019) both look at how internet access changes the collaboration patterns of firms with geographically dispersed establishments. To measure the impact of internet access, the papers have to reach back a long way, to the

1992-1998 era, when internet access was first beginning to roll across America. Forman and Zeebroeck (2012) show that after two establishments are connected to the internet, inventors in the connected establishments are more likely to be jointly listed on patents. In contrast, getting internet access doesn't seem to have any impact on the number of solo-inventor or geographically clustered inventor team patents, which suggests the internet's main advantage was in facilitating collaboration, not merely in increasing access to knowledge.

But it does that too. Studying the same era, Forman and Zeebroeck (2019) show that when two establishments are connected to the internet, patents by inventors from one establishment are more likely to cite patents by inventors from the other.

The preceding papers show internet access facilitates access to distant knowledge, but do not directly address the issue of whether this leads to less geographic clustering of inventive activity. Forman, Goldfarb, and Greenstein (2014) do. They compare the growth of regional patenting over 2000-2005 to patent levels over 1990-1995. Overall, they see a significant increase in the concentration of inventive activity: the counties with the most patents over 1990-1995 also had the fastest growth of patents over 2000-2005. What is novel, however, is that they show this effect is reduced by greater internet access.

For example, consider two counties with low internet connectivity in the 1990s. Specifically, among all US counties, suppose they are at the 25th percentile for internet adoption. If one of these counties was also in the 90th percentile of 1990s patents, it experienced annual patent growth in the 2000s that was 5.4% faster than a comparable county in the 25th percentile for 1990s patents.

Now consider two counties with high internet connectivity in the 1990s (the 90th percentile for internet adoption). If one of these counties was in the 90th percentile of 1990s patents, it experienced annual patent growth in the 2000s that was just 0.4% faster than a county in the 25th percentile. The innovation gap between counties grew much more quickly for counties without much internet access.

What's more, the difference between high-internet adopting and low-internet adopting counties is largest when we restrict attention to patents featuring distant collaboration among inventors. One interpretation of these results therefore is that people living in innovative counties in the 1990s didn't really need the internet to find potential collaborators, so it's presence or absence didn't matter that much. But people living in less innovative regions benefited a lot from internet access, because it allowed them to find good collaborators and participate in the innovation economy.

Summary

This section began by noting that geographic clustering of economic activity tends to be clustered in cities, but that this clustering has actually fallen along some dimensions over the last several decades. Geographic clustering can be caused by a variety of factors, sometimes grouped under the headings sharing, learning, and matching. This section has proposed that the learning advantage of geographic proximity (represented by S in equation 1) has been falling for a long time, due to improvements in communication and transportation technology.

6. The Rise of Remote Work

The previous four sections review a wide range of evidence to argue the four terms in equation (1) are all moving to make remote work more attractive. The productivity of individual colocated workers relative to remote ones (*A*) is quite close to or less than one, at least in some industries and as technology continues to advance, this is likely to be true in more industries. The ability to find high quality remote workers from large national labor markets (Q_r/Q_c) is rising as we move more and more of our social and professional lives online. The cost of colocated workers relative to remote ones is greater than one (W_c/W_r), and may be growing if home attachment continues to rise. And the knowledge spillovers that accrue to collections of colocated workers (*S*) is falling rapidly, and may even be gone.

If this theoretical framework is correct, then firms should also be aware of it and at some point we should begin to see an increasing share of positions go remote. In the presence of adjustment costs, this process may be slow to get going, but at some point it should still occur. In this section, I review evidence that, pre-covid-19, we were already seeing the rise of remote work.

The easiest place to start is simply by looking at census data, which asks whether respondents work from home. Figure 3 presents the share of non-farm workers working from home.

Figure 1. Share of Nonfarm Workers Working From Home, 1960-2017

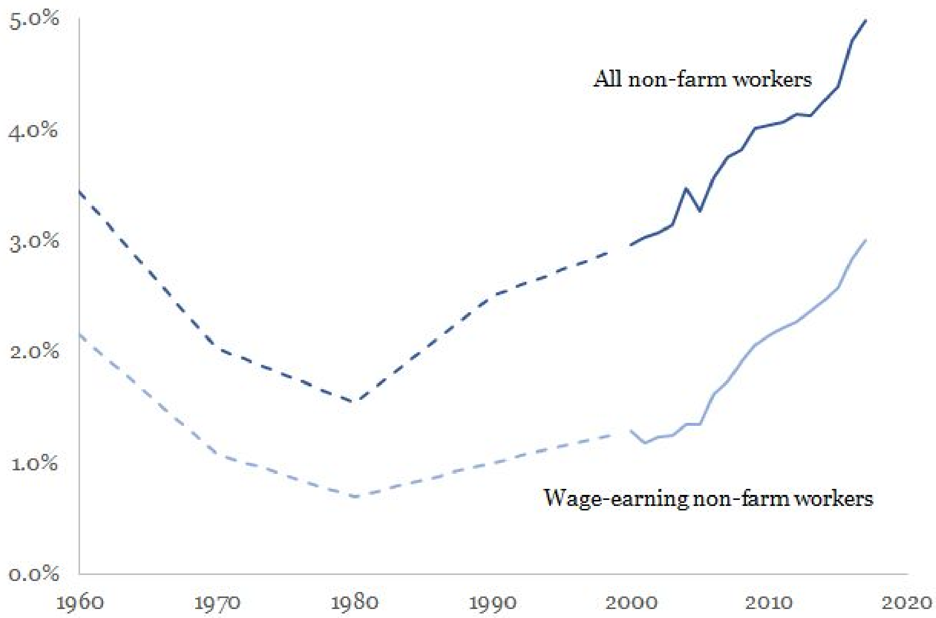

Note: Data from 1960-1990 is drawn from census means of travel to work data and available in 1960, 1970, 1980, and 1990 - other values are imputed. Data from 2000-2017 is drawn from the American Community Survey means of transportation to work data and is available annually. Data queried from IPUMs by Lyman Stone.

The share of people working from home has risen from a low of around 2.0% in 1980 to nearly 5.0% by 2017, and has accelerated in the most recent years. These workers can be divided into two types: self-employed workers and wage-earners, with the latter category a better match for the kind of remote work envisioned in this essay. Breaking out wage-earners suggests some of the increase in work-from-home between 1980 and 2000 was driven by an expansion of the self-employed. However, since 2000, most of the increase in work from home has been driven by increases in the number of wage-earners working from home; that is, from an increase in remote work.

Since these numbers focus exclusively on those who work from their home, they omit those who work from satellite offices, co-working spaces, coffee shops, and other alternatives. Accordingly, they understate the true extent of remote work. A 2019 survey commissioned by Upwork and the Freelancers Union finds a similar share of people who work from home as the census (about 5%), but found that including alternative work sites raises the share of remotes to 9.5%; including those who work remotely some of the time increases the percentage to about 36.1% (Ozimek 2019).

The ability to work remotely varies significantly by industry. Only 8.8% of leisure and hospitality workers have the option to sometimes work from home, but more than 50% of workers in the information, financial activities, and professional and business services industries do. Within industry, the option to work remotely is concentrated in management, business, and financial operations occupations (60.1%), and to a lesser extent in professional and related occupations (42.5%). Lastly, the option to work remotely is heavily concentrated in high earning and high education positions. Fully 51.9% of workers with a bachelor's degree or higher have the ability to work from home (some of the time), as do 61.5% of those in the top earning quartile (BLS 2019).

A variety of evidence suggests we are a long way from reaching the limits of remote work. To begin, the 2016 Understanding America Survey simply asked a representative sample of US adults whether their job could feasibly be completed from home. Among college educated workers, 41% say yes. Among those with a high school degree or less, the number is 14% (Mas and Pallais 2020). Alternatively, Dingel and Neiman (2020) use data from two Occupational Information Network surveys on "work context" and "generalized work activities" to estimate the share of jobs that can plausibly be done remotely today. Their approach is to rule out occupations as a candidate for remote work if their survey responses indicate the job cannot be performed online (for example, if the average respondent says they have to wear protective or safety equipment, they would consider the position unsuitable for remote work). Using this approach, they estimate 34% of jobs can be done remotely (though they interpret this number as an upper bound). Lastly, in 2017-2018, 28.8% of workers aged 15 and over could work from home at least some of the time (and 24.8% did at some point) (BLS, 2019). Together, these studies suggest in the current job landscape, an upper bound of one fifth to one third of jobs can be done completely remotely, though this number will be skewed toward high education positions.

There is also demographic momentum for an increase in remote work. Ozimek (2019) finds 69% of Gen Z and millennial hiring managers allowing team members to work remotely, compared to 58% of baby boomer hiring managers. And more than 40% of small business owners aged 18-34 planned to hire full-time remote workers, compared to just 10% of small business owners aged 50 and up. And though the differences are smaller, younger workers are typically more interested in working remotely or are already working remotely. All this suggests demographic changes, on their own, could make remote work more prevalent, even without the changes documented in this paper.

An alternative measure of our ability to work remotely is to look at trends in the extent of knowledge work conducted by geographically dispersed collaborators: patents and academic papers. Figure 4 plots the share of US-coauthored scientific publications with (1) authors belonging to multiple institutions and (2) international collaborations.

Figure 4. Share of US Scientific Articles with Multiple Institutions and International Institutions, 1996-2018

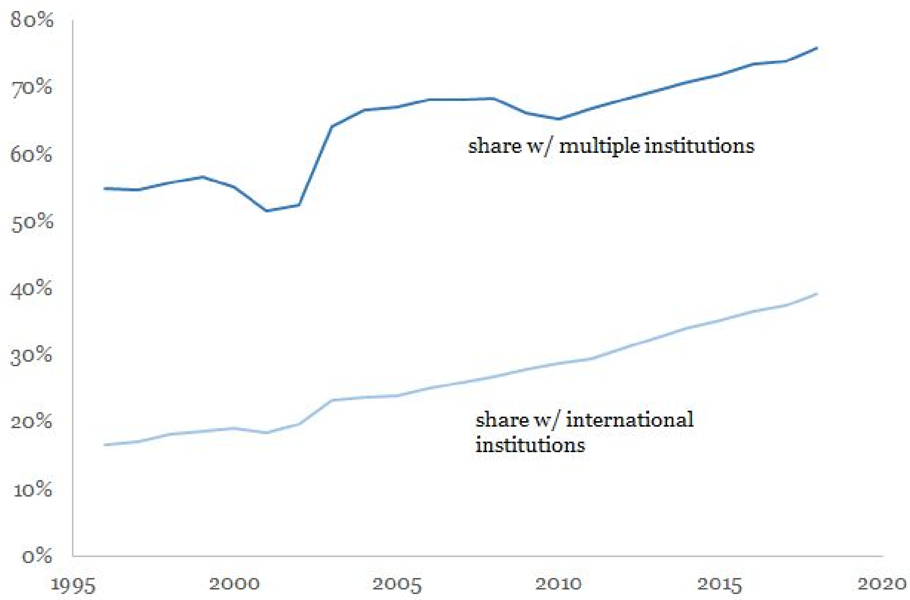

Note: Data from the National Science Foundation, Science and Engineering Indicators 2020. Articles are all journal articles listed in Scopus with at least one US-based coauthor.

The National Science Foundation data does not report the share of publications that are produced by colocated and remote collaborators, but two proxies are available. First, the share of papers with multiple institutions provides a rough proxy for remote collaboration, since the majority of papers written by distant collaborators will belong to different institutions. This measure has increased from 55% in 1996 to 76% in 2018. This likely exaggerates the extent of remote collaboration though, since two co-authors may belong to different institutions that are geographically close. A second proxy that does not share this issue is the share of papers with a US institution and at least one institution belonging to another country. This has risen from 17% to 39% over 1996-2018. This measure likely understates the extent of remote collaboration, since it omits collaborations between geographically distant US institutions. Note that both measures may also be biased if coauthors belong to different institutions but temporarily reside in the same place while collaborating, or if they belong to the same institution but reside in different places while collaborating (i.e., because one of them is on sabbatical). So long as this bias does not systematically increase over 1996-2018, the trend is still towards more remote collaboration.

Freeman, Ganguli, and Murciano-Goroff (2014)'s survey of papers in the fields of particle and field physics, nanoscience and nanotechnology, and biotechnology and applied microbiology, provide evidence that suggests the above trends reflect a genuine increase in remote collaboration. Freeman, Ganguli, and Murciano-Goroff (2014) map the institution of each co-author to an associated city. They define papers as "US-collocated" if all coauthors belong to the same (US) city. By this measure, the share of papers with at least one distant collaborator in their sample rises from 0.5 to just over 0.6 over 1990 to 2000, and then stays constant slightly above 0.6 through 2010. Figure 4 shows a similar trend for the overlapping time period: the share of papers with multiple institutions rose during the late 1990s before getting stuck around 67% between 2003 and 2011. From then on, however, the share of papers with multiple institutions steadily climbed to 76% by 2018.

As an alternative to academic collaboration, we can look at the location of inventors listed on US patents. The US Patent and Trademark Office's Patentview website lists the latitude and longitude of each inventor, which can be used to calculate the distance between inventors. Figure 5 displays the average distance between inventors listed on a patent for patents with 1-5 inventors (which account for the vast majority of patents) and for patents with exactly 2 inventors.

Figure 5. Average Distance (km) Between Inventors on a Patent, 1975-2015

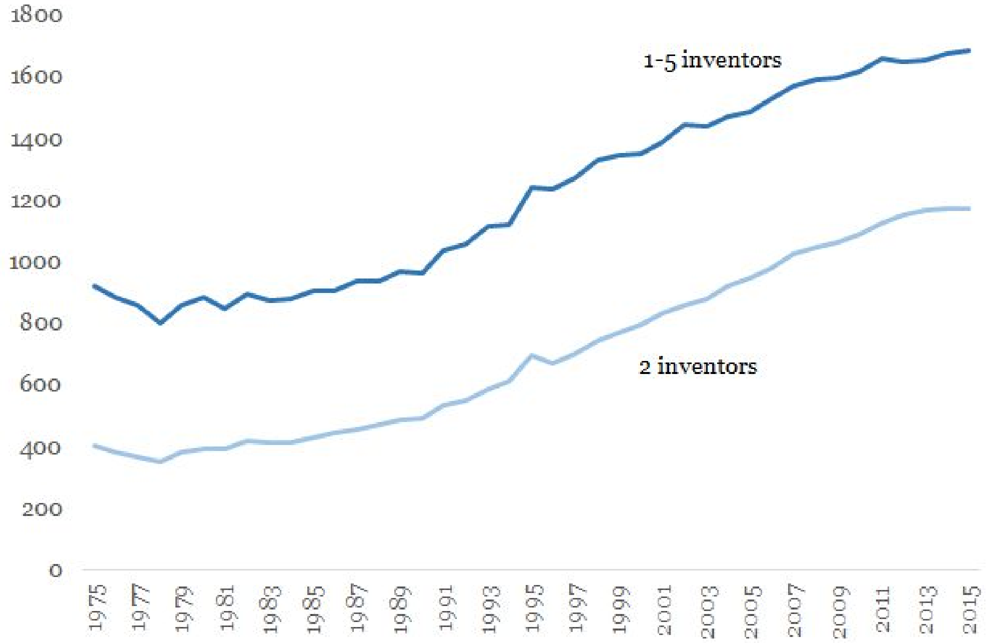

Notes: Distance is the average across all US patents applied for in a year of the average pairwise distance between all inventors listed on a patent, for patents with 1-5 inventors.

Patents with 1 inventor have a distance of 0. Data comes from patentsview inventor location data.

The average distance between inventors stayed between 800 and 1,000 km between 1975 and 1990, and then began to steadily climb toward 1,800 km by 2015. However, this aggregate measure conceals a composition effect as the number of inventors per patent has also steadily risen over this period (though this might also be driven by easier distant collaboration). However, even when attention is restricted to patents with two inventors, we see a steady rise in the average distance between collaborators, from approximately 400 km to 1,000 km between 1975 and 2015.

Alternatively, we can look at the extent of colocation by computing the share of patents with inventors who are unlikely to be collocated. Figure 6 reports the share of all patents with two inventors where the inventors are over 100 km apart and over 500 km apart.

Figure 6. Share of Two-Inventor Patents where Inventors are Greater than 100km and 500km apart, 1975-2015

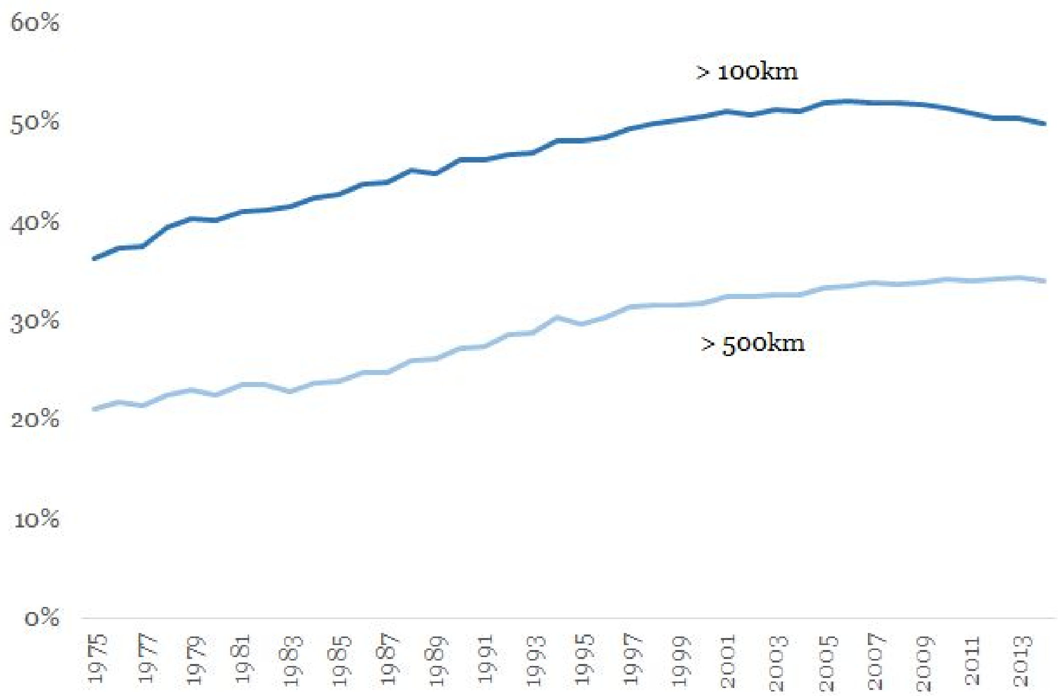

Notes: Data from patentview.org.

Figure 6 indicates it is indeed becoming more common for geographically distant inventors to collaborate on a patent. Inventors more than 100km apart accounted for 37% of two-inventor patents in 1975, and 50% of two-inventor patents in 2015. This is down slightly from a peak of 52% in 2007. However, the share of two-inventor patents with inventors more than 500 km apart rose from 22% to 34% over the same period.

Summary

Sections 2-5 presented a wide variety of reasons why remote work is becoming more attractive, relative to collocation. As predicted, remote work actually is becoming increasingly commonplace, as measured in a variety of ways. Remote collaboration on patents and academic research has grown to the point where it is no longer unusual. The share of positions that are performed from the home has begun to rapidly rise (from a very low base), especially since the mid-2000s. And these numbers seem to have a lot of scope to continue rising, given the estimates of the share of positions that can be performed remotely.

7. Discussion

The preceding has presented a set of reasons why remote work is likely to become more common. The individual productivity of remote workers appears to be quite high in many industries and will likely to continue to improve with technology. Technology for matching firms to geographically distant workers is also improving rapidly. And remote workers are likely to be cheaper than colocated ones. The spillover benefits of colocation have declined substantially, as better communication and transportation technology enables ideas to circulate more widely. Finally, many firms and workers seem to agree with this assessment, as the extent of remote collaboration is rising.

Promoting Remote Work for the Public Good
For the most part, the arguments for remote work presented in this paper are about private costs and benefits of remote work. They are likely to push the economy towards a greater use of remote work, whether or not the government adopts policies to promote remote work. However, there are also positive externalities from remote work, such that active government policy to promote it may be desirable. These go beyond their utility in combating any future pandemics.

First, remote work may be able to raise aggregate productivity for the economy. In most models of agglomeration economies, there are benefits to greater agglomeration - such as better matching or learning - but these are eventually offset by costs associated with congestion (Duranton and Puga 2003). These models pin down the predicted population of cities or regions by finding the population level for which the benefits of moving to the city for a potential new arrival (due to agglomeration effects) are exactly offset by the costs of moving (due to congestion costs). To the extent that there are positive externalities from agglomeration, in these models it can be optimal to adopt policies that encourage more people to move into cities, since that increases the size of the economy, which can be redistributed as desired (Hsieh and Moretti 2019, Rossi-Hansberg, Sarte, and Schwartzman 2019).

However, these models assume the benefits of agglomeration are only obtainable through physical proximity. As discussed in sections 3 and 5, learning and matching may not much depend on physical proximity anymore, if firms switch to a remote work equilibrium. If we can move into an equilibrium where we can obtain the benefits of agglomeration without congestion costs, then in principle we can benefit from (digital) agglomeration of many more workers, without the attendant congestion costs. In practice, this means a remote equilibrium could potentially support the excellent matches and exchange of knowledge of a very large and dense city, without the associated housing and commuting costs.

Second, remote work can reduce geographic inequality. The increased importance of agglomeration effects over the last several decades have led to economic prosperity for cities and economic decline in rural areas. This has had significant political and social cost (Wilkinson 2019, Case and Deaton 2020). While remote work is not a panacea, by decoupling where people live and work, it spreads economic activity more equitably and may reverse the tendency for economic activity to cluster in a small number of superstar cities.

Third, remote work may contribute to a reduction in carbon emissions by reducing commuting. This reduction will not be to zero for a number of reasons: it may be partially offset by more frequent long-distance travel to meet colleagues face-to-face; workers may still commute to coworking space or satellite offices; it may enable workers to live in smaller cities which are less energy efficient. Still, it is easy to imagine the net effect will be towards less emissions.

If there are indeed positive externalities associated with remote work, then government policy to promote remote work is appropriate. An examination of potential policies is beyond the scope of this essay, but I here suggest a few possibilities.

Policy #1: IT Infrastructure

Most obviously, remote work is only feasible when there is a robust underlying IT infrastructure. This means continued support for expanding broadband access to rural areas. An alternative approach would be to create community hubs for remote work, or to encourage coworking space companies to set up shop in small communities. Shared hubs could reduce the fixed cost of remote work by providing equipment for rent that may currently be too expensive for individual workers to own. They could also offer IT support and a more reliable internet connection.

Policy #2: Subsidies for Remote Work

It may be desirable for federal, state, and local governments to offer wage subsidies and other incentives for distant firms to hire local remote workers. This would be a micro version of the much larger tax breaks that are used today to try and lure businesses to invest locally. Essentially, the argument for subsidies is that there are positive externalities to remote work, so that it will be practiced at a lower than optimal rate without subsidies (at least initially).

Policy #3: Online Education and Training

Support for online degrees and worker training programs tailored to the needs of remote workers are also desirable. At present, those with the skills to work in the knowledge economy largely reside in cities, since that's where the jobs are. To skill up the kind of people who would most benefit from remote work, we need to offer online degree programs.

Online education has several virtues. First, it can serve as a screen on both employers and employees. The kinds of students who excel at online education are more likely to excel as remote workers too. Moreover, online education gives students an opportunity to try out the remote work lifestyle, while gaining a valuable skill. If it turns out they don't like it, they can transfer to a traditional university. To be effective screens, the courses should practice collaborative team-based pedagogy as much as possible. This would most develop the soft skills for working remotely, and would have the additional benefit of accelerating the development of norms for online collaboration and work.

Policy #4: Fostering Online Communal Spaces

Lastly, we need to support the continued emergence of communal spaces for online socializing and the exchange of ideas. This should occur across a diversity of spaces, to make remote work feasible for the broadest group of people possible. Supporting this work could take the form of grants for research, or experimentation with new forms of online community (analogous to grant support for the arts). At a minimum, the desirability of having places people want to hang out online should factor into the discussion around regulating big tech, which runs many of these platforms.

Covid-19 and Remote Work

But in the short run, the most important push for remote work is the global covid-19 pandemic. I am writing in the midst of a massive natural experiment about the capability of remote work. For the first time ever, nearly everyone capable of working remotely is doing so for an extended period of time. Surely, we will learn a lot about where it works and doesn't soon.

However, it is important to be aware of the ways in which the current shift to remote work diverges from the arguments made above. Firms switching to remote work may learn that the relative productivity of remote workers (*AS* in this paper) is higher than they believed. However, the current transition to remote work is far from ideal. It is rushed, often with minimal training, equipment, or coordination. Many remote workers are also taxed with caring for children who are home from school and daycare (me, for example). Some are ill, or caring for the ill. All of these factors may pull down the productivity of remote workers.

It is also occurring in an environment of extreme economic uncertainty. It may be harder for firms to learn the productivity of workers, given huge shocks to demand for goods and services. At the same time, for workers the experience is likely to feel more socially isolating than it would typically be, since it is occurring simultaneously with generalized social distancing.

Moreover, many of the potential benefits of remote work do not apply in this environment. The current cohort of new remote workers is drawn from the local labor market, with all the attendant costs of office space and colocated wages. Firms are not benefitting from better matches or lower salaries. Neither are workers benefiting from better matches or higher real wages.

However, despite all these caveats, it may still be that remote work becomes significantly more common after the emergency use of it during the pandemic subsides. Some firms will learn that positions can be done remotely, and will begin to hire remotely. Many firms and workers will also emerge from the pandemic with upgraded remote capabilities, ranging from digital infrastructure, to organizational strategies, to workplace norms. Perhaps most importantly, as suggested by equation (1), it may be that there are multiple equilibria in the type of work. Remote work faces a coordination problem: for firms, it's not worth investing in remote infrastructure if there are not many people looking for remote work, and for workers it's not worth looking for remote work if there are few firms offering it. Covid-19 might push us out of this equilibrium, into a new one.

References

Agrawal, Ajay, Alberto Galasso, and Alexander Oettl. 2017. Roads and Innovation. *Review of Economics and Statistics* 99(3): 417-434. doi.org/10.1162/REST_a_00619

Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2013. Digitization and the Contract Labor Market: A Research Agenda. *NBER Working Paper* 19525.

Agrawal, Ajay K., Nicola Lacetera, and Elizabeth Lyons. 2013. Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets? *NBER Working Paper* 18720.

Arora, Asish, Sharon Belenzon, and Honggi Lee. 2018. Reversed Citations and the Localization of Knowledge Spillovers. *Journal of Economic Geography* 18(3): 495-521.

Atkinson, Robert D., Mark Muro, and Jacob Whiton. 2019. The Case for Growth Centers. How to Spread Tech Innovation Across America. Brookings institute.

Autor, David H. 2019. Work of the Past, Work of the Future. *AEA Papers and Proceedings* 109: 1-32.

Bakhshi, Hasan, John Davies, and Juan Mateos-Garcia. 2015. The Net Effect: Using Social Media Data to Understand the Impact of a Conference on Social Networks. *Nesta*.

Balland, Pierre-Alexandre, Cristian Jar-Figueroa, Sergio G. Petralia, Mathieu P.A. Steijn, David L. Rigby, and César A. Hidalgo. 2020. Complex economic activities concentrate in large cities. Nature Human Behavior. https://doi.org/10.1038/s41562-019-0803-3

Bercovitz, Janet, and Maryann Feldman. 2011. The Mechanisms of Collaboration in Inventive Teams: Composition, Social Networks and Geography. Research Policy 40: 81-93.

Berkelaar, Brenda L. 2017. Different ways new information technologies influence conventional organizational practices and employment relationships: The case of cybervetting for personnel selection. *Human Relations* 70(9): 1115-1140.

Bettencourt, Luis M.A., José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B. West. 2007. Growth, innovation, scaling, and the pace of life in cities. *PNAS* 104(7) 7301-7306.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying. 2015. Does Working from Home Work? Evidence from a Chinese Experiment.

Boudreau, Kevin J., Tom Brady, Ina Ganguli, Patrick Gaule, Eva Guinan, Anthony Hollenberg, and Karim R. Lakhani. 2017. A Field Experiment on Search Costs and the Formation of Scientific Collaborations. *The Review of Economics and Statistics* 99(4): 565-576.

Brenčič, Vera, and John B. Norris. 2012. Employers' On-line Recruitment and Screening Practices. *Economic Inquiry* 50(1). https://doi.org/10.1111/j.1465-7295.2010.00324.x

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman. 2015. The Value of Hiring through Employee Referrals. *The Quarterly Journal of Economics* 130(2): 805-839.

Campos, Raquel, Fernanda Leon, and Ben McQuillen. 2018. Lost in the Storm: The Academic Collaborations That Went Missing in Hurricane Issac. *The Economic Journal* 128(610): 995-1018.

Case, Anne, and Angus Deaton. 2020. Deaths of Despair and the Future of Capitalism. Princeon University Press.

Catalini, Christian, Christian Fons-Rosen, and Patrick Gaule. 2018. How Do Travel Costs Shape Collaboration? NBER Working Papers 24780.

Chai, Sen, and Richard B. Freeman. 2019. Temporary Colocation and Collaborative Discovery: Who Confers at Conferences. *Strategic Management Journal* 40(13).

Choudhury, Prithwiraj, Cirrius Foroughi, and Barbara Larson. 2019. Work-from -anywhere: The Productivity Effects of Geographic Flexibility. Harvard Business School Working Paper 19-054.

Coate, Patrick, and Kyle Mangum. 2019. Fast Locations and Slowing Labor Mobility. *Federal Reserve Bank of Philadelphia Working Paper No.* 19-49.

Dahl, Michael S., and Olav Sorenson. 2010a. The Social Attachment to Place. *Social Forces* 89(2): 633-658.

Dahl, Michael S., and Olav Sorenson. 2010b. The Migration of Technical Workers. *Journal of Urban Economics* 67(1): 33-45.

Dingel, Jonathan, and Brent Neiman. 2020. How Many Jobs Can be Done at Home? *Becker Friedman Institute White Paper.*

Dong, Xiaofang, Siqi Zheng, and Matthew E. Kahn. 2018. The Role of Transportation Speed in Facilitating High Skilled Teamwork. NBER Working Paper 24539.

Dunbar, R.I.M. 2016. Do Online Social Media Cut Through the Constraints That Limit the Size of Offline Social Networks? *Royal Society Open Science.* https://doi.org/10.1098/rsos.150292

Duranton, Gilles, and Diego Puga. 2003. Micro-foundations for Urban Agglomeration Economies. NBER Working Paper No. 9931.

Economist. 2019. The Outlook is Dim for Americans without College Degrees. https://www.economist.com/finance-and-economics/2019/01/10/the-outlook-is-dim-for-americans-without-college-degrees

Entertainment Software Association. 2019. 2019 Essential Facts About the Computer and Video Game Industry. https://www.theesa.com/wp-content/uploads/2019/05/ESA_Essential_facts_2019_final.pdf

Faberman, Jason, and Marianna Kudlyak. 2015. What Does Online Job Search Tell Us about the Labor Market? *Economic Perspectives* 40(1).

Fitjar, Rune Dahl, and Andrés Rodriguez-Pose. 2016. Nothing is in the Air. *Growth and Change* 48(1). https://doi.org/10.1111/grow.12161

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2014. Information Technology and the Distribution of Inventive Activity. NBER Working Paper 20036.

Forman, Chris and Nicolas van Zeebroeck. 2012. From Wires to Partners: How the Internet Has Fostered R&D Collaborations Within Firms. *Management Science* 58(8): iv-1592.

Forman, Chris and Nicolas van Zeebroeck. 2019. Digital technology adoption and knowledge flows within firms: Can the Internet overcome geographic and technological distance? *Research Policy* 48(8). https://doi.org/10.1016/j.respol.2018.10.021

Freeman, Richard B., Ina Ganguli, and Raviv Murciano-Goroff. 2014. Why and Wherefore of Increased Scientific Collaboration. Chapter in The Changing Frontier: Rethinking Science and Innovation, NBER: 17-48.

Furman, Jeffrey, Markus Nagler, and Martin Watzinger. 2018. Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program. NBER Working Paper 24660.

Gajendran, Ravi S., and David A. Harrison. 2007. The Good, the Bad, and the Unknown about Telecommuting: Meta-Analysis of Psychological Mediators and Individual Consequences. Journal of Applied Psychology 92(6): 1524-1541.

Gee, Laura K., Jason Jones, and Moira Burke. 2017. Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network. *Journal of Labor Economics* 35(2): 485-518.

Gilrane, Veronia. 2019. Working Together when we're not together. Google blog. https://blog.google/inside-google/working-google/working-together-when-were-not-together/

Goldfarb, Avi, and Catherine Tucker. 2019. Digital Economics. *Journal of Economic Literature* 57(1): 3-43.

Granovetter, Mark S. 1973. The Strength of Weak Ties. *American Journal of Sociology* 78(6): 1360-1380.

Horton, John J. 2016. The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment. *Journal of Labor Economics* 35(2).

Hsieh, Chang-Ta, and Enrico Moretti. 2019. Housing Constraints and Spatial Misallocation. *American Economic Journal: Macroeconomics* 11(2): 1-39.

Iqbal, Mansoor. 2019. LinkedIn Usage and Revenue Statistics (2019). *Business of Apps*.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *The Quarterly Journal of Economics* 108(3): 577-598.

Johnson, Janna E., and Morris M. Kleiner. 2017. Is Occupational Licensing a Barrier to Interstate Migration? *NBER Working Paper* 24107.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. Understanding the Long-Run Decline in Interstate Migration. *International Economic Review* 58(1).

Karahan, Fatih, and Serena Rhee. 2017. Population Aging, Migration Spillovers, and the Decline in Interstate Migration. *Federal Reserve Bank of New York Staff Report No*. 699.

Karis, Demetrios, Daniel Wildman, and Amir Mané. 2016. Improving Remote Collaboration With Video Conferencing and Video Portals. *Human-Computer Interaction* 31(1).

Kim, E. Han, Adair Morse, and Luigi Zingales. 2009. Are elite universities losing their competitive edge? *Journal of Financial Economics* 93(3): 353-381.

Kroft, Kory, and Devin G. Pope. 2014. Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency? Evidence from Craigslist. *Journal of Labor Economics* 32(2): 259-303.

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco. 2020. Patent Citations Reexamined. *The RAND Journal of Economics*. https://doi.org/10.1111/1756-2171.12307

Kuhn, Peter, and Mikal Skuterand. 2004. Internet Job Search and Unemployment Durations. *American Economic Review* 94(1): 218-232.

Kwon, Hyuk-Soo, Jihong Lee, Sokbae Lee, and Ryungha Oh. 2019. Knowledge Spillovers and Patent Citations: Trends in Geographic Localization, 1976-2015. Cenmap working paper CWP58/19.

Laniado, David, Yana Volkovich, Salvatore Scellato, Cecilia Mascolo, and Andreas Kaltenbrunner. 2018. The Impact of Geographic Distance on Online Social Interactions. *Information Systems Frontiers* 20: 1203-1218.

Mas, Alexandre, and Amanda Pallais. 2017. Valuing Alternative Work Arrangements. *American Economic Review* 107(12): 3722-3759.

Mas, Alexandre, and Amanda Pallais. Forthcoming. Alternative Work Arrangements. *Annual Review of Economics*.

McCulloch, Gretchen. 2019. Because Internet: Understanding the New Rules of Language. Riverhead Books.

Mewes, Lars. 2019. Scaling of Atypical Knowledge Combinations in American Metropolitan Areas from 1836 to 2010. *Economic Geography* 95(4). https://doi.org/10.1080/00130095.2019.1567261

Molyneux, Logan, Krishnan Vasudevan, and Homero Gil de Zúñiga. 2015. Gaming Social Capital: Exploring Civic Value in Multiplayer Video Games. *Journal of Computer-Mediated Communication* 20(4): 381-399.

Monteiro, Natalia, Odd Rune Straume, and Marieta Valente. 2019. Does remote work improve or impair firm labour productivity? Longitudinal evidence from Portugal. NIPE Working Paper #14.

Moretti, Enrico. 2013. Real Wage Inequality. *American Economic Journal: Applied Economics* 5(1): 65-103.

Ozimek, Adam. 2019. Overboard on Offshore Fears. *Upwork Report*.

Packalen, Mikko, and Jay Bhattacharya. 2015. Cities and Ideas. *NBER Working Paper No.* 20921.

Pallais, Amanda. 2014. Inefficient Hiring in Entry-Level Labor Markets. *American Economic Review* 104(11): 3565-3599.

Rossi-Hanberg, Esteban, Pierre-Daniel Sarte, and Felipe Schwartzman. 2019. Cognitive Hubs and Spatial Redistribution. *Federal Reserve Bank of Richmond Working Paper Series* 19-16.

Sherman, Eliot L. 2019. Discretionary Remote Working Helps Mothers Without Harming Non-mothers: Evidence from a Field Experiment. *Management Science* doi/10.1287/mnsc.2018.3237

Singleton, David. 2019. Stripe's fifth engineering hub is Remote. *Stripe Blog*. https://stripe.com/blog/remote-hub

Steinkuehler, Constance A., and Dmitri Williams. 2006. Where Everybody Knows Your (Screen) Name: Online Games as "Third Places." *Journal of Computer-Mediated Communication* 11(4).

Stopfer, Juliane M., and Samuel D. Gosling. 2013. Online social networks in the work context. Chapter in The Psychology of Digital Media at Work, eds., Daantje Derks, Arnold Bakker. Routledge publishing.

Thompson, Peter, and Melanie Fox-Kean. 2005. Patent Citations and the Geography of Knowledge Spillovers: A Reassessment. *American Economic Review* 95(1): 450-460.

Trepte, Sabine, Leonard Reinecke, and Keno Juechems. 2012. The Social Side of Gaming: How Playing Online Computer Games Creates Online and Offline Social Support. *Computers in Human Behavior* 28(3): 832-839.

Wilkinson, Will. 2019. The Density Divide: Urbanization, Polarization, and the Populist Backlash. *Niskanen Center Research Paper*.

Wuestman, Mignon L., Jarno Hoekman, and Koen Frenken. 2019. The Geography of Scientific Citations. *Research Policy* 48(7): 1771-1780.

U.S. Bureau of Labor Statistics. 2019. Table 1. Workers who could work at home, did work at home, and were paid for work at home, by selected characteristics, averages for the period 2017-2018. *Economic New Release*.

U.S. National Science Foundation. 2019. Publications Output: U.S. Trends and International Comparisons, Table S5A-32. *Science and Engineering Indicators 2020*.

U.S. Patent and Trademark Office. 2020. Patentsview Website, location_inventor file.