Measuring the Impact of Remote [Knowledge] Work Using Big Data

Alan Kwan, Ben Matthies and Alexander Yuskavage[1]

[EARLY DRAFT - PLEASE DO NOT QUOTE]

Abstract

We develop a framework to measure remote work at the firm level using novel data of the daily internet activity for over 300,000 firms in the United States from 2019 to 2021. We observe whether employee internet activity originates from remote work locations or in-office and measure of the fraction remote work activity for each firm over time. Validating this classification, we document a 30% increase in remote IP traffic in March 2020 at the onset of the crisis and a negative covariance of -0.756 between the share of remote IP traffic in a county and mobile phone data on workplace visits. Next, we study the impact of remote work on firm performance using confidential tax filings data. Instrumenting remote work decisions with firm-level pre-pandemic commuting distance, we document an economically significant rise in output following a shift to remote work. Micro-evidence from employee reading patterns are consistent with a rise in overall productive activity. In the cross-section, we find that such benefits accrue primarily to firms with greater monitoring ability, lower monitoring costs, stronger worker incentives, tradeable sector, and urban firms. We also find evidence that remote work policies provide advantages to firms in the labor market as workers flow to firms which work remotely.

[1]Alan Kwan is with the University of Hong Kong, e-mail: apkwan@hku.hk (corresponding author); Ben Matthies is with the University of Notre Dame, e-mail: bmatthie@nd.edu; Alexander Yuskavage is with the U.S. Department of the Treasury, Office of Tax Analysis, email (alexander.yuskavage@treasury.gov). We thank Sumit Agarwal, Jose Barrero (discussant), Nicholas Bloom, Zhi Da, Steve Davis, Steven Dimmock (discussant), Michael Gibbs, Wei Jiang, Jin Li, Yukun Liu, Abhiroop Mukherjee, Yihui Pan (discussant), Wenlan Qian, Chang Sun, James Feyrer, John Bodian Klopfer, Alexandr Kopytov, Gedeon Lim, Erzo Luttmer, Paul Novosad, Sangyoon Park, Christopher Snyder, Yanhui Wu, and participants at Dartmouth College the Asia Online Corporate Finance Workshop, Hong Kong University, Hong Kong Applied Micro Workshop, NBER Summer Institute (IT and Digitization), National University of Singapore, SFS Cavalcade North America, the Stanford Conference on Remote Work, and the Young Scholars Finance Consortium for insightful comments. We thank the anonymous data provider for support. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors, and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury or the IRS. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository and only accessed by an employee of the Department of the Treasury, and all results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

The COVID crisis spurred the widespread adoption of remote work across the United States. As the pandemic subsides, firms face the decision about whether to permanently retain these remote work policies or to return to the office. Central to this decision is the effect of remote work on firm productivity. In this paper, we study the effect of remote work on firm outcomes using a measure of remote work activity of hundreds of thousands of firms in the United States over time. We link our remote work measure to administrative tax records to provide the first large-scale test of the relationship between remote work adoption and firm performance.

We build our measure of firm-level remote work using the internet activity of individuals in the United States. We collaborate with a technology company in the digital marketing space (the "Data Partner"). Our Data Partner collects over 1 billion observations each day of content consumption across several thousand publishers of major business content and news. The Data Partner helps these publishers to identify the employing firm of website visitors and measures the types of work-related content that each firm reads.[2]

We build upon the firm-linked internet activity dataset and develop a procedure to classify IP addresses, which allows us to determine whether each website visit originates from an office network, or, alternatively, a residential, mobile, or VPN network. Then, through the use of cookies, which are longitudinal identifiers of a browser session, we can track visitors across IP addresses which allows us to identify when an anonymized employee of a firm returns to a publisher and whether they do so from the office or from a remote IP such as their home. Using this data, we construct a firm-level measure of remote work as the percentage of internet activity each month belonging to an remote IP address.

We validate this measure using a battery of tests. First, we observe that the aggregate share of remote (non-office) web traffic in the U.S. increases dramatically at the beginning of March 2020 with an increase of over 30%. We also find that the shift towards remote work is higher

[2]The main purpose of classifying what businesses read is to identify what types of potential products and services the firm may be interested in, which is then sold to to marketing teams to help them identify potential customers.

in Democrat states and in industries that prior studies have identified as suitable for remote work (Dingel and Neiman (2020)). Additionally, we find our measure of remote work is significantly related to widely used mobility measures from the SafeGraph mobility dataset.[3] At the county-week level, the covariance between our remote work measure for workers located in a county and the ratio of mobile devices visiting their common-day-time location, a proxy for work behavior, is -0.756.[4] This number drops outside of work hours.

We also conduct a firm-level analysis of the determinants of the shift toward remote work. Our data contains 300,000 firm observations from March 2020 to December 2021 that meet our sample coverage requirements. We define our baseline measure of remote work as the fraction of internet traffic *during work hours* not at an office IP address for a given firm in a given month. We document an intuitive, positive relationship between our measure of remote work and basic public health risk factors. For example, we find that the prevalence of COVID-19 and mask mandates in a firm's location are positively correlated with a shift toward remote work. Firms situated in more urban and densely populated areas, as well as larger firms, exhibit a proclivity for remote work adoption. These factors persist in influencing remote work decisions in 2021. Overall, aggregate, industry, geographic, and firm-level analyses jointly suggest that our measure of remote work is economically reasonable.

Next, we study the impact of remote work on firm performance. Although there is a building consensus about the effects of remote work at the individual level (e.g. (Bloom, Liang, Roberts and Ying, 2015a),Emanuel and Harrington (2020)), it is far from clear whether this aggregates to the *firm* level. Two commonly discussed considerations made by firms are organizational monitoring costs and labor market positioning. In particular, the implementation of remote work policies can exacerbate the moral hazard problem faced by firms since the distance between managers and employees reduces manager oversight (Gumpert (2018)). The reduction in monitoring allows employees to shirk responsibility and lowers productivity (Alchian and Demsetz (1972); Holmstrom

[3] e.g., Charoenwong, Kwan and Pursiainen (2020);Chiou and Tucker (2020)

[4] We do not expect an elasticity of 1, since we likely capture knowledge workers, whereas mobile phone data may contain a subset of workers who are not knowledge workers.

(1982)), or reduce an employee's altruistic incentives to exert effort for the benefit of their coworkers and by diminishing external peer pressure (Kandel and Lazear (1992)), or prevent increased communication costs and suboptimal coordination across employees at the firm (Garicano (2000); Dessein, Galeotti and Santos (2016)). On the other hand, despite the potential costs of remote work, an often proposed explanation for the persistence of remote work policies is that workers demand them. For example, Barrero, Bloom and Davis (2021) report that 36% of surveyed workers would begin looking for a remote work job if asked to return to the office, with another 6% outright quitting. Other reasons might be commute distance and reduced distraction from interpersonal interactions in the office. Thus, firms may act expecting that they benefit from better workers by committing to a remote work policy, all else equal.

Ultimately, how the costs and benefits of remote work aggregate to firm-level productivity is an empirical question. To approach this question, we study the relationship between changes in remote work and firm output using tax filing data. These tests require exogenous variation in remote work, as firms may choose to engage in remote work according to their productivity or other business fundamentals. Accordingly, we instrument firm remote work decisions using commute distance. Using **pre-pandemic** information from our reading dataset, we can longitudinally observe an anonymized profile associated with a firm traversing across IP addresses.[5] In the first stage, we regress our firm-level measure of remote work on commute distance at the firm's locations in 2019. In the second stage, we regress our outcome variable on the instrumented remote work. In all specifications, we control for lagged remote work because lagged remote work (and thus the shift to remote work) may be correlated with the ex-ante level of remote work, and the lagged ex-ante remote work is also correlated with the ex-ante commute distance.

We posit that commute distance is idiosyncratic and largely predetermined before the COVID-19 pandemic, thus providing an exogenous source of variation in the propensity to work from home. Looking at summary statistics, we find that firms with higher commute distance have simi-

[5]For these profiles, we construct the common day-time location and night-time location associated with the employee, normalized for the location of their common-day-time location (i.e. compare those within Los Angeles to each other). We validate this measure by comparing to SafeGraph's phone distance measures, which are based partly on retail workers and primarily on store visitors (and only available for select establishments in our sample).

lar pre-period profitability. However, there is the potential concern that firms with longer average commute distances may have been differently prepared for remote work. In this case, our identification strategy would capture the local average treatment effect (LATE) - while this estimate may not generalize to all firms, it still provides causal variation on those firms that were better equipped to transition to remote work.

To implement these tests, we define the growth in firm income as the change in either sales receipts or total income from 2019 to 2021, deflated by 2019 assets.[6] In these specifications, we additionally control for lag profitability to mitigate concerns raised in prior studies which have documented a relationship between worker income and commute distance. Our reduced form estimates suggest that firms which engaged in increased remote work over this period saw lower sales growth, while our IV estimates reverse: that greater remote work is associated with significantly higher receipts and total income. A one standard deviation increase in remote work is linked to an increase in receipts of 28-38% and an increase in total income of 23-30%. These findings indicate that, on average, the shift to remote work during the COVID-19 pandemic has positively impacted profitability and top-line revenue growth. These estimates regarding selection and treatment concur with the findings of Bloom et al. (2015a) and Emanuel and Harrington (2020). In our context, the results might imply firms observing negative productivity shocks may choose to work remotely, whereas firms exogenously induced toward remote work do not see reduced productivity.

To complement these results, we examine worker-related activity using the reading dataset. We construct proxy variables for work-related reading. Broadly, the results rationalize the financial performance - we find that the fraction of time allocated toward websites that rarely have work-related content tends to rise when firms engage in remote work, whilst IV estimates suggest differently that such content consumption either drops or stays constant. These findings overall concur with the results from financial performance.

We present cross-sectional tests. First, we show that the increase in firm profitability is strongest subsets where remote work was likely most impactful: tradable industries and firms in urban en-

[6]Similar results obtain using other deflators, as well as simply doing log differences in total receipts

vironments. In addition, we examine two sets of variables corresponding to the monitoring cost proxies we presented earlier. We also examine two monitoring cost proxies: firm size, and the fraction of worker payout relative to total firm profitability or wages. The logic of these proxies is that in the absence of monitoring ability, workers' incentives align them toward firm profitability. We also expect smaller firms to have fewer monitoring costs. Our results generally are directionally consistent with our hypothesis that the effect of remote work is modulated by monitoring costs, although we do not find significant results for all proxies.

In additional cross-sectional tests, we discuss a number of additional tests and mechanisms. We study one factor affecting firms' adoption of remote work: worker demand. We find results, for example, that labor market concentration is negatively associated with remote work, possibly because workers have fewer negotiation options. Firms with more skilled, longer-tenured employees favor remote work, but those with more women and minorities do not, despite their ostensible preference for it, and firms with high remote work attract more workers.

We also present an analysis of labor flows across firms by constructing a firm-pair matrix of worker movements. Using data from Revelio Labs (underlied by LinkedIn), we find workers' latest jobs prior to March 2020, and the last job they had after 2020. We regress net hiring across firms against the difference in remote work practices across the firms. Our results indicate that a one-standard-deviation differential in remote work between firms is associated with a 9% difference in net hiring. This implies that workers flow from low-remote work to high-remote work firms, consistent with an argument about workers preferring to work remotely and firms drawing in workers with remote work. While these findings are only suggestive, they imply labor market positioning and worker demands may be a factor in firms' remote work decisions.

This paper most directly contributes to the expanding literature studying the recent rapid rise of remote work. Most studies of worker performance focus on individuals inside of firms (e.g. Gibbs, Mengel and Siemroth (2021), Bloom, Liang, Roberts and Ying (2015b) and Emanuel and Harrington (2020)).[7] Relative to this literature, we make two contributions. First, we develop a

[7]Also, Ben-Rephael, Carlin, Da and Israelsen (2021) and Ben-Rephael, Carlin, Da and Israelsen (2022) study the work patterns of executives and analysts respectively using logins to Bloomberg accounts.

measurement framework for remote work that contributes the first economy-wide measure of remote work. Second, coupling our unique empirical setting with administrative tax filings data, we provide new facts about the trade-offs and costs of remote work. We deliver the first firm-level estimates on productivity and output. Our approach deviates from prior survey approaches. Surveys are more appropriate for accurate and nuanced information about remote work patterns, but less appropriate for productivity: respondents self-select into answering the survey and require self-assessments of productivity, which can be problematic. However, our results generally concur with micro-level estimates which suggest positive productivity effects of remote work. Our estimates provide the first firm-level causal interpretation of the impact of remote work.

2 Data

2.1 Firm Internet Activity

We partner with a data analytics company from the marketing technology space, the "Data Partner." The Data Partner maintains a large network of partnerships with online publishers, focused primarily (but not exclusively) on business content and news. Contributors include thousands of major internet publishers. Most particpate anonymously but span a wide range of business functions such as technology, finance, marketing, legal, human resources, manufacturing, science, and general business. Participating publishers contribute to the Data Partner's pooled dataset via a technology mechanism, which shares information about web content consumption, including the external IP address of the network originating the HTTP request and the URL of content accessed. Overall, the platform aggregates around 1 billion content consumption events per day. From this large dataset, the Data Partner performs two steps: (1) it associates visitors with companies, when possible, and (2) it quantifies the "topics" of the content visitors read about. From these two steps, the Data Partner produces analytics are primarily sold to companies to facilitate sales and marketing – by identifying companies with heightened research interest in a specific business topic, they may narrow down likely customers. Participating publishers receive some of the Data Partner analytics in return. Our data in this paper are more granular than what is sold commercially, access to the data was designed to take special care with respect to confidentiality restrictions. Further, the Data Partner does not sell these analytics products to financial institutions for the purpose of financial trading.

To associate user sessions with the firms they work for, the Data Partner creates a profile through the use of first- and third-party cookies. This enables the publisher, and in turn The Data Partner, to observe when a visitor returns to a website. Over time, the Data Partner infers the association between the profile and their place of employment (Company) through a wide ensemble of industry-accepted methods. For example, user profiles are associated with a Company when visitors use a work email to log into a participating publisher's website. Another example is through IP addresses. That is, if a profile consistently logs onto a publisher website from a work-associated IP address, this gives a strong association the profile belongs to a particular company. The Data Partner also receives data from third-party sources who perform identity resolution of visitors. Through its proprietary processes, the Data Partner assembles these various sources of data and determines whether a reliable association between a profile and a company can be inferred, and when it can be, what that association is. Crucially, once a visitor has been associated with a Company reliably, the visitor is associated with that Company even though the visitor may traverse different IP addresses. This is the primary mechanism through which we are able to monitor transitions between different types of IP addresses, and thus whether the employee is remote or not. In total, Panel A of Table 1 shows there there are approximately 6.5 million companies observed worldwide in the dataset in any given year and about 1 billion interactions per day.

The Data Partner performs various content classifications. First, the Data Partner flags content for business relevance using a model trained against a hand-picked set of content that is related to sports, entertainment, adult, and other non-work content. After flagging content that is work-related, the Data Partner performs a topic analysis. The topics are business-related topics driven by their own view of important topics or customer requests. For example, to chart interest in Cloud Computing they have assembled a set of 80 to 100 articles that a human has labeled as being pertinent to Cloud Computing. We leverage the content analysis for one empirical test. Panel D of

Table 1 breaks this down by category of website ranking. The website ranks were obtained from Censys, which tracks the Alexa top 1 million on a regular basis. About a quarter of website traffic belongs to the Alexarank top 1,000, and about 25% falls outside of the Alexarank top 1 million.

Panel C of Table 1 lists the top 10 countries by number of website visits. The table also lists the number of unique IPs observed through the first eight months of 2020. The largest source of visitors is the United States, contributing 6.8 billion events per month in 2020, or roughly about one-fifth of the observations. The next three are India, Canada, and the United Kingdom, which have large English-speaking populations. In terms of IP addresses, approximately one-fifth of over 4 billion IPv4 addresses in the world are observed in the data set. Since websites, as well as some servers and internet-connected devices, are assigned IP addresses, the IP addresses we observe likely comprise an even larger fraction of IP addresses primarily used for human content consumption.

Panel E reports the content topic distribution covered by the cooperative. It suggests that while there is a tilt toward technology topics, this may reflect the internet as a whole and it is unclear if this suggests the data over-represents certain types of publishers. However, there is significant balance in topic coverage of non-technology topics: sales and marketing, legal issues, finance, business, human resources, and more.

There are some limitations to intent data worth discussing. One can only observe data that is obtained with consent from publishers. While the content covered in the cooperative covers a broad range of topics, there may be certain types of publishers that show up more often than others. In Panel B of 1, we classify websites based on Internet Advertising Bureau (IAB), which is common in the digital advertising industry. For example, IAB1 refers to Arts and Entertainment, while IAB12 refers to News. We gather IAB classifications from DMoZ, BuiltWith, and Webshrinker. A website can have multiple IAB classifications, so we use the lowest numbered one (e.g., given IAB 1 and IAB 12, I use IAB 1). 21.7% of content can be classified as news, with 8.81% additionally classified as business, 8.41% classified as finance, and 7.79% classified as technologies. 3.39% is classified as Science. This suggests that the largest pluralities of content consumed on this site belong to news and business categories, followed by tech. In addition, we observe the consumption of additional content, such as Arts and Entertainment, Hobbies, Non-Standard Content, and Travel, Sports, and Shopping, which are less likely to be business-relevant.

Appendix Figure A6 shows that firms in different industries generally read fairly similarly, with firms in the tech industry only reading about 20–25% more than other firms on tech-related topics. In other words, in a given firm, most firms tend to acquire information across multiple functional areas. Thus, any skews in the sample are unlikely to lead to a large skew in the coverage of the firms. Second, most of the publishers are from North America. For this reason and reasons related to how we classify IP addresses, we focus our tests on firms with establishments in the United States.

2.2 Other Data

Firm level data Tax data comes from the Internal Revenue Service's Compliance Data Warehouse (CDW). The CDW houses an exhaustive collection of tax returns, which covers both business and individual filings from 1996 to the present. We rely on three tax forms, each of which must be filed annually depending on the business type: Form 1120 for Corporations, Form 1120S for S-Corporations, and Form 1065 for Partnerships. From these tax forms, we obtain two sorts of information: identifying information (including the name, Employer Identification Number (EIN), and mailing address) and tax information (including the various components of income and deductions that the firm reports). Firms were matched to outside data based on their EIN, with a link being considered successful if both sources agreed on the firm name (with an allowance made for small variations) and the firm state. The tax data used include Total Assets (box D on Form 1120), Gross Receipts (line 1a), Total Income (line 11), Cost of Goods Sold (line 2), a measure of total compensation (lines 12, 13, 23, and 24), a measure of the costs used in calculating EBITDA (line 27 minus lines 17, 18, 20, and 21), and Total Deductions (line 27).

Technology The primary dataset of firms' technology comes from the Aberdeen Computer Intelligence Database (CiTDB), otherwise known in the literature as the Harte Hanks Database. It is a dataset of firms and their associated technology usage. Their underlying data provider is the Dunn and Bradstreet database, which underpins the National Establishment Time Series commonly used in the academic literature. For technology, we focus on the "product installs" database which identifies specific software associated with a firm. [8] To obtain ballpark figures for IT budget, Aberdeen began to employ models to increase coverage. Based on emails with Aberdeen, we learned that IT budgets are based on employees, revenues, industry, location, firm/establishment age, and a series of unknown adjustment factors, which likely do take into account actual IT usage behavior observed by Aberdeen. Many of the firm underlying numbers are from Dunn and Bradstreet, who themselves often model the underlying employees and revenue numbers. Therefore, we depart from prior research which has looked at budgets, and instead look directly at software which is both more granular and directly verified.

Restrictions We purchase the dataset used in Spiegel and Tookes (2021). The COVID Restrictions database is a hand-collected database of restrictions inclusive of business and residential mask mandates, gathering restrictions, business closures, and other types of restrictions.[9]

Mobility We use aggregated mobile phone data from SafeGraph, a company producing anonymized mobile phone location statistics covering 10% of U.S. mobile devices. The data include the monthly number of visits at points-of-interest identified by SafeGraph, and for a subset of time produces metrics of how often people who live in a given area stay at home (e.g., socially distance) versus go to work physically. In the U.S., SafeGraph observes 18.75 million devices, approximately 5.6% of the U.S. population and about 10% of mobile devices. According to its analysis of user characteristics, SafeGraph posits that its sample is representative of the U.S. population based on income characteristics, age, and demographics of its users. The data are widely used in studies of social distancing during the COVID-19 pandemic (for example, see Charoenwong et al. (2020),Chiou and Tucker (2020))

[8] After several discussions with Aberdeen, we learned that the primary data collection method for this portion of the Aberdeen dataset is third-party data vendors who leverage data sources as resumes and job postings to identify the technologies firms use. Our interpretation of this measure is therefore that this is the IT that employees use internally. Around 2010, Aberdeen substantially changed its coverage to include more firms. Whereas prior to that, The CiTDB (then Harte Hanks) reported the results of IT surveys sent to establishments, in 2010 its coverage changed tenfold and its approach changed to increase its overall coverage.

[9] It does not include restrictions related to educational / school closures.

Employment data We obtain employment data from Revelio Labs, which obtains employment profiles from various sources such as Linkedin. We obtain employment profiles, user demographics, skills, and demographics for a list of firms we identify as part of this study. Revelio is used by many other studies (e.g., Liang, Lourie, Nekrasov and Yoo (2021);Li, Mai, Shen and Yan (2020)) and more broadly, employment profiles are used across the literature (e.g., Jeffers (2019);Agrawal, Hacamo and Hu (2021)) to measure employment behavior. We also obtain job posting data from Burning Glass. This data has been used variously in studies such as Hershbein and Kahn (2018). The data provides occupational classifications, locations, and skills for firms from a wide variety of job boards.

3 Building a Remote Work Measure

3.1 IP address classification

To make a determination of whether a firm is working remotely, we calculate the fraction of the firm's IP traffic originating from a "remote" IP. We define a "remote IP" as an IP related to a VPN, residence, or mobile IP. We classify all IP addresses observed in the data set from January 2020 to August 2020. Importantly, the classification is based on **pre-pandemic information**. Some IP addresses show up later on without pre-pandemic information in our dataset, but are generally a very small fraction of total internet activity. Implicitly, the assumption we make is that the classification of an IP address to a particular type remains stable during the pandemic, or that any change in the usage of an IP address is not systematic in a way that introduces bias. [10]

We build our own classification rather than buy because commercial solutions are not suitable. They are expensive, tend not to provide historical information, and their methodologies are not disclosed. We develop two classification methods: a rule-of-thumb classification and a simple machine-learning based classification model. While we only need to identify office IP addresses, it

[10]One might be concerned that large shifts in internet traffic caused by remote work might lead to changes in the configuration of IP addresses. While IP addresses do get re-assigned regularly to different owners, it is not obvious why businesses hoping to set up infrastructure for remote workers would need to purchase multiple new IP addresses. IP addresses have no limit to the number of users and services that can operate from them.

is helpful to further classify IP addresses to ensure the IP address classifications are reasonable. For example, we would expect residential IP addresses to look very different from mobile IP addresses.

First, in the rule-of-thumb approach, we collect several indicators that imply an IP belongs to a particular type. The rule-of-thumb approach covers approximately one-fifth of the IPs but over 60% of traffic. An example of a rule-of-thumb is that many IP address have domain name service (DNS) records revealing a network name, such as *vpn.examplewebsite.com*. One can likely assume that any such site is a VPN.

After classifying IP addresses via rule-of-thumb, we classify the remainder using a simple machine learning model. We construct a ground truth dataset of 10,000 IP addresses. [11] To classify IP addresses, we utilize five sets of features: (1) port-scans, whether specific ports are open on the device, (2) content-based filters, which ask whether the IP address consumes more work-related content, (3) scale features, which ask what is the scale of reading activity at the IP address in terms of article reads and browsers observed (4) temporal features, defined as the fraction of activity which occurs during work hours, and (5) geographic features, which characterize the *geographic distribution* of profiles. For example, if an IP address visitors are associated with several different states or countries, one can reasonably infer it is a mobile network or VPN and not a single household. The majority of these features are based on the reading data set, but port scans and DNS records are collected by third-party cybersecurity firms (Censys and Rapid 7). We describe these details in Appendix A.1. We also display summary statistics of the labeled data sets in Appendix Table A13.

We sample down to reduce imbalance in the panel and execute a random forest model against the data, splitting it 70-30. We report additional machine learning metrics in A. The classification

[11] For residential labels, we use Timewarner and Comcast IP address blocks. For business addresses, the data provider has a set of known IP addresses from the Fortune 1000 which they certify are corporate IPs. For mobile networks, we define mobile networks as US mobile address networks where at least half of the observed events in the Event data set in 2019 were mobile, Maxmind's connection-type database lists it as a mobile connection, and the reverse RDNS entry in February 2020 indicated it was part of the mobile network of a major U.S. telecom (e.g., AT&T). Fourth, we classify VPNs of two types: residential proxies, which tend to service IP addresses from all of the world across a wide variety of users (who may originate from different companies), and business VPNs, which are generally used by members of one company. Empirically, these two types of VPNs will have different behaviors, which we separately model.

accuracy across five classes using a random forest model is around 90%, with high sensitivity and specificity for all classes except for a relatively lower but passable sensitivity of 0.73 for business VPNs. In other words, while the model will accurately classify an IP address that is not belonging to a VPN, it will be less accurate identifying VPNs.[12]

3.2 Summary statistics by IP address type

In Table 2 we report summary statistics of IP addresses *after* we classify them based on their activity prior to the pandemic from January 2019 to February 2020. In Appendix section A, we present additional summary statistics for the interested reader. The summary statistics here are based on IPs we have labeled after running our classifications. On the set of cleanly labeled data, we present additional summary statistics that are inputs into the machine learning model on *labeled*. The statistics are more detailed and are meant to illustrate the typical differences in characteristics between mobile, residential, VPN, and business IPs. We also explore machine learning accuracy and performance metrics.

Based on our classification, the vast majority of IP addresses are most likely residential IP addresses. Of approximately 760 million IP addresses observed, 13 billion events per month in 2020 originate from 474 million residential IP addresses. The next largest category is mobile phones, followed by businesses. VPNs are the smallest category, with a small fraction of content coming from them.

The summary statistics of behavior observed on different IP addresses is reasonable. While residential IPs are the largest share of IP addresses we observe, they are the least active. For each residential IP, we observe roughly 36 events per month. Other types of IP addresses, conversely, VPNs – which are likely shared by multiple users – have double the events per IP. Meanwhile, business IPs are even larger. This is sensible given that business establishments could be much larger than households.

[12]We discuss additional features of our classification in Section A.4 in the Appendix.

4 Validation Tests of our Remote Work Measure

4.1 Time-series

First, we examine the time-series pattern in the composition of traffic across IP addresses of different types at the onset of the pandemic. Figure 1 plots the fraction of total traffic in our sample from what we classify to be a remote IP (mobile, VPN, or residential). The left-out group is an office IP. We observe a sharp jump in mobile traffic from 60% prior to the pandemic to just under 80% in April. Although survey estimates suggest that around 1 in 10 workers worked remotely prior to the pandemic (Bloom (2020), the pre-pandemic baseline in our sample is approximately 60%. The difference likely arises because of two reasons. First, not all internet traffic originates from a business. For example, retirees, students, or workers on vacation may be accessing the internet. Second, even prior to the pandemic, much of work-hours internet traffic may have originated from mobile phones. These would be considered remote both prior to and during the pandemic. While this means we cannot measure the absolute fraction of remote work, we are able to measure the change in remote work. The fact that the gap closes by half resonates with statistics that approximately half of Americans were working from home during 2020 Bloom (2020). However, given that we cannot measure the ”level” of remote work accurately, our econometric tests include two-way fixed effects to help account for differences in the level of remote work that exist prior to the pandemic.

Next, in Figure 2, we present the same plot decomposed by IP type: residential IPs, VPNs, businesses, and mobile networks. We normalize the within-class totals to be comparable, otherwise residential and mobile traffic would outsize other classes. The surges apparent in Figure 1 correspond primarily to drops in the business sector, particularly small businesses, as well as rises in the residential IP addresses. There is a temporary blip in mobile IPs, but mobile phones generally hover around their pre-pandemic mean. As noted before, mobile phones might be used during work hours, thus obfuscating the location of the worker. VPNs also surge, but as noted before, summary statistics reveal that VPNs are a small percentage of total internet traffic.

4.2 Geography

First, we examine the shift to remote work by state. Figure 3 presents a heatmap, where colors represent the average percentage change in remote work from the baseline period (December 2019 to February 2020) to a comparison period at the start of the pandemic (March 2020 to May 2020). Firms are weighted based on the number of reading events in the baseline period (a firm that read 10,000 articles has twice the weight of a firm which read 5,000). The heatmap is colored with blue representing a larger shift toward remote work and red representing a relatively low shift to remote work. Areas with relatively high shifts to remote work include the northeastern United States as well as California. Texas, which Forbes acclaimed as one of the states with companies friendly to remote work or working from home, is slightly above average among states. [13]

Next, we benchmark our remote work measures against mobile phone data from SafeGraph, which is the most popular dataset to study mobility during the pandemic (e.g. Charoenwong et al. (2020)) We conduct a county-week level test whereby we correlate the fraction of internet traffic originating from a remote IP address to the percentage of mobile phones based in that county engaging in social-distancing behavior or work-like behavior. We estimate the following specification:

$$\%Traffic^{remoteORVPN} = \alpha_{county} + \alpha_{week} + \beta_1 MobiityMeasure_{county,week} + + \varepsilon_{county,week}$$

The specification employs county and week fixed effects. Therefore, it can be interpreted as how the change in the level of internet traffic *within* a specific county to a remote IP or VPN varies with changes in mobility in the same week. We define mobility as the percentage of devices engaging in what SafeGraph work-like behavior. Per industry standard, SafeGraph defines work-like behavior as visiting one's "common day-time location," a place where a device is typically observed during 8am to 4pm on a workday. For most people this is thought to be their workplace. Secondly, we define social distancing as $\frac{PhonesCompletelyAtHome}{DeviceCount}$. We would expect the percentage of

[13] Please see: https://www.forbes.com/sites/alexandratalty/2020/06/26/work-from-home-california-texas-named-as-best-states-for-remote-work/?sh=4a27684c383c

traffic on a Remote IP to be positively related to social distancing and negatively related to work-like behavior. In addition, we separately conduct the above tests on VPNs to assess whether VPNs could drive the results.

Table 3 reports the results of the specification. Column 1 shows that the remote traffic measure covaries with social distancing with a coefficient of 0.207. This can be interpreted as an elasticity given that both measures range from 0 to 1. That is, a one percentage increase in social distancing is positively correlated with a 20.7 basis point increase in the fraction of internet traffic coming remotely. Column 2 reports ”work-like behavior” with the sign of the two coefficients positive as expected. The elasticity of remote traffic during work hours to ”work-like behavior” displayed through mobile phones is -0.756. This suggests the elasticity of remote IP traffic is almost, but not quite, one for one, which is reasonable.[14] Further, comparing the coefficient on *Social Distancing* to the coefficient *Worklike Behavior* suggests that social distancing has a smaller relationship than worklike behavior. This is sensible – if one is not at work, they may or may not be social distancing. Also, many people socially distancing might be unemployed adults, teenagers, or elderly and not be associated with a workplace. Column 3 pits these two measures together, showing that as expected, our measure of remote traffic to business sites is more strongly correlated with workplace attendance than social distancing. Column 4 examines VPN-traffic. The coefficient on *Social distancing* and *Worklike Behavior* suggest that there is a rise in VPN traffic when people engage in social distancing. However, the economic magnitudes are smal. This suggests that the VPNs we observe in the data are economically small in relation to remote work in general. Columns 5 and 6 study the night-time. This test serves as a placebo: the coefficient drops from drops from -0.756 to about one-third. It does not drop to zero as people may work weekends or otherwise work outside the 9-4 p.m. work hours tracked by SafeGraph.

[14]It is reasonable to observe an elasticity below 1 for two reasons. First, there is noise in both measures, which introduces attenuation bias. Second, mobile data may capture workplace attendance by people who are not knowledge workers, which our dataset is biased toward. A factory or restaurant worker, for example, may have a mobile phone, and thus be observed as going to work, but read little during the day. Therefore, the economic magnitudes are sensibly large.

4.3 Industry

At the industry level, Table ,4 displays the shift to remote work across industries. We focus on knowledge workers within firms, so we do not expect to quantitatively capture the full extent of differences in remote work across industries. However, knowledge workers with colleagues that remote work less (e.g., retail or factory workers) should also be less likely to engae in remote work. We find patterns largely consistent with prior research such as Dingel and Neiman (2020) have found. Agriculture, retail, and labor-intensive industries rank lowest in remote work uptake, while knowledge-based industries rank highest.

5 Determinants and Consequences of Remote Work

5.1 Empirical Design

In this section, we study both the determinants and consequences of remote work. In terms of sample selection, we include all firm-months that satisfy the following criteria: (1) in a given month, have 100 work-time observations (9 a.m. to 5 p.m.), (2) average at least 1,000 observations per month whether on the weekend, or weekday, during work hours, or otherwise, for at least half of months from 2019 until February 2020, and half of months from March 2020 to end of December 2021, and (3) included in a firmographic database such as Axiomatic's Form 5500 database or the Aberdeen Computer Intelligence database or others where we observe the company domain (our main identifier from the remote work data) and firm characteristics such as business name, address and industry. This is because for our main tests, we need to use name and address or EIN to link to the Social Security Administration database.

To measure remote work ias defined above, $Remote_{it}$ as the fraction of internet traffic during work hours that is not at an office IP address where i indexes firms and t indexes months. Our main focus in the **change** in remote work from 2019 to 2020/2021. While remote work may lead organizations to work outside of the normal workday, internet traffic associated with a firm during work hours is more likely to be representative of where employees are working than activity outside of work hours. For each firm i, we construct $\Delta Remote_i$ as the average fraction of internet traffic during work hours that is not at an office IP address for the firm from April 2020 to December 2020 minus the fraction of remote traffic during 2019.

$$\Delta Remote_i = \eta_j + \sum_{\forall k} \beta_k X_k + + \varepsilon_i$$

where i indexes firms, η_j is a set of industry fixed effects and state fixed effects. We also conduct analysis where remote work is an explanatory variable, with the following specification:

$$Y_f = \eta_j + \Delta Remote_i + \sum_{\forall k} \beta_k X_k + + \varepsilon_i$$

The obvious problem with the above empirical specification is that the choice of remote work could be endogenous to unobserved business circumstances. In the case of, for example, firm output, firms may observe their expected profitability and use this information to determine whether it is optimal to go to the office. To instrument remote work, we employ commute distance based on employee-level internet reading activity. The intuition of the commute distance measure is that employees with longer commute times face higher costs and hassles associated with commuting to work, so they are more likely to adopt remote work options if available. During the pandemic, when commute distance is de-stigmatiized, we posit that thos who have a higher commute distance are likelier to experiment with remote work, all else equal.

To construct the remote work measure, we leverage the dataset on internet activity. Within each firm, we observe individual user sessions: the internet activity of individual employees in the firm. We observe the approximate location of each worker's IP address during non-working hours and the location during working hours (Monday through Friday between 9 a.m. and 5 p.m.). This allows us to calculate the approximate commute distance from each employee's home to the office for each firm in the United States.[15] We calculate the haversine distance between the approximate home location and the office location for each user session in each firm during 2019. For each firm, we compute the 25th and 75th percentiles of the commute distance across all sessions during 2019.

[15] We do not observe any personal identifying information about any employee. We also do not observe precise locations of employee residences – we observe the approximate neighborhood of each employee.

We calculate the average commute distance of the middle 50th percentiles to obtain $commute_i$, the average commute distance of employees at firm i. We describe and validate this commute distance measure in Appendix A.5 against the SafeGraph dataset.

Commute distance has a few threats to identification. First, commute distance varies by geography. However, as we normalize the commute distance measure *within geography*, we effectively remove this as a source of variation. Second, commute distance could be correlated with worker characteristics, which in turn could lead to differential exposure to demand shocks during the pandemic. We control for lag remote work, which captures the differential ability to remote work prior to the pandemic, as well as the ex-ante commute distance, as well as firm size. This mitigates the correlation between commute distance and firm types. In addition, in Table A16 we present statistics of the difference between high and low commute distance firms. Firms that remote work more had similar pay, but were smaller and had better IT all else equal. They had similar income-payout ratios, which is a key statistic we focus on in later analyses.

5.2 Baseline determinants of remote work

We begin by analyzing certain determinants of remote work. For key variables, we present summary statistics in Table A17. Each column in Table 5 presents the results of a different specification of the regression above. We report the coefficient estimates with standard errors in parentheses below. Column 1 presents the univariate specification with COVID as the explanatory variable. The coefficient estimate on COVID is positive and significant at the 10% level, indicating that higher COVID cases per capita in the firm location is associated with an increase in firm remote work. Column 2 interacts Business closures with COVID. It suggests the effect of COVID is no longer significant when interacting with Business Closures, suggesting the shift toward remote work from COVID is driven by changes in workplace policy.

Column 3 presents results including additional sets of covariates. The coefficients on Population Density, Urban percent, and size are positive and each significant at the 1% level. These variables subsume the role of COVID itself and are intuitive: higher population density and urbanization lead to more remote work. Column 4 adds the variable Commute Distance. The coefficient estimate on Commute Distance is positive and significant at the 1% level. Firms whose employees ex-ante worked farther from the office are more likely to work remotely in 2020. This finding is consistent with spatial models of work-from-home which posit that firms whose employees face longer commutes have a greater benefit to remote work which increases the likelihood of implementing a work-from-home policy. The coefficient estimate on prior remote work is negative and significant at the 1% level. Column 5 repeats the analysis of Column 4 but changes the outcome variable to be the shift in remote work to 2021. It shows that many of the features we analyzed previously have a persistent and economically similar effect as during 2020, implying that some of the concerns about public health factors or commute distance remain significant beyond the immediate onset of the pandemic.

5.3 Coordination Costs and Moral Hazard

A central concern surrounding remote work is the potential for these policies to exacerbate moral hazard by reducing employee monitoring. Remote work creates distance between managers and employees which reduces manager oversight on day-to-day employee activities, inducing slacking (Alchian and Demsetz (1972); Holmstrom (1982)). Also, remote work precludes face-to-face interaction which can reduce coordination, making even hard-working employees less effective(Garicano (2000); Dessein et al. (2016)). Accordingly, we study factors associated with coordination costs and moral hazard by constructing measures of firm information technology (IT) and managerial layers. Information technologies may reduce the coordination costs of remote work by providing collaboration tools such as video or audio communication, workflow delegation and organization such as Slack, and the ability to more easily share data and files. They may also reduce employee slacking by providing stronger monitoring of employee efforts while working from home. We use data from Aberdeen to construct several measures of information technology prevalence at each firm and the change from 2019 to 2020. We first measure $\frac{Laptops}{PC}$, which is the ratio of laptops (presumably mobile devices) to total computers. Second, we the count of the number of unique technologies employed by the firm, "$ITindex$", and similarly construct "$\Delta ITIndex$" as the change in "IT index" from 2019 to 2020. Third, we measure the software budget per employee.

Panel A of Table 5 presents our results. Column 1 presents results from the specification running the change in remote work from 2019 to 2021 on $\frac{Laptops}{PC}$. We find the intuitive result that a greater fraction of laptops is correlated with more remote work. In Column 2, we the count of the number of unique technologies employed by the firm "$ITindex$", and "$\Delta ITIndex$" the change in "IT index" from 2019 to 2020. The coefficient on IT Index is positive and significant at the 1% level, indicating that firms with more existing information technology infrastructure implement more remote work policies at the onset of the crisis. The coefficient on the change in the IT index is also positive and significant at the 1% level, indicating a positive association between the adoption of new information technology and transitioning to more remote work. Column 4 studies software budget. We find that this alternative measure of IT strongly correlates to increased remote work.

Next, we examine the relationship between remote work activity and firm managerial structure. Our analysis is motivated by studies of monitoring such as Gumpert (2018) who show empirically that firms with establishments further from headquarters staff more managers at these locations. Given concerns about employee slacking, firms with stronger management structures would be more likely to implement work-from-home policies. To test this idea, we construct two measures of the ratio of managers to employees at each firm: the number of job postings for managers versus the total number of job postings based on BurningGlass data for each firm in 2019; and the number of employees in management roles compared with the total number of employees based on LinkedIn data in 2019 for each firm. We run the following specification:

$$\Delta Remote_i = \eta_j + \beta Manager_i + Controls + \varepsilon_i$$

Where $\Delta Remote_i$ is the change in the fraction of remote internet activity for firm i between 2019 and 2020, η_j are fixed effects for industry j, and $Manager_i$ is a measure of manager to worker ratio. Panel C of Table 5 presents the results of these specifications. Columns 1 and 2 present results using the Burning Glass manager ratio with and without controls. The coefficient on the manager ratio variable is positive and significant at the 5% level or higher in both specifications. Similarly, the coefficient estimates on the LinkedIn manager ratio in Columns 3 and 4 are also positive and significant at the 1% level. These findings indicate that firms with greater managerial layers tend to engage in more remote work, whether measured by their pre-pandemic hiring patterns (inferred from Burning Glass job postings for managerial workers) or by pre-pandemic share of LinkedIn employees in manager roles. This finding is consistent with the idea that firms with a high ratio of managers to employees are better able to monitor employee effort, which alleviates the slacking concerns that arise when firms switch to remote work.

In the Appendix, we study firm investments to job postings and investments in technology as outcome variable and reduce moral hazard by increasing employee monitoring. We find evidence overall supportive of the notion that firms make investments to ameliorate these monitoring frictions. Taken together, our results affirm theories of monitoring costs: unambiguously firms with greater ability to monitor engage in more remote work. Firms appear to invest in monitoring capabilities alongside remote work.

5.4 Firm Level Financial Performance

5.4.1 Main effect

Table 7 presents the results of our empirical analysis, which uses firms as the unit of observation and considers two dependent variables: changes in Receipts and Total Income. Total income accounts for costs. Both 2021 and 2019 are scaled by pre-period assets. Columns 1-2 and 5-6 present OLS findings, while 3-4 and 7-8 present instrumental variable estimates. These regressions indicate a strong relationship overall between remote work and income, once accounting for the selection into remote work. OLS results suggest that firms that tend to engage in remote work tend to be less productive, but accounting for endogeneity to remote work uncovers a positive relationship with remote work. The OLS results are not significant for OLS and total income, but are for gross-sales. However, they are positively significant for the instrumented remote work variable. These results might suggest that firms that choose to engage in remote work observe negative productivity shocks - that is, when one is less busy, they choose not to be in the office. However, the marginal firm being forced into remote work became more productive, all else equal, over this period. The estimates in Column (4) suggests a one standard deviation increase in remote work corresponds to roughly a 9% increase relative to the residual standard error of the outcome variable. The magnitude in Column (8) is very similar. The notion of selection and remote work resonates with findings from microstudies of individuals, whereby individuals tend to be more productive when engaging in remote work, but perhaps due to various factors such as career concerns, these more productive workers tend to *choose* working from the office (e.g.Emanuel and Harrington (2020)).

As robustness to choice of scaling, in Table A18 we present an alternative specification of the outcome variable, which provides qualitatively similar results. This suggests our conclusions are not sensitive to the choice of deflator.

5.4.2 Micro-evidence from Internet Activity

To provide complementary evidence of our firm-level financial data, we study reading patterns within the firm, focusing on the relationship between remote work and various measures of reading behavior. In broad strokes, the reading patterns comprise evidence which broadly speaking rationalizes the findings from the firm-level financial data. Like before, we focus on the difference between our reduced form and instrumental variable estimates. The outcome variable is the fraction of reading on workdays devoted to non-work-related content, as defined by the characteristics of the websites. Appendix Table A19 suggests that the fraction of business-relevant reading is positively correlated to firm sales and spending, unconditionally, suggesting it is a potential proxy for work-related activity. The standard deviation of the Remote Work variable is 0.175.

Columns (1) and (2) consider work-related content as reading from websites with less than 1% of work-relevant content (corresponding to 49.5% of content consumption events), while Columns (3) and (4) define work-related content as websites with less than 5% work-relevant content (corresponding to 60% of content consumption events). Our findings indicate that there is a significant relationship between remote work and reading behavior, but this relationship depends on the reduced form versus instrumental variable estimates. When firms choose to work remotely, we see an increase in the fraction of reading devoted to non-work-related content with less than 1% of work-relevant content by 0.175 * 0.0986 = 0.01725, or 1.73% (Column 1) or 0.175 * 0.1083 = 0.01895, or 1.90% (Column 3). The instrumental variable estimates of Column (2) reveals a decrease by 0.175 * (-0.0801) = -0.01402, or 1.40% (Column 2), and a decrease by 0.175 * (-0.0211) = -0.00369, or 0.37% (Column 4).

5.4.3 Cross-sectional Analysis: Demographic Factors

As a sanity check, we first examine the tradable and non-tradable designations from Mian and Sufi (2012). We expanded their classification as their notion of tradable and non-tradable corresponds to 1/6th of firms in our sample, but we find similar results restricting to only those they covered.[16] The results suggest the impact of the shift toward remote is absent for non-tradable firms and positive for tradable firms, in line with expectation.[17] We also examine rural versus urban. The majority of firms in our sample is urban. Correspondingly, it is these areas where the positive effects from remote work occur.

5.4.4 Cross-sectional Analysis: Monitoring Costs

We examine two sets of variables corresponding to the monitoring cost proxies we presented earlier in Table 6. We also examine two monitoring cost proxies: firm size, and the fraction of worker payout relative to total firm profitability or wages. The logic of these proxies is that in the absence of monitoring ability, workers' incentives align them toward firm profitability. We also expect smaller firms to have fewer monitoring costs. Our results generally are directionally consistent with our hypothesis that the effect of remote work is modulated by monitoring costs, although we do not find significant results for all proxies.

IT and Managerial Layers: We study the software budget per employee and fraction of employees who hold managerial titles as proxies for monitoring cost, consistent with our prior analysis. Our analysis of IT and managers reveals moderate but negative relationship between moni-

[16] To augment the rest, we used ChatGPT4 to classify them as tradable or non-tradable, and then manually adjudicated the classification.

[17] When using logs instead of financial ratios, we find the non-tradable firms to actually be negatively affected.

toring costs sand remote work. In Columns (1)-(4) we present the low monitoring cost samples, and we find that in these samples, the effect of remote work is large. It is significant for IT but not for managers. In Columns (5)-(8) the relationships between remote work and output in the low-IT sub-samples is insignificant and in some cases marginally negative. While our proxies show some noise, it is clear from these results that the results generally point toward higher monitoring ability leading to a greater effectiveness of remote work. In terms of sensitivity analysis, our results are stronger if we withhold controls for size, which is sensible because the relationship between IT and size, and manager ratios and size, might be correlated. Our analysis using other proxies for IT and job postings as a proxy for managers shows similar but weaker relationships.

High versus low-incentives: In the absence of the ability to monitor, firms may rely on financial incentives as a substitute for monitoring. We obtain two measures from the tax data. One advantage of calculating these from the tax data is that we do not need to rely on noisy external datasets with sub-optimal coverage. The first measure is the fraction of wage over income, which is a payout ratio as a percentage of the firm's overall gross profit. We find similar results look at wage over sales. We also look at employee pay relative to total pay. In our split sample analyses, we find that high-incentive (high employee payout) comprise the entirety of the effect. The results are noticeably absent in firms with relatively low incentives.

Size: Firms with smaller size may have lower monitoring costs, all else equal. We split the sample based on size. We find that the positive effects of remote work are clearly stronger for small firms, although they are not absent for large firms. Looking at columns (1) and (5), we see the impact of shifting toward remote work is double that of large firms in terms of gross receipts. However, for total income, we find that overall profitability does not significantly increase for large firms.

Overall, our cross-sectional results point to intuitive variation across sub-samples. Small firms in urban areas and in tradable sectors should benefit more from remote work, given that it is precisely in urban areas that commute distance is greater, and in tradable sectors that remote work is feasible. While there are reasons to believe both large and small firms could benefit from remote work, it appears as though in the short-run small firms benefit more, owing perhaps due to the relative ease with which coordination happens inside small firms.

5.5 Other Analyses

In Appendix section C.3, we present results examining the relationship between remote work and labor markets. Utilizing metrics like labor market tightness and Herfindahl-Hirschman indices, we find a negative correlation between labor market concentration and remote work (Table A20). Our analysis of employee demographics reveals that firms with more skilled and longer-tenured employees are likelier to adopt remote work (Table A21). Furthermore, we observe a significant flow of labor from firms with low remote-work adoption to those with high adoption (Table A22). Overall, our findings indicate that employees with greater labor market leverage have more favorable remote work opportunities, and firms gain workers if they engage in remote work.

6 Conclusion

In this paper, we develop a new methodology for measuring remote work activity. Collaborating with a data analytics firm from the digital marketing space that has visibility into 20% of IP addresses on the internet, we form a classification of internet networks around the world into remote IPs versus office IPs. We use this measure to document a set of stylized facts about the adoption of work-from-home across firms in the United States. We find a sharp increase in remote work at the start of the COVID crisis, which persists years later. We then study the drivers of remote work decisions as well as the impact of remote work on measures of firm productivity. We join the data with confidential tax filings and utilize an identification strategy based on the commute distance of employees at each firm. We document that an increase in instrumented remote work is associated with an increase in several measures of firm performance and profitability. Using data on the websites visited by employees and patterns of employee internet activity we find micro-evidence consistent with the firm-level increase in productivity. We also that remote work is associated with firm investments into IT tools and increased managerial hiring, indicating the importance of monitoring.

While our findings suggest the baseline effect of remote work on firm level output, a number of interesting questions remain. First is whether remote work - for all of its top-line financial benefit, poses any cost to the firm in terms of building up culture or innovation, which are long-term drivers of firm value. Second, there is more to explore regarding the heterogeneity of which firms benefit and lose from remote work. Finally, even if the average effects of remote work reflect a large effect of remote work on firm output, a reasonable question is how these effects aggregate: do some firms win and lose, off-setting all gains, and how much of that loss is merely re-allocation. We plan to study many of these questions in upcoming future work.

References

Agrawal, Ashwini, Isaac Hacamo, and Zhongchen Hu, "Information dispersion across employees and stock returns," *The Review of Financial Studies*, 2021, *34* (10), 4785–4831.

Alchian, Armen A and Harold Demsetz, "Production, information costs, and economic organization," *The American economic review*, 1972, *62* (5), 777–795.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor market concentration," *Journal of Human Resources*, 2022, *57* (S), S167–S199.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, "Let me work from home, or I will find another job," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2021, (2021-87).

Ben-Rephael, Azi, Bruce Carlin, Zhi Da, and Ryan D Israelsen, "Analyzing the Analysts: Evidence from their Bloomberg Usage," *Available at SSRN*, 2022.

_ , Bruce I Carlin, Zhi Da, and Ryan D Israelsen, "Uncovering the hidden effort problem," Technical Report, National Bureau of Economic Research 2021.

Bloom, Nicholas, "How working from home works out," *Institute for Economic Policy Research (SIEPR). Policy Brief June*, 2020.

_ , James Liang, John Roberts, and Zhichun Jenny Ying, "Does working from home work? Evidence from a Chinese experiment," *The Quarterly Journal of Economics*, 2015, *130* (1), 165–218.

_ , _ , _ , and _ , "Does working from home work? Evidence from a Chinese experiment," *The Quarterly Journal of Economics*, 2015, *130* (1), 165–218.

Charoenwong, Ben, Alan Kwan, and Vesa Pursiainen, "Social connections with COVID-19–affected areas increase compliance with mobility restrictions," *Science advances*, 2020, *6* (47), eabc3054.

Chiou, Lesley and Catherine Tucker, "Social distancing, internet access and inequality," Technical Report, National Bureau of Economic Research 2020.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, "In search of fundamentals," in "AFA 2012 Chicago Meetings Paper" 2011.

Dessein, Wouter, Andrea Galeotti, and Tano Santos, "Rational inattention and organizational focus," *American Economic Review*, 2016, *106* (6), 1522–36.

Dingel, Jonathan I and Brent Neiman, "How many jobs can be done at home?," *Journal of Public Economics*, 2020, *189*, 104235.

Emanuel, Natalia and Emma Harrington, ""Working" remotely?: Selection, treatment and the market provision remote work"," Technical Report, Harvard mimeo 2020.

Garicano, Luis, "Hierarchies and the Organization of Knowledge in Production," *Journal of political economy*, 2000, *108* (5), 874–904.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth, "Work from Home & Productivity: Evidence from Personnel & Analytics Data on IT Professionals," *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2021, (2021-56).

Gumpert, Anna, "The organization of knowledge in multinational firms," *Journal of the European Economic Association*, 2018, *16* (6), 1929–1976.

Hershbein, Brad and Lisa B Kahn, "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings," *American Economic Review*, 2018, *108* (7), 1737–72.

Holmstrom, Bengt, "Moral hazard in teams," *The Bell journal of economics*, 1982, pp. 324–340.

Huang, Peng, Nicholas H Lurie, and Sabyasachi Mitra, "Searching for experience on the web: An empirical examination of consumer behavior for search and experience goods," *Journal of marketing*, 2009, *73* (2), 55–69.

Jeffers, Jessica, "The impact of restricting labor mobility on corporate investment and entrepreneurship," *Available at SSRN 3040393*, 2019.

Kandel, Eugene and Edward P Lazear, "Peer pressure and partnerships," *Journal of political Economy*, 1992, *100* (4), 801–817.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan, "Measuring corporate culture using machine learning," *The Review of Financial Studies*, 2020.

Liang, Chuchu, Ben Lourie, Alexander Nekrasov, and Il Sun Yoo, "Voluntary Disclosure of Workforce Gender Diversity," *Available at SSRN*, 2021.

Mi, Xianghang, Xuan Feng, Xiaojing Liao, Baojun Liu, XiaoFeng Wang, Feng Qian, Zhou Li, Sumayah Alrwais, Limin Sun, and Ying Liu, "Resident Evil: Understanding residential ip proxy as a dark service," 2019.

Mian, Atif R and Amir Sufi, "What explains high unemployment? The aggregate demand channel," Technical Report, National Bureau of Economic Research 2012.

Spiegel, Matthew and Heather Tookes, "Business restrictions and Covid-19 fatalities," *The Review of Financial Studies*, 2021, *34* (11), 5266–5308.

Figure 1: Reading activity during work hours versus night-time hours

In this graph, we plot the number of work-time (Monday-Friday 9 a.m. to 6 p.m. local time) internet reads by type of IP address during the year 2020. Our classification for IP addresses is described in section 3.1, involving a mixture of hand-classification and a machine learning algorithm. Specifically, we collapse residential and mobile IP addresses together and plot the share of aggregate reading activity coming from these types of IP addresses during night and work hours. Internet traffic in this graph is global.

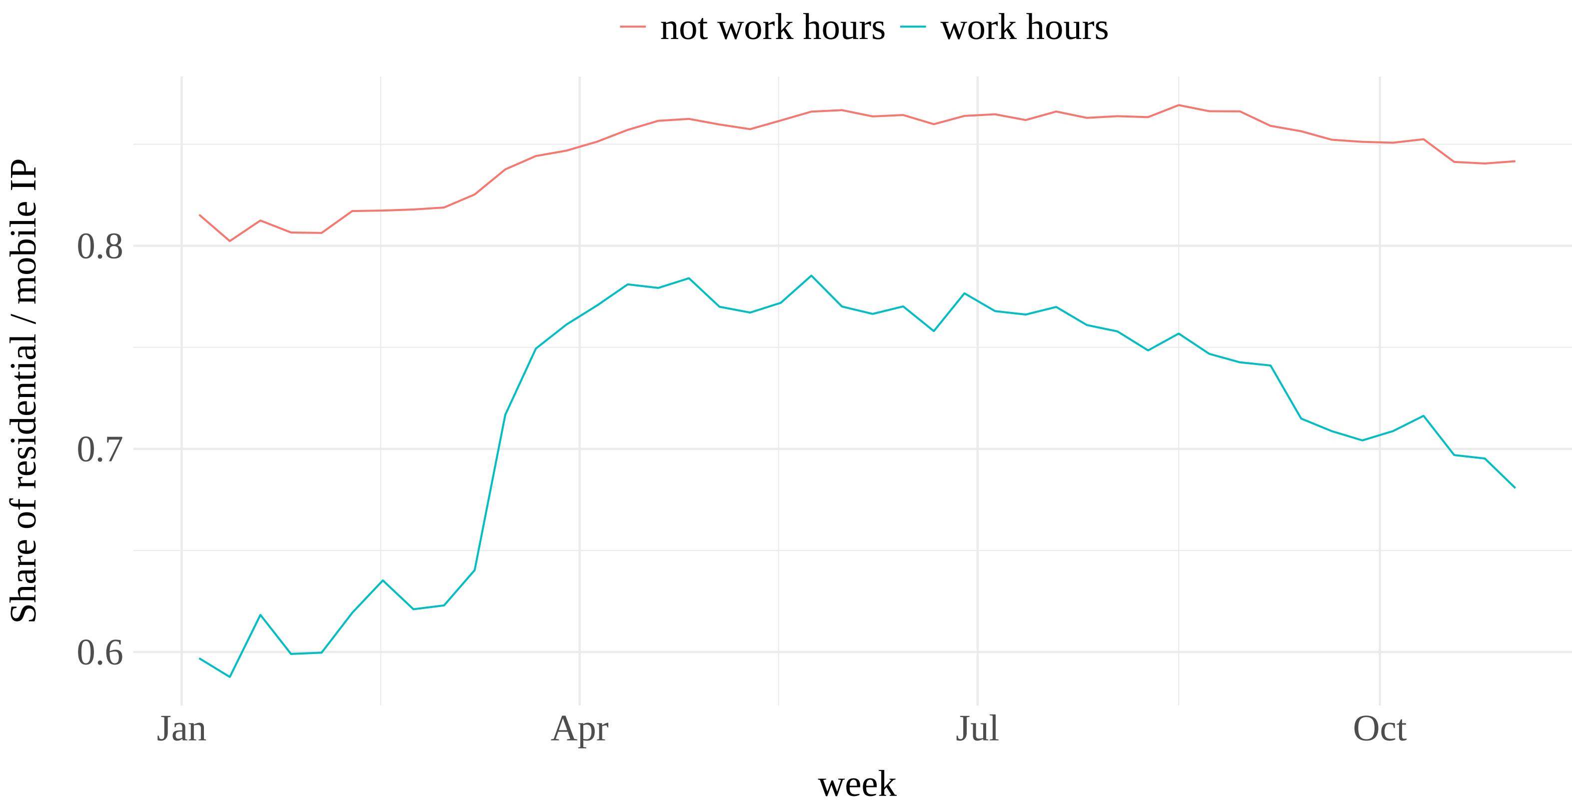

Figure 2: Reading activity during work hours during pandemic times

In this graph, we plot the number of work-time (Monday-Friday 9 a.m. to 6 p.m. local time) internet reads by type of IP address during the year 2020. Our classification for IP addresses is described in section 3.1, involving a mixture of hand-classification and a machine learning algorithm. Internet traffic in this graph is global.

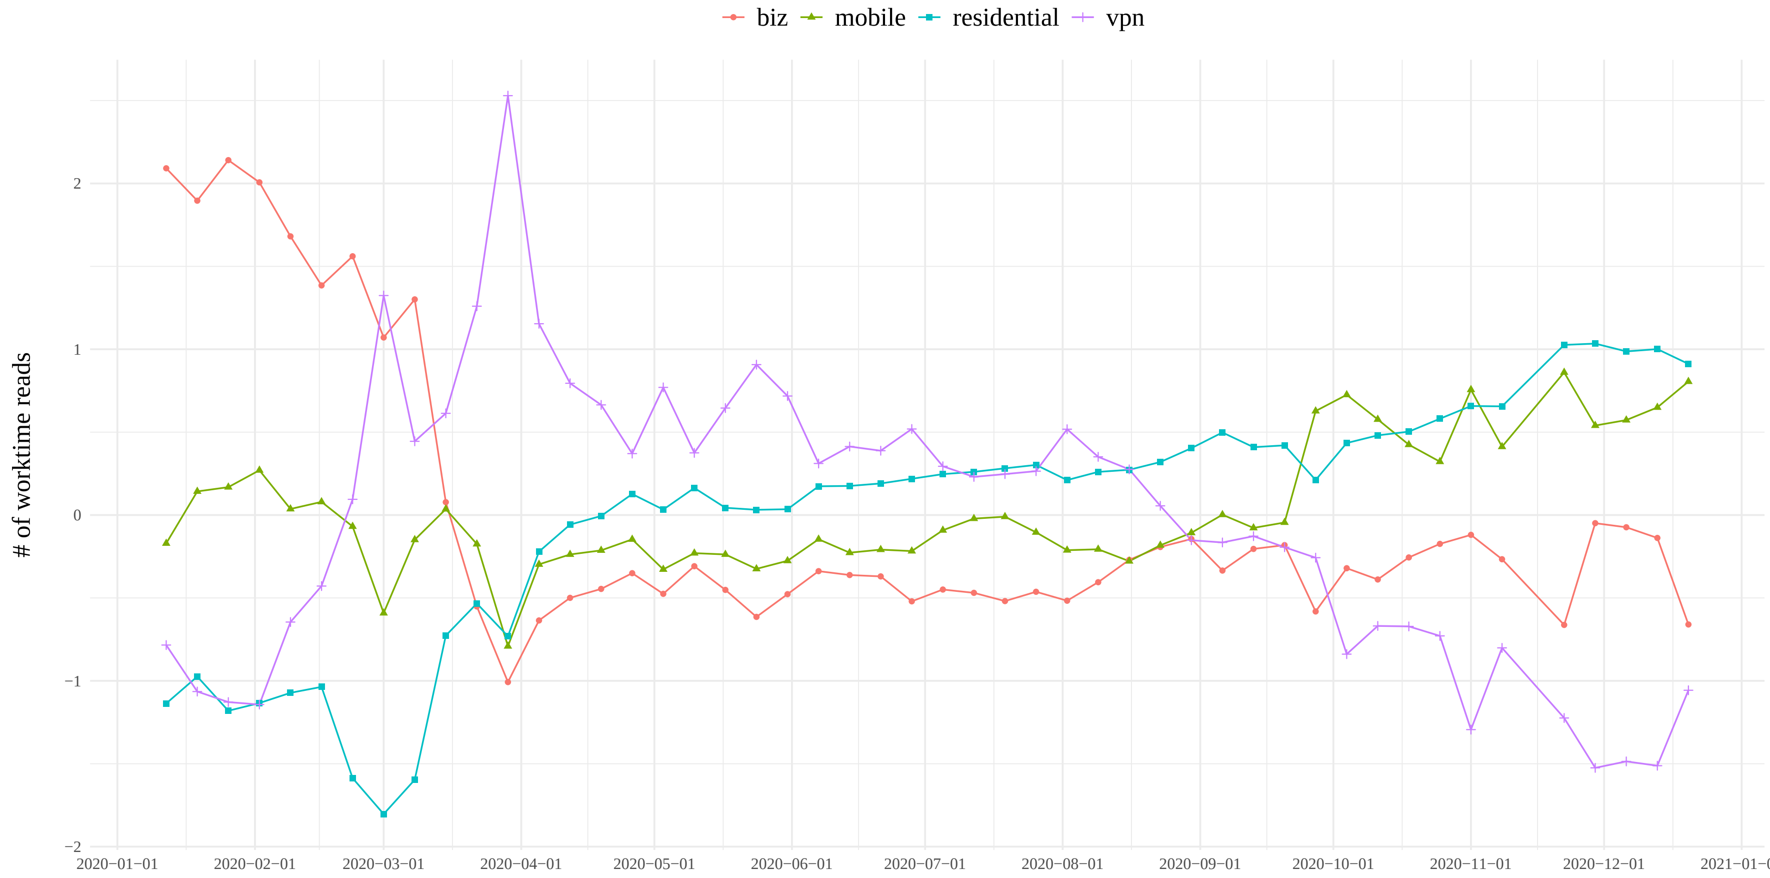

Figure 3: Δ remote work by state

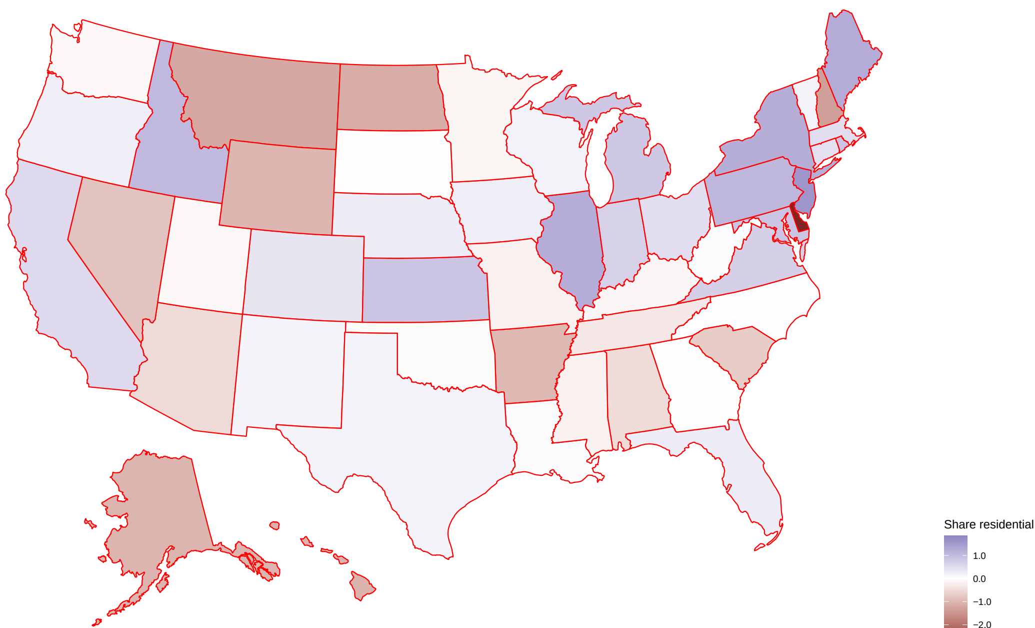

Table 1
Sample summary statistics

In this table, we present summary statistics on the composition of the internet activity dataset. Panel A provides data on the quantity of consumption. *Event Count* is the number of content consumption events observed by the data provider on a daily basis. Panel B provides analysis of top-level domains by IAB classification for the top 15 observed IAB classifications. Panel C reports a breakdown of country of visitor origin for the first eight months of 2020. Panel D reports a breakdown of Alexarank ranking by fraction of content consumption. Panel E reports the distribution of topics into themes by the data provider's content classification taxonomy.

Panel A: Size of the dataset

Year	Filtered Event Count (MM)	# Email Domains
2019	1009	6,828,914
2020	895	6,589,901

Panel B: Breakdown by IAB of publisher (top 15)

IAB	% Total	IAB name
IAB12	21.77	News
IAB3	8.81	Business
IAB13	8.41	Personal Finance
IAB19	7.79	Technology, Computing
IAB9	6.37	Hobbies, Interests
IAB25	5.48	Non-Standard Content
IAB1	5.27	Arts, Entertainment
IAB5	4.22	Education
IAB22	4.07	Shopping
IAB17	3.68	Sports
IAB15	3.39	Science
IAB7	3.08	Health/Fitness
IAB8	2.45	Food / Drink
IAB11	2.37	Law, Gov't, Politics
IAB20	1.88	Travel

Panel C: Monthly 2020 Breakdown for Top 10 countries

country	IPs	Events (mill)
United States	153301073	6824.8
India	14072322	1159.4
Canada	18422743	806.9
United Kingdom	33419330	775
Brazil	39320056	728
Italy	18604597	696.2
Indonesia	4201970	622.9
Mexico	12811176	492.3
France	30277970	471.1
Japan	49520154	469.7

Panel D: Breakdown by Alexarank of publisher

Alexarank	% Total Reading
$[1-1000]$	27.87
$[1000-10000]$	17.59
$[25000-50000]$	7.32
$[50000-500000]$	5.56
$[500000-10^6]$	15.91
$>10^6$	25.75

Panel E: Topic Distribution by Category

Theme name	**# topics**	**Theme name**	**# topics**
Events and Conferences	60	Human Resources	321
Government	87	Healthcare	327
Biotechnology	106	Energy/Construction/Manufacturing	339
Consumer Technology	149	Marketing	523
Business	301	Finance	561
Legal	305	Company	1025
		Non-consumer technology	1849

Table 2
Classified IPs summary statistics

In this table, we present summary statistics which describe the number of events by our classification of IP addresses. We describe the methodology for creating this classification in Appendix Section A. VPN refers to either business or residential VPN, SMB refers to small business, Business refers to corporations, and mobile refers to mobile network. Panel B segments by time of day. Daytime is 9 a.m. to 5 p.m. and work-days are Monday through Friday.

Panel A: By type of network

Label	Events	# IPs (mill.)
vpn	194.887	2.757
residential	13,189.210	474.996
biz	2,398.067	23.656
mobile	3,348.885	6.714
smb	1,182.044	9.175

Panel B: By time of day

Block of time	biz	mobile	residential	smb	vpn
Weekday Day	0.48	2.97	5.77	2.47	0.01
Week Other	0.44	2.97	6.25	2.43	0.01
Work hours	2.29	8.34	16.18	21.21	0.02
Morning of Workday	0.29	2	3.75	2.30	0.01
Night of Workday	0.77	4.63	9.99	4.42	0.01

Table 3
County-level Mobile Geo-location and % Remote and VPN Internet Traffic

In this table, we document the correlation between the fraction of internet activity at remote IPs *during work hours* and Social distancing metrics come from Safegraph. Work hours are 9-5 p.m. local time on weekdays. The outcome variable is the percentage of county-level traffic originating from a remote network (residential or mobile). *Social Distancing* is defined as the percentage of devices completely at home. *Worklike Behavior* is the percentage of devices in a county leaving their common evening location to go to a common daytime location, which is an approximation of a workplace. Cluster robust standard errors are presented at the county level.

		Remote		VPN	Remote	VPN
	(1)	(2)	(3)	(4)	(5)	(6)
Social Distancing	0.207***		0.028	0.0003***	−0.032**	0.0004***
	(0.022)		(0.023)	(0.0001)	(0.014)	(0.0001)
Worklike Behavior		−0.756***	−0.743***	0.0002	−0.214***	0.0003*
		(0.037)	(0.040)	(0.0001)	(0.023)	(0.0002)
Work Hours?	✓	✓	✓	✓		
County FE?	✓	✓	✓	✓	✓	✓
Week FE?	✓	✓	✓	✓	✓	✓
Observations	272,378	272,378	272,378	272,378	275,184	275,184
R^2	0.700	0.705	0.705	0.248	0.675	0.254

Table 4
Industry and Remote Work

NAICS	NAME
11	Agriculture, Forestry, Fishing and Hunting
22	Utilities
92	Public Administration (91 in the United States, 92 in Canada[4])
44	Retail Trade
23	Construction
48	Transportation and Warehousing
33	Manufacturing
45	Retail Trade
62	Health Care and Social Assistance
72	Accommodation and Food Services
32	Manufacturing
31	Manufacturing
71	Arts, Entertainment, and Recreation
49	Transportation and Warehousing
42	Wholesale Trade (41 in Canada,[3] 42 in the United States[2])
21	Mining, Quarrying, and Oil and Gas Extraction
51	Information[notes 1]
56	Administrative and Support and Waste Management and Remediation Services
81	Other Services (except Public Administration)
53	Real Estate and Rental and Leasing
52	Finance and Insurance
61	Educational Services
55	Management of Companies and Enterprises
54	Professional, Scientific, and Technical Services

Table 5
Baseline Firm-Level Determinants of Remote Work

In this table, we present determinants of the shift of remote work between 2019 and 2020 (or 2021). The unit of observation is a firm. $COVID$ is the per capita number of cases at establishments associated with the firm, weighted by the estimated employment of each establishment (defined by Aberdeen's Computer Intelligence Database). *BusinessClosure* corresponds to the frequency of business closures in 2020 in counties where a firm has establishments, defined by the Spiegel and Tookes (2021) database. *PopulationDensity* and $UrbanPercentage$ are the people per square mile of urban land area and Urban percentage is the percentage of the landmass in a county that is classified as Urban by the United States Census. *Size* is the number of employees in Form 5500. $CommuteDistance$ is the average distance computed for the firm's employees using the reading data. $ihs(Remote_{Prepandemic})$ is the level of pre-pandemic reading remote-work activity. Robust standard errors clustered by firm are reported in parentheses.

		$\Delta Remote$2019$\rightarrow$ 2020			$\Delta Remote$2019$\rightarrow$ 2021
	(1)	(2)	(3)	(4)	(5)
COVID	0.0916***	0.0011	-0.0162	-0.0301	
	(0.0281)	(0.0787)	(0.0306)	(0.0264)	
Business Closure		1.016***			
		(0.3890)			
Business Closure $\times$ COVID		0.5812			
		(0.3732)			
Population Density			0.8034***	0.6627***	0.9273***
			(0.0397)	(0.0330)	(0.0321)
Urban Pct			0.3805***	0.3398***	0.2574***
			(0.0390)	(0.0336)	(0.0424)
ihs(No Employees in 2019)			0.2752***	0.1454***	0.0038
			(0.0179)	(0.0151)	(0.0187)
Commute Distance				0.8989***	0.8646***
				(0.0302)	(0.0363)
ihs(Remote$_{Prepandemic}$)				-50.42***	-59.23***
				(0.1706)	(0.2023)
R^2	0.01233	0.01238	0.01571	0.29817	0.27135
Observations	318,706	318,706	294,911	292,110	300,764
Industry fixed effects	✓	✓	✓	✓	✓

Table 6
Monitoring Costs and a Shift Toward Remote Work

Panel A: Technology as a proxy for Monitoring Costs

In this table, we present determinants of the shift of remote work between 2019 and 2021. The unit of observation is a firm. The *IT index* is the count of number of unique technologies recorded in Aberdeen in 2019. $ITindex\Delta 2019 \rightarrow 2020$ is the change in technologies from 2019 to 2020. $\frac{Laptops}{Employees}$ and $\frac{Desktops}{Employees}$ are the number of laptops and desktops per employee. Software budget per employee takes estimates of software and employees from Aberdeen. Robust standard errors clustered by firm are reported in parentheses.

	$\Delta Remote 2019 \rightarrow 2021$		
	(1)	(2)	(3)
Population Density	1.023***	1.002***	1.014***
	(0.0384)	(0.0394)	(0.0384)
Urban Pct	0.3341***	0.2708***	0.3266***
	(0.0488)	(0.0498)	(0.0488)
ihs(No Employees in 2019)	0.1618***	-0.0692***	0.1354***
	(0.0224)	(0.0251)	(0.0219)
$\frac{Laptops}{PC}$	1.903**		
	(0.8690)		
IT Index		0.8322***	
		(0.0366)	
IT Index $\Delta 2019 \rightarrow 2020$		0.3538***	
		(0.0322)	
Log Software Budget Per Emplyoee			0.2120***
			(0.0329)
R^2	0.01871	0.02073	0.01882
Observations	302,093	290,741	302,086
Industry fixed effects	✓	✓	✓

Panel B: Managerial Layers as a Proxy for Monitoring Costs

In this table, we present determinants of the shift of remote work between 2019 and 2021. The unit of observation is a firm. Managerial ratios are based on pre-pandemic hiring (Columns 1–2) based on job postings from Burning Glass, or based on workers associated with a firm in 2019. Robust standard errors clustered by firm are reported in parentheses.

	$\Delta Remote 2019 \rightarrow 2021$			
	(1)	(2)	(3)	(4)
Ratio of Manager$_{JobPostings}(Prepandemic)$	1.818***	1.453***		
	(0.2556)	(0.2305)		
Ratio of Managers$_{Linkedin}(Prepandemic)$			1.671***	2.667***
			(0.1535)	(0.1491)
Controls		Yes		Yes
R^2	0.01365	0.27356	0.01369	0.27070
Observations	135,635	123,085	196,942	153,866
Industry fixed effects	✓	✓	✓	✓
State fixed effects	✓	✓	✓	✓

Table 7
Financial Performance: Firm Income

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported in line 1a of Form 1120. Total Income, from line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $\frac{GrossSales2021}{Assets2019} - \frac{GrossSales2019}{Assets2019}$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present reduced form results, while even-numbered columns present IV results. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses. Controls include pre-pandemic level of remote work and log assets in 2018. Kleibergen-Paap robust F-statistics are reported.

	Total Income or Loss				Gross Sales			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Remote Work	-0.047	-0.030			-0.164**	-0.147**		
	(0.033)	(0.044)			(0.057)	(0.058)		
$\widehat{\text{Remote Work}}$			1.206**	1.208**			1.588**	1.695**
			(0.604)	(0.603)			(0.795)	(0.791)
F-stat			339.352	343.756			339.352	343.756
N	80,262	80,262	78,980	78,554	80,262	80,262	78,980	78,554
R^2	0.007213	0.017213	0.000959	-0.001816	0.008938	0.017831	-0.001263	-0.004987
Controls	✓	✓	✓	✓	✓	✓	✓	✓
State FE	✓		✓		✓		✓	
HQ County FE		✓		✓		✓		✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

Table 8
Microevidence: Content Consumption Patterns

In this table, we present cross-sectional regressions of firm-level reading patterns. The instrumental variable is Commute Distance is the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday day-time location and night-time location. The outcome variable is the fraction of reading on workdays that is devoted to non-work related content, defined by the characteristics of the website. Columns (1) and (2) define work-related as reading of websites with less than 1% of work-relevant content (corresponding to 49.5% of content consumption events) and websites with less than 5% (corresponding to 60%). Standard errors clustered by firm are reported in parentheses. Kleibergen-Paap F-statistics are reported.

	Definition 1, Shirking		Definition 2, Shirking	
	(1)	(2)	(3)	(4)
$\widehat{RemoteWork}$	0.0986***	-0.0801***	0.1083***	-0.0211
	(0.0018)	(0.0243)	(0.0019)	(0.0262)
Controls?	Yes	Yes	Yes	Yes
Yes				
F-stat?		1351.52		1351.52
R^2	0.01660	-0.02084	0.01834	0.00181
Observations	327,637	324,239	327,637	324,239
Industry fixed effects	✓	✓	✓	✓
State fixed effects	✓	✓	✓	✓

Table 9
Robustness: Tradeable/Nontradeable Subsamples

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. We split our sample into two groups, Tradeable and Nontradeable, based on the NAICS 4-digit industry classifications. We run our analysis on each subsample separately and report the results below. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported in line 1a of Form 1120. Total Income, from line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $\frac{GrossSales2021}{Assets2019} - \frac{GrossSales2019}{Assets2019}$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present IV results for the sample of Tradeable firms, while even-numbered columns present IV results for the sample of Nontradeable firms. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses. Controls include pre-pandemic level of remote work and log assets in 2018. Kleibergen-Paap robust F-statistics are reported.

	Sales (1)	TI (2)	Sales (3)	TI (4)	Sales (5)	TI (6)	Sales (7)	TI (8)
$\widehat{\text{Remote Work}}$	3.415***	2.302**	-0.404	0.331	1.897**	1.330**	-1.184	-0.360
	(1.248)	(0.923)	(1.530)	(1.198)	(0.864)	(0.653)	(1.572)	(1.424)
Sample	Tradable		Non-Tradable		Urban		Rural	
F-stat	101.654	101.654	133.213	133.213	302.633	302.633	43.476	43.476
N	29,524	29,524	28,833	28,833	70,760	70,760	7,751	7,751
R^2	-0.043432	-0.033310	0.003983	0.001948	-0.005368	-0.001850	-0.032734	-0.026212
Controls	✓	✓	✓	✓	✓	✓	✓	✓
HQ County FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

Table 10

Cross-sectional analysis, Software and Linkedin-Implied Managers as Proxies for Monitoring Ability

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported on Line 1a of Form 1120. Total Income, derived from Line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes in logged quantities from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $log(\frac{GrossSales2021}{Assets2019}) - log(\frac{GrossSales2019}{Assets2019})$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present reduced form results, while even-numbered columns present IV results. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses. Controls include the pre-pandemic level of remote work and log assets in 2018 . Kleibergen-Paap robust F-statistics are reported.

	Sales		NI		Sales		NI	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\text{Remote Work}}$	2.616**	1.120	1.762*	1.187	0.277	-0.301	-0.206	0.008
	(1.235)	(1.348)	(0.977)	(1.327)	(1.141)	(0.949)	(0.839)	(1.221)
Sample	High Monitoring Ability				Low Monitoring Ability			
Measure	IT	Manager	IT	Manager	IT	Manager	IT	Manager
F-stat	148.429	126.866	148.429	130.350	156.025	101.534	156.025	101.534
N	34,572	21,167	34,572	21,143	35,734	23,271	35,734	23,271
R^2	-0.033320	-0.003112	-0.035841	-0.012291	-0.008556	-0.000771	-0.007115	-0.000292
Controls	✓	✓	✓	✓	✓	✓	✓	✓
HQ County FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

Table 11
High Incentive versus Low Incentive

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. We split our sample into two groups, Tradeable and Nontradeable, based on the NAICS 4-digit industry classifications. We run our analysis on each subsample separately and report the results below. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported in line 1a of Form 1120. Total Income, from line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $\frac{GrossSales2021}{Assets2019} - \frac{GrossSales2019}{Assets2019}$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present IV results for the sample of Tradeable firms, while even-numbered columns present IV results for the sample of Nontradeable firms. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses. Controls include pre-pandemic level of remote work and log assets in 2018. Kleibergen-Paap robust F-statistics are reported.

	Sales (1)	TI (2)	Sales (3)	TI (4)	Sales (5)	TI (6)	Sales (7)	TI (8)
$\widehat{\text{Remote Work}}$	3.295**	2.712***	0.168	-0.295	3.165**	2.663***	1.350	0.327
	(1.306)	(0.999)	(1.144)	(0.872)	(1.393)	(1.010)	(1.202)	(0.918)
Sample	High Incentive		Low Incentive		High Incentive		Low Incentive	
Payout Metric	$\frac{Wage}{Income}$				$\frac{EmployeePay}{Employee+OfficerPay}$			
F-stat	130.769	130.769	164.727	164.727	90.932	90.932	145.868	145.868
N	35,224	35,224	36,125	36,125	30,573	30,573	32,596	32,596
R^2	-0.030189	-0.029422	0.002451	0.003364	-0.030513	-0.035002	-0.007305	0.000003
Controls	✓	✓	✓	✓	✓	✓	✓	✓
HQ County FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

Table 12
Size-Subsamples and Profitability

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. We split our sample into two groups, Large and Small, based on the median of 2018 Total Assets. We run our analysis on each subsample separately and report the results below. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported in line 1a of Form 1120. Total Income, from line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $\frac{GrossSales2021}{Assets2019} - \frac{GrossSales2019}{Assets2019}$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present IV results for the sample of Large firms, while even-numbered columns present IV results for the sample of Small firms. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses.Controls include pre-pandemic level of remote work and log assets in 2018. Kleibergen-Paap robust F-statistics are reported.

	Sales		TI		Sales		TI	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\text{Remote Work}}$	3.235**	2.607*	2.593**	2.104*	1.393**	1.321**	0.453	0.460
	(1.412)	(1.429)	(1.103)	(1.112)	(0.614)	(0.630)	(0.365)	(0.376)
Sample	Small				Large			
F-stat	186.183	180.138	186.183	180.138	152.700	148.877	152.700	148.877
N	37,180	36,510	37,180	36,510	42,102	41,456	42,102	41,456
R^2	-0.025310	-0.012209	-0.025046	-0.010087	-0.028064	-0.026379	-0.013735	-0.013942
Controls		✓		✓		✓		✓
HQ County FE	✓	✓	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

A Additional Descriptions of IP Classification

A.1 Data on IP addresses

In order to classify IP addresses into various categories, we collate data from a variety of sources to either label data or obtain features to inform a classification model.

One can obtain various data points on the entire IPv4 space at any point in time. Currently, the internet operates on two protocols: IPv4 and IPv6. IPv4 was introduced by the Defense Advanced Research Projects Agency (DARPA) in 1981.The IPv4 protocol has $2^3$2 IP addresses, or just over 4 billion. IPv6 has 2^{32} address bits and are reverse compatible with IPv4 addresses. One minor limitation worth noting is that we only have IPv4 addresses in our database, but such IP addresses are likely dominant in the dataset. According to this Network World article, only 30% of IP addresses are reachable via IPv6 as of 2020. Statistics about IPv6 adoption likely do not count out IPv6 adoption among the IPv4 address pool. At a minimum, to get online, the owner of an IP address must register information with the RIR. All websites have an IP address and a forward DNS record is necessary to resolve a website such as *microsoft.com* to an IP. An IP address must be registered to an ISP. Therefore, one can learn the ISP, and for a large fraction of IPs, the owner of the IP address names itself. This is commonly referred to as a reverse RDNS. For example, the IP address 147.8.240.42 identifies itself as *vpn2fa.hku.hk*.

One may also ping an IP address directly on various ports. A computer may have an open remote desktop protocol (port 3389, typically) or mail service (sMTP on port 25, for example). The absence of a response may mean the computer is offline, is not listening on that port, or is not listening to the IP address making contact on that port (perhaps only listening to trusted devices). The responses may indicate the presence of services that are unlikely to be observed on a residential network but more likely an office or corporate network, such as a corporate mail server. We received researcher access to data from Rapid7 and Censys, two leading cybersecurity research firms who conduct regular censuses of the IPv4 space.[18]

[18]From Rapid7, we retrieve the reverse DNS table every week from 2019 through 2020. We also obtain data on

We also collected data on IP addresses that were used as part of residential proxy networks. Residential proxies are typically residential IP addresses wherein exchange for a free service – most commonly, the users in these networks obtain a free app or VPN – the companies operating residential networks are allowed to pass a small amount of traffic through them. These networks are used by consumers for VPN and by corporations to study test websites from the end users' vantage point or to scrape data. In the absence of evidence to the contrary, one can reasonably infer that a a residential proxy IP address is almost always residential, and further one can reasonably infer that /24 blocks (e.g., all IP addresses with the same last octet of an IP address) are also very likely residential.[19] we obtained data from two sources. Luminati.io was kind enough to allow us to record IP addresses on its network. Supplemental to Luminati.io, we also obtained data from Mi et al. (2019), which provides a snapshot of residential proxy networks the study procured through the middle of 2018. To the extent that ISPs continue to hold onto the same residential blocks two years later (even if the IPs are reassigned among household and other customers), this data should still serve valid.

A.2 Hand classification

As the machine learning model cannot cover every possible IP address (particularly those whose features were not previously observed), and to reproduce descriptive analysis using independently derived classifications, we hand-classify IP addresses. Using information as of February 2020, we hand-classify IP addresses based on several criteria. First, we define a remote IP as a VPN if (1) its reverse DNS entry mentions VPN in the sub-domain but not the domain. That is, vpn.hku.hk would be a VPN, while hku.vpnservices.hk would not be. Second, we define an IP as an office location as an IP which the data provider has classified as a business if it is consistently rated as a business per the company's algorithm. The company uses its proprietary dataset to classify IP addresses based on the temporal patterns of content consumed (e.g., is the content

ports open from Censys on various common services, such as databases (MySQL, PostgreSQL), mail servers, remote access protocols and so forth.

[19] As argued by Mi, Feng, Liao, Liu, Wang, Qian, Li, Alrwais, Sun and Liu (2019), typically ISPs register /24 blocks as a bundle

consumed 9-5 on a workday?), the device types observed (is it a desktop or a mobile phone or tablet?), and scale features (do we observe hundreds, thousands or millions of reads from the IP?). The IP classification model which it uses commercially to identify IP addresses of small businesses and large corporations is based on a truth set of Forune 1000 companies' offices. Based on their historical data, we classify a business as one where as of Feb 2020, the company has consistently believed the IP address belongs to a business over 90% times over of the past 14 months. One potential bias in this approach is that the original algorithm could classify as a business a person working from home prior to the pandemic. From our perspective, this would induce attenuation bias as such a business would have never "transitioned" to a remote IP.

Fourth, we define a network as a mobile network if it satisfies two of the following criteria: (1) over the course of 2019 it had at least 50% mobile users, or if during the first two months of 2020 it had 90% mobile users, (2) MaxMind's connection type database classifies it as a mobile connection, (3) the reverse DNS of mobile network resolves to a known mobile carrier. Of course, many mobile carriers over landlines as well.

Fifth, if an IP address is not classified by any of the above criteria, we classify something as residential if it belongs in the same /24 block (e.g., the 256 addresses denoted XXX possible within a 3 octet prefix a.b.c.XXX) as a residential proxy of some kind. Residential proxy services are global in nature and are typically entered into residential proxy networks because their users desire a VPN or other free service. In return for this free service, their computers are used to pass small amounts of traffic requests. This access is then sold by providers such as Luminati.io to corporations for purposes such as ad verification or data scraping. We obtained residential proxies from Luminati.io as well as the Github page of Mi et al. (2019), which lists data they scraped from various proxy networks as of 2018. The study by Mi et al. (2019) examine residential proxies and based on a labeled dataset of residential IP addresses and cloud IP addresses, they build a classification model, and show that based on their model, over 96% of IP addresses they observe on residential proxy networks are best classified as residential IP addresses.

A.3 Machine learning classifier

We aim to model the following labels:

- Business IP: Our ground truth labels come from office IPs maintained by the business provider. However, this may capture an IP address which behaves like business IP prior to the pandemic (such as a home office used before the pandemic).
- Business VPN: Different than a consumer VPN, this VPN is used by businesses to access work-relevant content.
- Residence: A home IP address.
- Residential VPN: A home IP address that belongs to a residential proxy network. We presume a user on this type of network is not working in an office. Distinguishing between this and a normal residential IP serves only to aid modeling, so as to reduce confusion between VPNs for business, VPNs for residential use, residential IPs more broadly, and mobile IPs.
- Mobile network: A mobile network used by many different phones.

We label data from the following sources. For business IPs, we use a list of office IPs maintained by the data provider. Business VPNs are those who are classified consistently in 2019 as VPNs by the company's business IP prediction algorithm (90% of 26 weeks) and the reverse DNS entry of the VPN indicates it is a VPN in the sub-domain (e.g., vpn2fa.hku.hk). The rationale is that this is clearly both a VPN and the IP has been flagged consistently as behaving like a business for a prolonged period of time. The ground truth for residential IPs is based off of Charter and Time Warner IP address blocks maintained by the company. Mobile IPs are those which across 2019 had over 200 mobile users per day. Residential VPNs were those who showed up in the dataset provided by Mi et al. (2019).

With these labels, we build features to make predictions. We do multi-class prediction. The features that are entered into the model are as follows:

- *Internet census data*: The number of ports open, the specific common ports open collected by Censys, such as MySQL. The data are taken from the first week of February.
- *Temporal patterns* From the data provider's historical dataset, we can observe the patterns of when content is observed. The data provider calculates the fraction of content observed during work hours versus local time. We take the average of these values across 2019 and the first two months of 2020.

- *Device type features*: Browser requests are made in a format that fits a browser device) that belong to a particular OS, which implies the device type. For example, the user agent

  $$Mozilla/5.0(iPhone;CPUiPhoneOS12_2likeMacOSX)....$$

  is a prefix of a user agent that implies an iPhone user.

- *Scale-based features*: The amount of content consumed at an IP address. At a network shared by many people, presumably the amount of content consumed will be larger. The device type is implied by the user agent, and the data provider tracks the fraction of user agents belonging to mobile devices versus desktops.

- *Profile-based features*: These are derived for our study. At the monthly level – in turn aggregated across our historical period of 14 months prior to March 2020 – we calculate (1) the share of the top profile in an IP. In a residential IP, the share of the top profile should be larger than in an organization or mobile network. (2) the geographic dispersion of profiles (distinct cities, states, and metros associated with profiles). The novel intuition is that mobile and VPN networks are likely to come from disperse locations. (3) The fraction of profiles with an associated firm, which should characterize business networks. (4) The graph degree.

For our labeled dataset, we report summary statistics in Table A13 below. The summary statistics are sensible for each category. The fraction of profiles observed on a business VPN belonging to the top domain, whatever it is, is 94%. For a business VPN, it is 68%. This suggests either the business VPN users are either un-matched or the VPN is perhaps shared across organizations. Conversely, at residential and mobile VPNs the fraction of users associated with the domain drops dramatically. Relative to Businesses, business VPNs are generally more geographically disperse, coming from a larger percentage of cities. Similarly, relative to regular residential addresses, residential VPN users come from a wider range of cities. Mobile networks and VPNs operate at a much larger scale than our residential or business labels, given that pool more users from a broader geography. This is reflected in the final two columns. First, profiles per day reveals approximately the number of users of an IP address (assuming that user stays within a single browser) and second, the total number of events reveals the scale of events observed at particular IP addresses by my data provider. VPNs are 2.8 to 4.6 times larger than their corresponding non-VPN counterparts, and mobile networkers are an order of magnitude larger. Overall, these summary statistics suggest a sensible breakdown of IP addresses characteristics across the labeled data, giving comfort that the labels are realistic.

Using these features, we employ a 70/30 split of my training and test dataset. My MCC score for this prediction is 89.3%. For most classes, our ability to detect them is above 90%, measured by the high sensitivity in the panel. This is our preferred measure because it measures the fraction of those correctly classified. The most common mistake made by the algorithm is that it predicts something is a business when it is in fact a business VPN – the implication of this mistake is that we will underestimate working from home to some degree.

A.4 Addressing Potential Concerns with Classifications

While the classification performance appears reasonable, the labeled data come from the U.S. There is good reason to believe IPs are assigned differently in different regions. IP address assignment is controlled by one of five network operators, one for each continent. IP addresses tend to be assigned differently based on the ratio of IPs to population, the share of IPs used for different purposes, and network-operator specific practices. In the U.S., it is common for a household to have its own IP address. In some countries, IPs may be shared by small towns. Thus, while we may be able to observe internet activity of some firms worldwide, we restrict our study to firms in the United States.

Another issue might be biases contributed by VPNs. From our perspective, someone on a VPN is working remotely. There is little reason to use a VPN when in the office. The question is whether VPNs can be easily identified and whether they matter. Firsst, VPNs do not always tunnel traffic. VPNs create latency as one must bounce the traffic to another server, and bandwidth could be costly (for example, some VPN providers charge in proportion to bandwidth). They may also want to protect the privacy of the user. Thus companies have an incentive to ”leak” traffic that is not strictly necessary. Since publishers we study are public content, they may fall into this category. Suppose further we identify VPN IPs, there are two adduitional scenarios to consider. First, for some VPNs, the VPN traffic has its own IP address and that VPN IP address is just the address that will be observed by the publisher. Second, the VPN is simply a gateway and the actual traffic is routed to another IP. This is called ”network address translation” (NAT). We would only be concerned if the resulting IP address is indistinguishable from the office IP. This might happen in a university for example, because a VPN might be used to access secure resources that are white-listed for a particular IP (such as a library resource). We are unaware of any statistics about the relative frequency of these scenarios.~[20]

On the issue of VPNs, we do not foresee large concerns for drawing inferences. First, VPN usage spikes in the beginning of the pandemic, but subsides quickly after the first few months. Firms adjust to perhaps tunnelling only specific traffic necessary to VPNs or offer other secure access solutions that make it easy to operate without a VPN. Second, half of the traffic we observe in our dataset pertain to mobile devices. We conjecture that a work device such as a laptop is likelier to be VPN-connected but mobile devices are less likely. Third, our estimates suggest that the VPNs we can observe are relatively rare in the data. These are VPNs which have their own IP, and the traffic that comes from them is not translated via NAT into the office IP. NAT VPNs would have to be at least 10–50 times more common in the data than non-NAT VPNs to be economically significant. Finally, if we cannot detect remote work, this serves as attenuation bias which simply works against being able to observe within-firm variation in remote work.

A.5 Estimating Commute Distance

Based entirely on pre-pandemic information, we estimate commute distance as follows. First, profiles are browser sessions that remain persistent. For a profile associated with a firm (either through a consistent recurrence on a company-associated IP, or a third-party profile database) we can observe every IP address the profile accesses an affiliated member website from. Websites are given associations with cities, and the centroid of that city is used to derive a rough latitude/longitude. We estimate commute distance as the distance between the median lat/lon of the weekday during workhours versus the median location (lat/lon) of the weekend based entirely on pre-pandemic information.

While not the main focus of our study, we believe our measure to be economically sensible

[20] HKU's VPN gateways as noted before are 147.8.240.41 and 147.8.240,42. That is, when one wants to know what IP address the gateway name *vpn2fa.hku.hk* resolves to, the internet RIR will point to these two IPs, but the actual IP used on the VPN will be different.

measure of commute distance because it correlates well with SafeGraph data, which is widely used in the literature to study human mobility patterns. SafeGraph's mobility dataset is based on precise coordinates provided by a phone's global positioning system. SafeGraph's Monthly Patterns dataset describes visitations to establishments in the United States, and has a set of data which describes the visitors to an establishment based on their common day-time location and common-night time location. To calculate commute distance, we first calculate the distance between the establishment location and the common night time location of the visitors, which is often used in this industry as an approximation of home location. We then match the establishment, where possible, to parent company. We merge these two datasets to allow us to compare reading-implied commute distance versus Safegraph commute distance, when SafeGraph data is available. While Safegraph's data is inclusive of both retail visitors and employees, and our dataset only consider employees, we would still expect there to be a positive correlation between these measures. As expected, as we show below in Figure A4 we find is that there is a positive relationship between our crude IP based measures. We also confirm this with econometric results.

Our metric has the advantage of (1) being available for many firms in our sample, and (2) being able to distinguish visitors from employees of a firm. However, our problem is that the locations we have are imprecise.

Figure A4: Comparing our measure of commute distance to mobility data

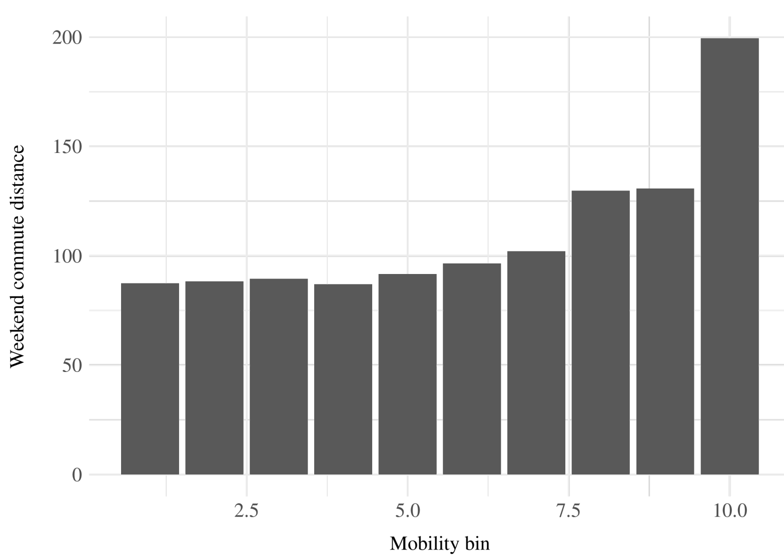

Table A13
Machine learning summary statistics

Panel A: Summary statistics for training sample

label	% Top Domain	City State	State to N	% Mobile	Outdegree	Profiles per day	% Topic total	Events total
biz	0.94	2.07	0.01	0.19	1,169.76	5.83	13.10	1,154.03
biz vpn	0.68	4.61	0.01	0.16	5,437.72	25.63	14.70	5,252.18
residential	0.22	1.51	0.04	0.72	106.91	1.75	10.81	216.43
residential vpn	0.17	2.37	0.06	0.62	424.34	3.62	5.97	586.13
mobile	0.15	98.97	0.19	0.93	13,388.15	145.86	8.48	22,960.17

Panel B: Machine learning accuracy statistics

	Business	Business VPN	Mobile	Residential	Residential VPN
Sensitivity	0.8991	0.731	0.9743	0.9765	0.90244
Specificity	0.9595	0.9782	0.9856	0.9802	0.99545
Balanced Accuracy	0.9293	0.8546	0.98	0.9783	0.94894

B Additional Results

In this section, we present additional summary statistics and robustness tests regarding our main analysis using financial statements.

Table A14
Variance Decomposition

In this table, we run a variance decomposition at the firm-month level. The outcome variable is the firm-month level of ”remote work.” Each row corresponds to incremental combinations of fixed effects.

Fixed Effects	R-squared
Month	0.085
Firm	0.410
Firm + Month	0.493
NAICS-by-Month	0.110
Firm + NAICS-by-Month	0.497
HQ County-by-Month	0.100
Firm + HQ County-by-Month + NAICS-by-Month	0.503

B.1 Investments in IT and Managerial Labor

Building upon our determinants analysis, we investigate the hypothesis that firms ameliorate monitoring costs by investing in IT and managers. Table A15 presents our results. Column 1 presents the first stage regression, while the remaining columns present results of the second-stage regressions of outcome variables on instrumented remote work from the first stage. Columns 2 and 3 present the results using the change in ”IT Index” between 2019 and 2021, the number of unique technologies employed by the firm, as the outcome variable. Column 3 includes additional controls: the log number of employees at the firm and the level of the ”IT index” in 2019. The coefficient estimates are positive and significant at the 1% level in both specifications. An exogenous increase in remote work generates investment in information technologies.

Columns 4 and 5 present results with the inverse hyperbolic sign of the percent of job postings after March 2020 that are related to a managerial role. Column 5 includes the log number of employees and the manager hiring percentage in 2019 as controls. The coefficient estimates on instrumented remote work are positive and significant at the 1% level in both specifications. An exogenous increase in remote work generates an increase in firm hiring efforts for managerial positions. Both sets of results indicate that remote work incentivizes firms to increase monitoring, both through investments in information technology and through increased hiring of managers.

Table A15
Firm Investment in Technology and Managerial Layers

In this table, we present consequences of the shift of remote work between 2019 and 2021. The unit of observation is a firm. Column 1 shows the first stage of the instrumental variable regression for Column 2. Commute Distance is the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday day-time location and night-time location. IT upgrade refers to the change in the $ITindex$ between 2019 to 2021, defined as the count of unique technologies used at the firm. $ihs(\%ManagerJobPostings)$ is the inverse hypoerbolic sine of the percentage of job postings after March 2020 that are related to a managerial role. Industry refers to two-digit NAICS. Standard errors clustered by firm are reported in parentheses. Kleibergen-Paap robust F-statistics are reported.

	$\widehat{RemoteWork}$	IT upgrade		$\Delta PctManagerJobPostings$	
	(1)	(2)	(3)	(4)	(5)
Commute Distance	0.0428***				
	(0.0012)				
$\widehat{RemoteWork}$		0.0072***	0.0776***	0.0194***	0.1659***
		(0.0004)	(0.0065)	(0.0048)	(0.0623)
F-stat?			1285.32		574.24
R^2	0.36188	0.02508	-0.04116	0.00577	-0.00234
Observations	310,193	313,444	310,193	115,630	114,783
Industry fixed effects	✓	✓	✓	✓	✓
State fixed effects	✓	✓	✓	✓	✓

B.2 Additional Profitability Results

Table A16
Commute Distance Groups and Summary Statistics

CD Group	Size	Average pay	Past $\frac{wage}{receipts}$	Past $\frac{wage}{receipts}$	$\frac{Software}{Employee}$	Prepand. WFH
1	0.5062144	74508.71	0.2009115	0.2081269	2305.501	0.4926
2	0.4428209	72933.35	0.3229421	0.2183591	3161.546	0.5475
3	0.3609949	70965.52	0.3270097	0.2168521	3070.584	0.565
4	0.321263	76162.62	0.3270097	0.229341	3457.065	0.572

Table A17
Key Statistics

Mean	SD	
Remote Work	0.1178	0.14377
Total $\frac{Totalincome}{Assets}$	-0.0998745	1.842476
$\frac{Grossreceipts}{Assets}$	-0.2090664	2.479732

Table A18
Firm Income (Logs)

In this table, we present the effects of the shift to remote work between 2019 and 2021, using firms as the unit of observation. Receipts represents Gross Receipts, which are the raw income from goods and services sold, as reported on Line 1a of Form 1120. Total Income, derived from Line 11 of Form 1120, includes all income sources minus the Cost of Goods Sold. The dependent variables are calculated as changes in logged quantities from 2019 to 2021, using 2019 firm assets as a constant deflator (i.e. $log(\frac{GrossSales2021}{Assets2019}) - log(\frac{GrossSales2019}{Assets2019})$). Assets come from Box D on Form 1120. The dependent variables are winsorized at the 5 percent level. The first/last four columns display results using Receipt/Net Income as the dependent variable. Odd-numbered columns present reduced form results, while even-numbered columns present IV results. We instrument remote work using the firm commute distance, the pre-pandemic commute distance estimated using our dataset based on the difference between common weekday daytime and nighttime locations. Industry refers to two-digit NAICS. Standard errors, clustered by firm, are reported in parentheses. Controls include the pre-pandemic level of remote work and log assets in 2018. Kleibergen-Paap robust F-statistics are reported.

	Total Income or Loss				Gross Sales			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Remote Work	-0.054***	-0.055***	0.287	0.272	-0.066***	-0.065***	0.335**	0.342**
	(0.013)	(0.013)	(0.184)	(0.183)	(0.011)	(0.011)	(0.159)	(0.158)
$\widehat{\text{Remote Work}}$			0.287	0.272			0.335**	0.342**
			(0.184)	(0.183)			(0.159)	(0.158)
F-stat			324.263	328.071			310.963	313.578
N	77,946	77,526	-0.005807	-0.008377	75,769	75,380	74,575	74,190
R^2	0.016219	0.016114	76,713	76,297	0.021003	0.019942	-0.014641	-0.018124
Controls	✓	✓	✓	✓	✓	✓	✓	✓
State FE	✓		✓		✓		✓	
HQ County FE		✓		✓		✓		✓
Industry FE	✓	✓	✓	✓	✓	✓	✓	✓

C Additional results

C.1 Additional Descriptive Analysis of the Data

In this section, we provide additional summary statistics about the dataset. First, an assumption of our analysis is that reading activity largely describes activity of the workday, not breaks from it.

To help interpret the data, we examine the distribution of content consumption activity across the workday using a special event-level dataset from the first half of 2018. It is not commercially available nor available for our full sample, but it provides illuminating descriptive statistics. In it, we observe reading data at the (anonymized) article level, including timestamp and topic. The timestamp is localized to the estimated location based on IP address.

One interpretation of internet browsing is that it may be for leisure, rather than work. Although we cannot ascertain the exact motive for a given instance of content consumption, on average, we presume reading for work would be higher during work hours, while reading for leisure would be higher outside of work hours. In Figure A5, we show three panels plotting various 15-minute distributions of reading across the workday. Panel A shows the country-level statistics for workers located in Portugal and Spain. Comparing the Spanish and Portuguese workday interestingly uncovers a pattern consistent with the famous anecdote of the "afternoon siesta," where Spanish workers have a recess after lunch and return to work hours later, while Portugese workers exhibit a dip in reading during the normal lunch hour. While information workers who consume English articles are likely a non-representative sub-sample of their respective populations, it is interesting to see this anecdote manifest itself in the data. By this novel estimate, the Spanish siesta impacts reading activity by roughly 10%.

On a more serious note, Panel B shows the intra-day reading for an anonymous finance firm, and Panel C presents that of an anonymous technology firm. Both are publicly listed in the U.S. We break down reading by topic-theme. The first notable aspect is that reading is industry-relevant, as the top topics in Panel B (the financial firm) relate to companies/finance and the top bar in Panel C (the tech firm) relates to technology. The tech firm shows a notably sharper jump at the start of the workday, with finance/company-related reading increasing around 8 a.m., ahead of technology. One explanation might be that investors at the financial firm consume content on the markets before they open, whereas technologists do not. Panel C reveals that relative to Panel B, the distribution of hours is more widely distributed throughout the day. This would be consistent with tech firms having more flexible hours, while finance firms adhere to financial market trading hours. Broadly, the temporal distribution of reading activity fits the active workday, not pauses from the workday.

Next, in Table A6 we show the allocation of attention across topic themes by industry. We make two important observations. First, firms for whom a topic cluster is clearly relevant lead the category. For example, industry 52- finance – read the most finance articles, spending roughly 19.4% of their time in this category and 15.9% of their time reading about specific companies. Industries 51 and 54 – the information and professional scientific services industries – read the most about technology, registering numbers of 25Does a technology firm really only spend 25% of its time reading about technology? It is hard to tell if this is representative of task allocation at a technology firm or skews in the dataset. However, it is interesting that across industries, the difference between industries tends to be within the same magnitude. For example, within tech, the lowest observed percentage by industry is 17.7%, reflecting a less than 50% difference relative to firms in industry 51. One might conclude, therefore, that in aggregate the time spent across firms in different industries is more

C.2 Full sample relationship between reading and performance

In this section, we examine the impact of reading activity on future spending *unconditionally*. This is to establish baseline correlations supporting the interpretation that specific types of reading which are re-allocated during and arguably because of the pandemic. Appendix Table A19 and Appendix Table **??** present the results. Appendix Table A19 runs two tests.

First, at a monthly level, does reading correlate with future business to business spending? The business spending data comes from Cortera, which provides category-level spending for a large subset of U.S. publicly-listed firms. We only have the ability to delineate business-relevant and other reading from October 2017. *Business Reading* is defined as the count of articles that mont, non-business reading is the count of articles which do not have a business topic associated with them. % Business is the fraction of total reading that is associated with a topic. The outcome variable is the asset-scaled business-service spending two months later. Columns 1 and 2 report analogous results. First, the total *level* of reading in both specifications loads positively. This supports the basic premise of the data provider that a higher amount of content searched for is positively related to business spending. This finding is not necessarily surprising in that consumer search literature has documented that many people use the internet prior to making purchases.[21] However, what is unique about this dataset is the ability to differentiate other types of reading for the same subject. It suggests that non-business search is unrelated to spending.

Similarly, Columns 3–5 explore the impact of business relevant reading on sales. The unit of analysis here is the firm-quarter. Although the stated purpose of the dataset is not to support productivity, one would expect as a basic premise that more business-relevant reading could be related to higher sales turnover for two reasons. First, to the extent that mechanically, a higher level of reading is generated by more employees, sales is proportional to headcount, there may be a mechanical relation between sales through the number of employees or work hours. Second, the information being gathered may be useful for making decisions that drive productivity, even holding work hours or headcount constant. For example, employees may be engaging in the purchase of CRM software to acquire new customers or database software to automate specific repetitive tasks. Either interpretation is benign to our inferences regarding remote work. Columns 3 and 4 show that business-relevant reading is more tightly linked to sales than business irrelevant reading. Column 3 shows that both types of reading are related to sales, although the business relevant reading is more significant by at least a standard error. However, Column 4 which controls for assets eliminates the significance of non-business reading. Whatever interpretation one has, it suggests the robustness of the relation between sales and business-relevant reading in particular. Column 5 also mitigates the size effect by employing the same ratio-specification as Column 3. The inference obtains again that business relevant reading drives business actions. Thus, for the average

[21] For example, see Huang, Lurie and Mitra (2009) or Da, Engelberg and Gao (2011).

firm, one may expect that all-else equal, positive firm outcomes are non-decreasing in business relevant reading.

The above analysis can be interpreted as an extensive margin analysis where the extensive margin is defined as the degree of business relevant reading. One interpretation of the intensive margin is the composition of that reading, and in particular how much is *on a new topic* relative to before. Presumably, when a firm is producing intangible capital, it must read about new ideas. Firms that read about the same ideas are non-adaptive but rather routine.

Table **??** explores reading on the top 200 ideas versus not as defined previously in **??** and its impact on intangible capital formation versus other types of business spending. We define the fraction of reading or total level reading on topics outside the top 200 the firm has read the prior quarter versus reading on any other business relevant topic. Columns 1–4 define the outcome variable as R&D and SG&A at the quarterly level. Columns 5–6 examine CAPX. Column 7–8 examines business service spending at the monthly level as in Table A19. We find that reading on new ideas, either defined as a fraction or in levels, is positively correlated to R&D and SG&A, two outcomes typically described as intangible capital. The magnitudes are larger for R&D, a component of innovation, when accounting for the fact that the sample average firm does around one-third of the R&D than SG&A. The interaction between reading of any kind or reading on new topics is comparatively weaker for CAPX and business spending, and are significant in one-half of the specifications. These correlations are liable to multiple interpretations, but we contend one may find the reduced significance plausible and sensible because capital expenditure and business services spending are less likely than RD to be tied to new products or innovations for the firm, and thus less likely to demand attention toward new ideas from firm employees. In other words, the fact that reading on new topics is more tightly linked to RD and not more recurring, perfunctory expenses may be a signal that new topics supports innovative, creative activity of the firm.

Figure A5: Anecdotes of the workday

Panel A: The afternoon siesta

For a subset of (anonymized) article-level data from January to June 2018, we show the temporal distribution of reading in local time for employees located in Portugal and Spain. The lines represent the probability density activity of reading at 15-minute intervals in percentages (e.g., 1 refers to 1 percent).

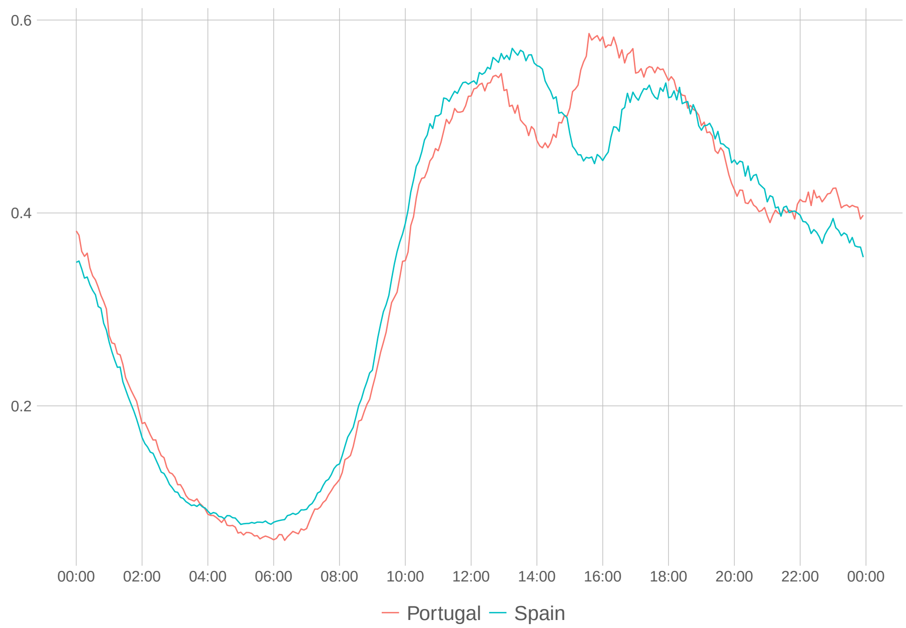

Panel B: Anonymous tech firm

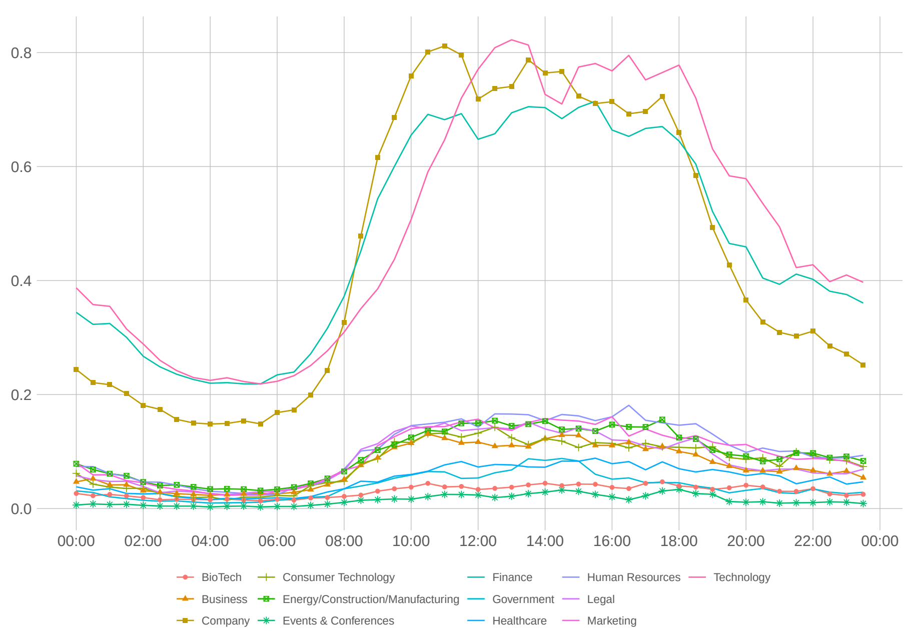

Panel C: Anonymous finance firm

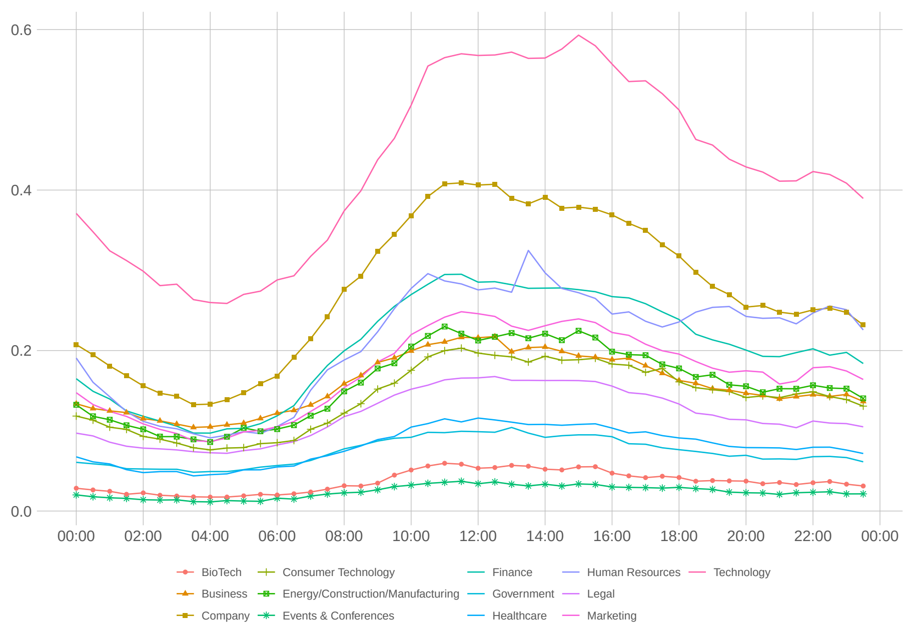

Figure A6: Industry by Topic Theme

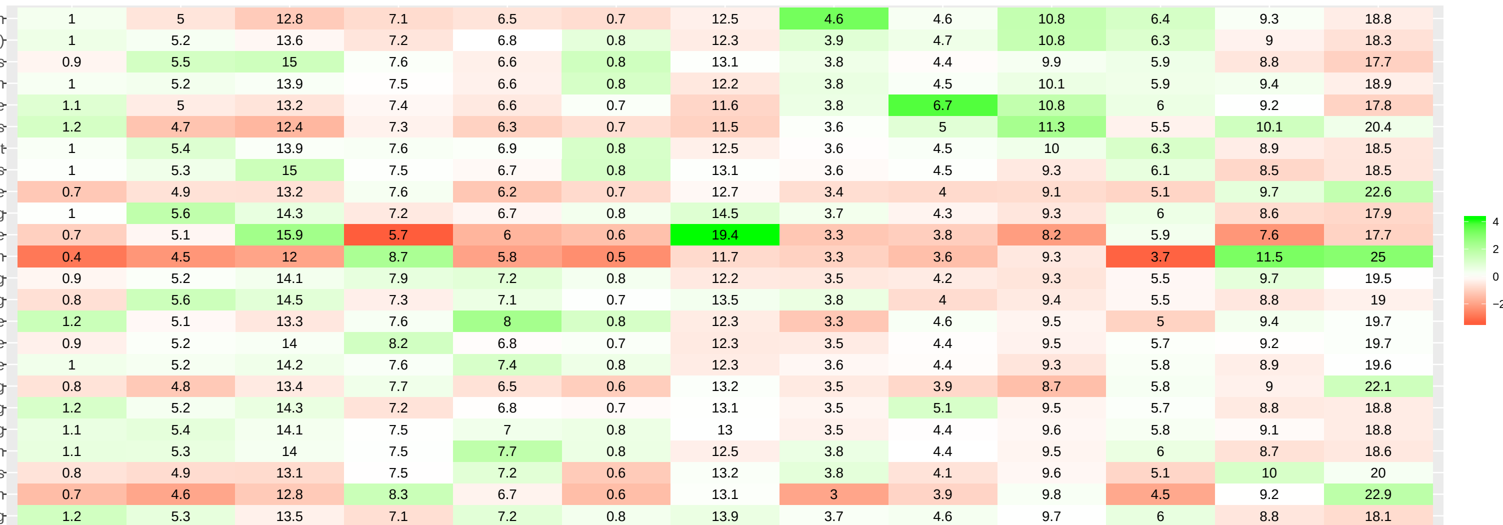

	BioTech	Business	Company	Consumer Technology	Energy/Construct/Manuf	Events & Conferences	Finance	Government	Healthcare	Human Resources	Legal	Marketing	Technology
92 Public Admin	1	5	12.8	7.1	6.5	0.7	12.5	4.6	4.6	10.8	6.4	9.3	18.8
81 Other Services (except Public Admin)	1	5.2	13.6	7.2	6.8	0.8	12.3	3.9	4.7	10.8	6.3	9	18.3
72 Accommodation & Food Services	0.9	5.5	15	7.6	6.6	0.8	13.1	3.8	4.4	9.9	5.9	8.8	17.7
71 Arts, Entertainment, & Recreation	1	5.2	13.9	7.5	6.6	0.8	12.2	3.8	4.5	10.1	5.9	9.4	18.9
62 Health Care & Social Assistance	1.1	5	13.2	7.4	6.6	0.7	11.6	3.8	6.7	10.8	6	9.2	17.8
61 Educational Services	1.2	4.7	12.4	7.3	6.3	0.7	11.5	3.6	5	11.3	5.5	10.1	20.4
56 Admin & Support & Waste Mgmt	1	5.4	13.9	7.6	6.9	0.8	12.5	3.6	4.5	10	6.3	8.9	18.5
55 Mgmt of Companies & Enterprises	1	5.3	15	7.5	6.7	0.8	13.1	3.6	4.5	9.3	6.1	8.5	18.5
54 Prof, Scientific, & Technical Service	0.7	4.9	13.2	7.6	6.2	0.7	12.7	3.4	4	9.1	5.1	9.7	22.6
53 Real Estate & Rental & Leasing	1	5.6	14.3	7.2	6.7	0.8	14.5	3.7	4.3	9.3	6	8.6	17.9
52 Finance & Insurance	0.7	5.1	15.9	5.7	6	0.6	19.4	3.3	3.8	8.2	5.9	7.6	17.7
51 Information	0.4	4.5	12	8.7	5.8	0.5	11.7	3.3	3.6	9.3	3.7	11.5	25
49 Transportation & Warehousing	0.9	5.2	14.1	7.9	7.2	0.8	12.2	3.5	4.2	9.3	5.5	9.7	19.5
48 Transportation & Warehousing	0.8	5.6	14.5	7.3	7.1	0.7	13.5	3.8	4	9.4	5.5	8.8	19
45 Retail Trade	1.2	5.1	13.3	7.6	8	0.8	12.3	3.3	4.6	9.5	5	9.4	19.7
44 Retail Trade	0.9	5.2	14	8.2	6.8	0.7	12.3	3.5	4.4	9.5	5.7	9.2	19.7
42 Wholesale Trade	1	5.2	14.2	7.6	7.4	0.8	12.3	3.6	4.4	9.3	5.8	8.9	19.6
33 Manufacturing	0.8	4.8	13.4	7.7	6.5	0.6	13.2	3.5	3.9	8.7	5.8	9	22.1
32 Manufacturing	1.2	5.2	14.3	7.2	6.8	0.7	13.1	3.5	5.1	9.5	5.7	8.8	18.8
31 Manufacturing	1.1	5.4	14.1	7.5	7	0.8	13	3.5	4.4	9.6	5.8	9.1	18.8
23 Construction	1.1	5.3	14	7.5	7.7	0.8	12.5	3.8	4.4	9.5	6	8.7	18.6
22 Utilities	0.8	4.9	13.1	7.5	7.2	0.6	13.2	3.8	4.1	9.6	5.1	10	20
21 Mining, Quarrying, & Oil & Gas Extraction	0.7	4.6	12.8	8.3	6.7	0.6	13.1	3	3.9	9.8	4.5	9.2	22.9
11 Agriculture, Forestry, Fishing & Hunting	1.2	5.3	13.5	7.1	7.2	0.8	13.9	3.7	4.6	9.7	6	8.8	18.1

Theme

Table A19
Business Relevant Reading and Spending

	B2B Spend		Sales		
	(1)	(2)	(3)	(4)	(5)
Business Reading	0.112***		1.609***	0.768***	
	(0.025)		(0.175)	(0.148)	
Non Business Reading	−0.010		1.416***	−0.039	
	(0.021)		(0.169)	(0.136)	
% Business		−0.070			−1.037
		(0.146)			(1.211)
% Business * Total Reading		0.099**			0.723**
		(0.044)			(0.330)
Total Reading		0.050**			2.524***
		(0.024)			(0.206)
Assets				−10.303***	
				(0.587)	
Firm FE?	Yes	Yes	Yes	Yes	Yes
Time FE?	Yes	Yes	Yes	Yes	Yes
Observations	64,016	64,016	36,433	36,433	36,433
R^2	0.767	0.767	0.894	0.909	0.894
Adjusted R^2	0.757	0.757	0.882	0.898	0.882
Note:			*p<0.1; **p<0.05; ***p<0.01		

C.3 Labor and Remote Work

An often proposed explanation for the recent rise in remote work is that workers demand it, such that firms that offer remote work options benefit from advantages in the labor market (Barrero et al. (2021)). To test this hypothesis, we conduct three analyses. First, following studies such as Azar, Marinescu and Steinbaum (2022), we examine the relationship between labor market concentration (using job postings). We revisit the same specification used in Table 5 examining the shift to remote work from 2019 to 2021. The unit of observation is a firm. Our hypothesis is that labor market concentration leads to less remote work, as this gives workers fewer options to move to another employer.

To test this hypothesis, we use Bureau of Labor Statistics data on the unemployment and vacancies in each county over the year 2020. Tightness is the ratio of unemployed to vacancies. We also use Burning Glass data to construct three alternative measures of labor market concentration. First, at the industry level, $HHI-NAICS$ is the Herfindahl-Hirschman index of job postings at the Employer level, calculated at the NAICS six-digit industry level, based on data from 2007 to 2019 from Burning Glass Technologies. We construct two occupation-level measures. $HHI-SOC4$ is calculated at the four-digit Standard Industrial Occupation level, and weighted for each firm in proportion to the number of job postings by the firm for that occupation from 2007 to 2019. Likewise, $HHI-SOC6$ is calculated at the six-digit Standard Industrial Occupation level.

We report our analysis in Table A20. Column 1 studies labor market tightness, revealing a positive relationship between tightness in the areas surrounding the establishments of a county and remote work. Columns 2–4 show labor market concentration for each one of our incremental measures, and Column 5 uses the SOC-6 measure but adds industry fixed effects. All specifications control for our baseline determinants: size, urban area share, and population density. We find that across all of our measures, the relationship between labor market concentration and remote work is negative. Based on the coefficient estimates, a one-standard-deviation increase in SOC-6 labor market concentration is associated with a 2% decline in remote work from 2019 to 2021.

For descriptive interest, we also examine what types of workers most prefer remote work. We extract LinkedIn data from Revelio Labs to characterize how firm remote-work decisions relate to employee demographics. We regress the change in remote work from 2019 to 2021 for each firm on the firm demographics. In each specification, we control for the percentage of employees that are managers. Table A21 presents the results of specifications run using different demographic variables. In Column 1 we include the number of active employees at the firm in 2019. The coefficient is positive and significant at the 1% level. Column 2 presents the results of the specification run using the percentage of employees at each firm who were employed at the same firm three years before the start of the pandemic. The coefficient estimate is positive and significant at the 1% level, indicating that firms with longer tenured employees are more likely to implement work-from-home policies and maintain them throughout the pandemic. Column 3 presents results from the regression using a measure of average employee skill-level at each firm. The coefficient on employee skill-level is positive and significant at the 1% level. Firms with more skilled employees exhibit a greater increase in remote work. We document a negative relationship between the change in remote work and the percentage of Caucasian employees and the fraction of male employees, as shown in Columns 4 and 5. Lastly, Column 6 presents a specification run using all demographic variables. We find that the signs of the coefficients on each variable remain the same and all coefficients remain significant at the 1% level. While these results are only indicative, the positive relationship between remote work and employee skill-level, age, and tenure with the firm are consistent with the idea that employees with labor market leverage obtain more favorable remote work opportunities.

Second, we present an analysis of labor flows across firms. We report this analysis in Table A22. The main finding in this analysis is that workers flow from low remote-work firms to high ones. In this analysis, the unit of observation is a firm-pair matrix. Using the Revelio data, we define all pairs of firms for which we have data on remote work and also data from Revelio on employees. For each employee, we form a panel of the last job the employee had prior to March 2020, and the last job they had after 2020. Firm pairs are only counted once (e.g., if flows from firm A to B are counted, the reverse is not counted). Given that 150,000 firms would yield 11.2 billion pairs, we only count those pairs where there was at least one transitioning employee.

Our outcome variable is $NetHiring$. If an employee joins Firm B from Firm A, we consider that to be adding to $NetHiring$ whereas if the employee leaves the Firm A for Firm B, that subtracts from $NetHiring$ for Firm A. To control for labor demand on both sides, we scale the labor flows by the total number of inflow to each firm (e.g., for a person going to a firm, scale by the number of employees that firm hires and for a person leaving a firm, scale by the number of leavers of that firm). $\Delta Remote_{dest}$ is the shift from remote work by the firm at the destination firm from 2019 to 2021. $\Delta Remote_{origin}$ is the shift at the origin firm over the same period. Standard errors are clustered by pair.

Thus, the variable of interest, $Remote_{dest} - Remote_{origin}$, captures the difference in the level of remote work over this period. Column 1 shows the relationship between remote work differentials and net hiring without any fixed effects. A standard deviation movement of the variable of interest is 0.23, thus the coefficient implies a 10.7% increase relative to the residual standard error of the outcome variable. Column 2 adds destination fixed effects, while Column 3 adds origin fixed effects. In all specifications, the coefficient on $Remote_{dest} - Remote_{origin}$ is positive and significant at the 1% level. The coefficient implies that a one- standard-deviation differential in remote work between firms is associated with a 9% difference in $NetHiring$.

In our next test, we normalize our remote work measures by the pre-pandemic average of remote work at the firm-level to address concerns that remote work may be difficult to compare across firms due to technical aspects of our measure. Panel B repeats this analysis by looking at simply the difference in shifts (e.g., $\Delta Remote_{dest} - \Delta Remote_{origin}$). We find similarly significant results. These results indicate a significant flow of workers from low remote-work to high remote-work firms since the onset of the pandemic.

Table A20
Labor Market and Determinants of Remote Work

In this table, we present determinants of the shift of remote work between 2019 and 2020 or 2019 to 2021. The unit of observation is a firm. $Tightness$ is a measure of labor market tightness (unemployed to vacancies) from the Bureau of Labor Statistics. $HHI-NAICS$ is the Herfindahl-Hirschman index of job postings at the Employer level calculated at the six digit industry level based on data from 2007 to 2019 from Burning Glass Technologies. $HHI-SOC4$ is calculated at the four-digit Standard Industrial Occupation level, and weighted for each firm in proportion to the number of job postings by the firm for that occupation from 2007 to 2019. Likewise, $HHI-SOC6$ is calculated at the two-digit Standard Industrial Occupation level. Industry refers to six-digit NAICS. Standard errors clustered by firm are reported in parentheses.

	$\Delta Remote 2019 \rightarrow 2021$				
	(1)	(2)	(3)	(4)	(5)
(Intercept)	42.04***	42.08***	42.04***	42.01***	
	(0.1475)	(0.1477)	(0.1474)	(0.1475)	
Tightness	0.3499***				
	(0.0313)				
HHI-NAICS		-0.4311***			
		(0.0306)			
HHI-SOC4			-0.2363***		
			(0.0371)		
HHI-SOC6				-0.4100***	-0.4197***
				(0.0439)	(0.0657)
Controls?	✓	✓	✓	✓	✓
R^2	0.24803	0.24827	0.24785	0.24796	0.25774
Observations	303,500	302,432	303,702	303,702	303,686
Industry fixed effects					✓
State fixed effects					✓

Table A21
Labor and Demographics

In this table, we present determinants of the shift of remote work between 2019 and 2021. Demographics are inferred from employees at the firm as of 2019. Data come from Revelio, whose underlying data on worker resumes primarily comes from LinkedIn. Demographics are estimated using the data provider's model. The unit of observation is a firm.

	$\Delta Remote$2019$\rightarrow$ 2021					
	(1)	(2)	(3)	(4)	(5)	(6)
% Managers	2.81***	2.83***	2.72***	2.90***	2.86***	3.23***
	(0.145)	(0.148)	(0.145)	(0.145)	(0.146)	(0.160)
No. Active Employees, Linkedin 2019	0.384***	0.413***	0.434***	0.355***	0.391***	0.538***
	(0.020)	(0.021)	(0.020)	(0.020)	(0.020)	(0.024)
% Employeed 3 Years Before Pandemic		0.128				0.748***
		(0.208)				(0.246)
Average Employee Skill			0.335***			0.529***
			(0.034)			(0.037)
% Caucasian Employees				-2.81***		-3.64***
				(0.232)		(0.272)
Male Share					-1.85***	-2.50***
					(0.227)	(0.259)
Age Above 50 Share						-0.022
						(0.220)
R^2	0.26877	0.26932	0.26925	0.26951	0.26929	0.27347
Observations	159,351	157,420	159,351	159,344	158,998	145,813
Industry fixed effects	✓	✓	✓	✓	✓	✓
State fixed effects	✓	✓	✓	✓	✓	✓

Table A22
Labor flows

In this table, we present an analysis of labor flows across firms. The unit of observation is a firm pair matrix. Firm pairs are only counted once (e.g., if flows from firm A to B are counted, the reverse is not counted). Using the Revelio data, we define all pairs of firms for which we have data on remote work and also data from Revelio on employees. For each employee, we form a panel of the last job the employee had prior to March 2020, and the last job they had after 2020. If an employee joins the new firm, we consider that to be adding to $NetHiring$ whereas if the employee leaves the firm, that subtracts from $NetHiring$. To control for labor demand on both sides, we scale the labor flows by the total number of inflow to each firm (e.g., for a person going to a firm, scale by the number of employees that firm hires and for a person leaving a firm, scale by the number of leavers of that firm). $\Delta Remote_{dest}$ is the shift from remote work by the firm at the destination firm from 2019 to 2021, $\Delta Remote_{origin}$ is the shift at the origin firm over the same period. Standard errors are clustered by pair.

Panel A: Difference in levels of Remote Work

	Net Hiring		
	(1)	(2)	(3)
(Intercept)	0.0030***		
	(0.0002)		
$Remote_{dest} - Remote_{origin}$	0.1043***	0.0845***	0.0903***
	(0.0012)	(0.0015)	(0.0015)
Observations	946,105	946,105	946,105
R^2	0.01127	0.45943	0.42931
Destination fixed effects		✓	
Origin fixed effects			✓

Panel B: Difference in Δ of Remote Work

	Net Hiring		
	(1)	(2)	(3)
(Intercept)	0.0028***		
	(0.0002)		
$\Delta Remote_{dest} - \Delta Remote_{origin}$	0.0196***	0.0119***	0.0158***
	(0.0008)	(0.0010)	(0.0011)
Observations	946,105	946,105	946,105
R^2	0.00088	0.45601	0.42540
Destination fixed effects		✓	
Origin fixed effects			✓