Digitization and Employment in the Pandemic: Evidence from Seventy Billion Emails

Sida Peng, Peichun Wang, Eric Auerbach, Hongwei Liang, Andy Wu

November 22, 2021

What is the effect of the COVID-19 pandemic on IT adoption and utilization? We construct and analyze a proprietary and comprehensive dataset of the digital technology adoption and utilization of 6706 public and private US firms and 65 million employees as reflected by their use—over seventy billion emails—of the most popular enterprise productivity suite in the world from January 2019 to January 2021. We find that the pandemic disrupted both the levels at which firms employ workers utilizing digital technology and the levels at which digital workers utilize digital technology, with significant heterogeneity across firms. While larger firms and firms with higher levels of pre-pandemic digitization sustain digital employment and utilization through the pandemic and return fully to their pre-pandemic trends of employment growth, smaller firms and those with low pre-pandemic digitalization continue to suffer persistently from the pandemic.

1 Introduction

An extensive literature studies the causes and consequences of information technology (IT) adoption (*1, 2*). At the individual employee level, IT improves productivity through enhanced connectivity (*3, 4*), especially across intra-organizational boundaries (*5*). At the firm level, IT generally improves performance (*6–9*), particularly for firms with complementary resources like supporting business processes (*10–12*). Given these returns to both employee productivity and firm performance, the literature emphasizes the importance of understanding the determinants of IT adoption. Examples include the existence of complementary knowledge and business processes (*13, 14*), firm size and organizational slack (*15, 16*), and the availability of support services (*17*).

However, few empirical studies of IT adoption consider the impact of external market and societal-level factors (*2, 18*). An economic recession, for example, may depress the demand for goods and services, putting pressure on firms to either adapt digitally or pull back to focus on core traditional, non-digital capabilities. For individual workers, a shock like a recession can have a heterogeneous impact across individual workers, like the 2008 financial crisis which disproportionately impaired older male workers in manufacturing (*19*).

The main contribution of this paper is to study the effect of the COVID-19 pandemic on IT adoption. Since the start of the pandemic, researchers have moved rapidly to conduct surveys and leverage administrative data to understand the impact of the pandemic (*20–22*). In the field of information systems, (*23*) leverages data on open-access preprint research studies to show that, while the pandemic lockdown benefited overall research productivity, female academic productivity dropped relative to male academics.

The COVID-19 pandemic presents a unique opportunity to understand the role of external factors on IT adoption (*24*). First, the need to dismiss in-person workplaces forced firms into using technology for remote work (*25, 26*). Without the pandemic, these firms may not have been incentivized to adopt this technology because of inertial factors like firm culture.[1] We find evidence below further suggesting that organizational factors like slack resources and a general cultural disposition toward digital technology influence a firm's ability to effectively adopt digital technology.

Second, the pandemic drove an economic recession whereby resource-constrained firms reduced ongoing expenses, often by cutting down on employment (*27*).[2] Several studies have documented the heterogeneous effects of the COVID-19 recession on firm performance (*28–30*). Our work focuses specifically on firms' digital employment and how firms with different levels of slack resources and digital capabilities retained their digital workforce through the pandemic.

We leverage a unique dataset on the universe of individual-level email activities from a large sample of US firms using proprietary data from Microsoft, the largest provider of enterprise software services in the world. Compared to other data used in prior studies, our data covers a broad set of many major US firms both public and private, as opposed to prior studies concerning a single firm (*3, 31*) or only publicly traded firms (*6*). Moreover, we are able to track individual employees at each firm, allowing us to accurately generate a timely measure of employment during the pandemic compared to self-reported data (*16*) or data collected by a government agency (*22*). Finally, our use of email volume serves as a proxy for employee-level IT utilization (*5,31,32*); prior literature emphasizes the importance of this granularity especially to distinguish employee-level utilization from organization-level adoption (*6, 10*).

To evaluate the impact of the pandemic on firms' digitization and employment, we use an

[1]Business Wire. (March 19, 2020). Gartner HR survey reveals 88% of organizations have encouraged or required employees to work from home due to coronavirus. `https://www.businesswire.com/news/home/20200319005102/en/Gartner-HR-Survey-Reveals-88-of-Organizations-Have-Encouraged-or-Required-Employees-to-Work-From-Home-Due-to-Coronavirus`

[2]U.S. Department of Labor. (2020, April 23). News release: Unemployment insurance weekly claims. `https://oui.doleta.gov/press/2020/042320.pdf`

interrupted time series design. We find that the pandemic in general put significant pressure on firms to adopt digitization tools. As a general pattern, we find that all types of firms in the study experience a drop in their digital workforce during the peak of the pandemic. They also experience a dramatic increase in the use of digital technology during and through the pandemic. But once we dig into these general patterns, clear heterogeneity across firm type becomes apparent.

Across a broad set of firms—varying drastically in both size and pre-pandemic digitization levels—we document the differential impact of the COVID-19 pandemic on their digital workforce. In particular, we find evidence of a K-shaped pattern in digital adoption, along the lines of a Matthew effect, that favors those already better-positioned for the digital transformation independent of the pandemic. Two key patterns emerge that favor pandemic-driven digitization among firms with organizational characteristics already better-equipped for digitization.

First, we find that larger firms sustain more of their workforce through the pandemic and return fully to their pre-pandemic trends of employment growth. At the peak of the impact of the pandemic—around the summer of 2020 when we document the lowest general levels of digital employment—smaller firms suffered much more, experiencing a much larger negative shock to their workforce growth, from which they do not fully recover by the end of the study. We argue that this difference occurs because large firms have the slack and organizational resources to invest rapidly in digital technology and generally sustain their operations during an economic downturn. We document empirical evidence for this argument by showing that large firms are better able to take advantage of this transition to digital technology, with significant persistent gaps in workforce digital technology utilization emerging during and then persisting through the tail of the pandemic.

Second, we find that firms with higher levels of pre-pandemic digitization both sustain more employment through the pandemic and have a better recovery at the end of our study. We argue that firms already ahead on digitization are better-prepared to adapt to the unique circumstances of the pandemic and, further, that their pre-pandemic digital transformation, potentially due to factors like culture, leadership, and tacit organizational routines, facilitates digitization in the pandemic period. Indeed, those firms with high levels of pre-pandemic digital utilization among employees experience even faster growth in digitization during the pandemic, leading to an even larger persistent gap between digital and non-digital firms.

The rest of the paper is organized as follows. Section 2 introduces the data. Section 3 describes our empirical methodology. Section 4 presents the results of our empirical analysis. Section 5 discusses these results in a broader context and concludes.

2 Data

We use Microsoft Office data to measure the change in employment as well as the digital activities for major US enterprises, where the term ”enterprise” refers to the firm or organization contracting with Microsoft. The core of Microsoft Office is the Exchange Server and Active Directory, a set of tools that an enterprise can use to manage employee accounts across the organization. This dataset provides us with several unique advantages to study the impact of the COVID-19 pandemic on major US companies.

First, our data has broad coverage of major US enterprises. Microsoft Office is a major productivity application used by companies all over the world, including over 95% of fortune 500 companies and many smaller companies, by 2019. In contrast, most of the IT productivity literature focus on publicly traded firms (*6*). Unlike most of the studies on pandemic impact which are based on survey or self-reported data, our data provides a richer characterization of technology adoption in the US economy.

To form our sample, we select the set of Microsoft Office US customers with at least 1,000 purchased users in March 2019, and we keep track of their activity from March 2019 to De-

Figure 1: Total number of unique email accounts over time

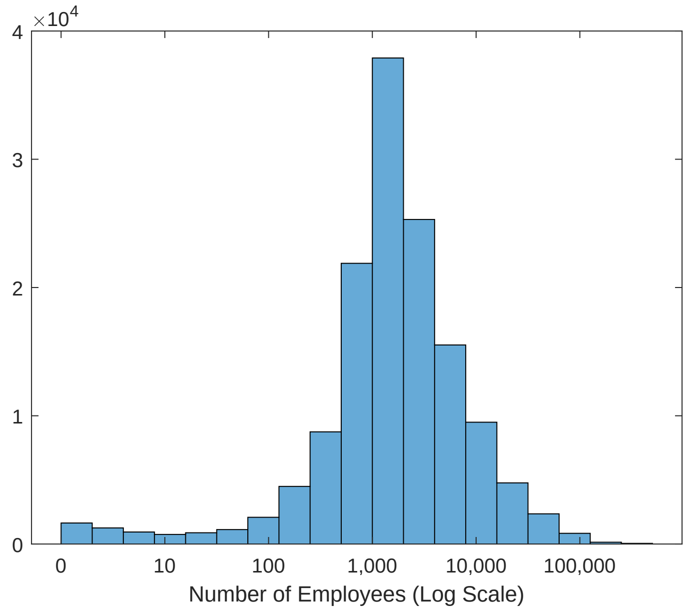

Note: The figure shows the distribution of the log number of active email accounts for each firm-month. It measures the number of digital employees in the firm. Mean: 4523; median: 1507; standard deviation: 1295; min: 0; max: 473613.

cember 2020.[3] We observe 6706 firms across 21 industries, with a total of 65 million unique email accounts during our sample. We observe the monthly level of email accounts and the total number of emails sent both within and across companies, with a total number of 69.3 billion emails covered during our sample. Figure 1 shows the distribution of the log number of active email accounts for each firm-month, which we use to proxy for the number of digital employees. Figure 2 presents the average monthly email volume per employee account at the firm level. Employees at a firm send on average 81 emails per month. The maximum we observe is 778 emails per average user-month.

Second, our data provides an accurate and timely measure of digital employment. For each

[3]Seats refer to the number of user accounts purchased. Our analysis focuses on the number of actively engaged accounts. We choose the threshold because 1,000 seats is a threshold required in Microsoft's Enterprise Agreement to qualify for additional discounts, reducing the possibility of multiple contracts with Microsoft from the same customer company.

Figure 2: Distribution of average monthly email volume per employee

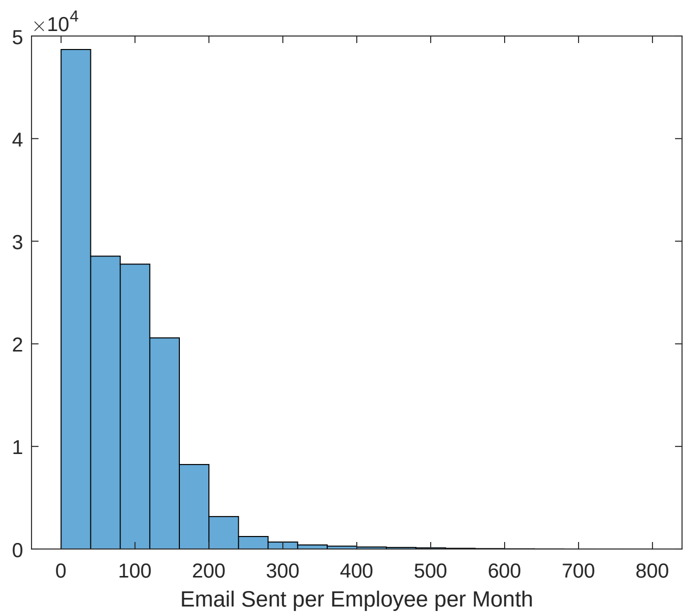

Note: The figure shows the frequency distribution of the firm-month-level average monthly email volume per employee account. Mean: 81; median: 70; standard deviation: 68; min: 0; max: 778.

Microsoft Office enterprise, a (digital) employee is defined as an active email account associated with the enterprise's Enterprise Agreement (EA) license. The Microsoft EA license bundles productivity tools such as Outlook, Word, Excel, and PowerPoint with security and compliance products such as Azure Active Directory and protection against phishing email and malware to provide a one-stop solution for enterprise customers. The Azure Active Directory is an account control tool, which integrates with the enterprise's internal HR platform to facilitate access control. As a result, we take the number of users within an enterprise to be an accurate measure of the number of core employees. Compared to traditional sources of employment data from government census or other surveys, our data is more timely and allows researchers and policymakers to quickly analyze the impact of a recession such as the COVID-19 pandemic.

Third, we use the volume of email communications as a measure of the employees' digital activities. Other measures of employees' use of digital tools have been considered in the literature: (*6*), digital collaboration tool (*5*), and networked communication (*3*). These measures characterize the familiarity and importance of digitization in the firms' workforce but are limited in the sense that they are difficult to compare across industries. We focus in contrast on measures of email utilization. Email is relatively old but broadly used in firms' distinct functions and can be directly compared across industries. Furthermore, Microsoft Exchange is one of the most widely used email services among US enterprises. Other studies that use email volume as a measure of digitization include (*3, 31, 32*).

3 Methodology

We use an interrupted time series design to evaluate the impact of the pandemic on firm employment and technology adoption. The ideal approach would be to simply compare the outcomes of firms that were affected by the pandemic to a control group of similar firms that were not affected. Unfortunately such a control group does not exist because all of the firms in our study were impacted by the pandemic at essentially the same point in time. Our approach uses instead the outcomes of the firms before the pandemic occurred to construct a control. Specifically, we fit a dynamic linear regression model using the pre-pandemic outcomes, and we use the fitted values from this model in the post-pandemic time periods as the control. We use a dynamic linear regression model because it allows for the outcomes to trend upwards or downwards even before the pandemic occurs. This identification strategy is known in the event study literature as an interrupted time series design (see for instance (*33*)). See the appendix for a formal discussion.

4 Results

We first present visual evidence of the estimated impacts (i.e., $\hat{\Delta}_{it}$'s) of the pandemic on different types of firms and then quantify these impacts in a regression model. The two outcomes (i.e., Y_{it}'s) of interest are firm employment, measured using the number of active email accounts, and digital adoption, measured using the average number of emails sent per employee at the firm level. Motivated by the theories of slack resources (*15, 16*) and absorptive capacity (*13, 14*), we then explore heterogeneity of the treatment effect along two dimensions: firm size and pre-pandemic level of digitization. Specifically, we consider a firm to be large if its pre-pandemic employee headcount is above the median of its industry and small otherwise; we define a firm to be already digital if its pre-pandemic average employee email count is above its industry median and non-digital otherwise.

4.1 Visualizations

To focus on the main mechanisms of interest, we compare segments of firms where the effects of the pandemic should be the most salient and observable. In other words, we specifically want to identify the firms for which pre-pandemic size and digitization are most relevant for whether and how the firm adapts to the pandemic. With respect to understanding the differential impact arising from pre-pandemic digitization, we focus on smaller firms that lack the resources to easily adapt in the midst of the pandemic as a large firm may be better-positioned to do. Thus, we compare non-digital small firms to digital small firms. Under a similar logic, to understand the differential impact arising from firm size, we focus on non-digital firms that lack the existing digital capabilities to adapt to pandemic conditions, which distinctly punished less-digital firms. Thus, we compare small non-digital firms to large non-digital firms.

Employment Figure 3 (first panel) compares, among small firms, changes in employment during the pandemic between digital and non-digital firms prior to the pandemic. The pandemic has a severe impact on small firm employment in general and non-digital small firms in particular. The impact of the pandemic on employment is most severe in July 2020, when non-digital small firms suffer over an 8% drop in employment. Small firms of both types then begin rebounding immediately after that trough although, as of December 2020, small firms have yet to return to their expected level of employment based on pre-pandemic trends. The first panel of Figure 3 makes clear that more-digital firms adapt better to the pandemic, losing less employment at the peak and recovering more quickly, thereby demonstrating the importance of digitization on firm performance, especially in a pandemic where only digital interactions are permitted.

In the second panel of Figure 3, we compare, among non-digital firms, the amount of the digital workforce retained by large and small firms. We find that large firms significantly outperform small firms in retaining their workforce, both at the worst times in July 2020 and towards the end of the study period. As of December 2020, large firms appear to fully recover back to their expected level of employment, where they might have been had there been no pandemic. It appears that small firms are disproportionately impacted by the pandemic and that the divide between large and small is exacerbated through the pandemic.

Digitization Figure 4 (left panel) compares, among small firms, changes in digitization between those that are already more digital prior to the pandemic versus those that are not. The pandemic increases the level of digitization for all firms relative to their expected baseline level based on pre-pandemic trends. We interpret this result as the pandemic accelerating digital adoption. While vitally important for small firms, this digital gain from the pandemic differs across firms with different levels of digital preparedness, demonstrating a divergence further

Figure 3: Impact of the COVID-19 pandemic on firm employment

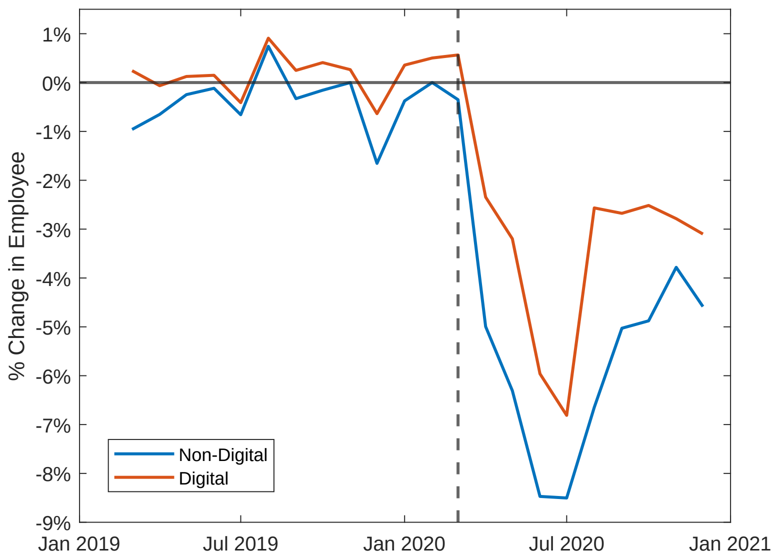

Note: The first panel compares, among small firms, changes in employment relative to pre-pandemic trend between firms with high and low levels of digitization before the pandemic. The second panel compares, among the pre-pandemic less-digital firms, changes in employment relative to pre-pandemic trend between large and small firms.

widening the gap between the prior leaders and laggards in digital adoption before the pandemic. We hypothesize that this could be due to either slack resources created from the better preparedness for the pandemic or firms' different internal culture and leadership driving the different levels of digital adoption.

Finally, the first panel of Figure 4 examines, among non-digital firms, whether the pandemic has differential impact on digitization for small firms vs. large firms. In theory, while large firms might have more slack resources to invest in technological transformations, small firms might be nimbler in adopting technologies. We find that, however, again large firms disproportionately benefit from the pandemic in terms of digital adoption.

4.2 Regression Analysis

While Figures 3 and 4 illustrate the different trends in the raw data, they potentially mask additional heterogeneity in how different industries react to the pandemic. To address this issue, we provide results from a regression analysis. Specifically, we consider the following specification:

$$y_i = \alpha + \beta_1 \cdot Large_i + \beta_2 \cdot Digital_i + \beta_3 \cdot Large_i \cdot Digital_i + \Gamma_i + \epsilon_i \tag{1}$$

where y_i denotes the average estimated impact of the pandemic for firm i across pandemic time periods on either firm employment or digitization, i.e., $y_i = \frac{1}{T}\sum_t \hat{\Delta}_{it}$ where T is the number of months during the pandemic in our sample (see appendix section 1 for details). $Large_i$ and $Digital_i$ are dummy variables indicating whether firm i's pre-pandemic employment or digitization level, respectively is above its industry median. Γ_i is a full set of industry fixed effects.

Table 1 presents the results. We present the combined effects of a firm being "large" or "digital," e.g., a large digital firm shows 17% more employment ($\beta_1 + \beta_2 + \beta_3$) than a small non-digital firm, whereas it is 17.9% (β_1) and 16.3% (β_2) for large non-digital and small digital

Figure 4: Impact of the COVID-19 pandemic on firm employment

Note: The first panel compares, among small firms, changes in digitization relative to pre-pandemic trend between firms with high and low levels of digitization before the pandemic. The second panel compares, among the pre-pandemic less-digital firms, changes in digitization relative to pre-pandemic trend between large and small firms.

firms, respectively. To show that these differences are not a coincidence of the idiosyncratic errors, we also test whether $\beta_1 + \beta_2$, $\beta_2 + \beta_3$, β_1, β_2 are statistically different from zero.

Table 1: Regression results

		Digital	**Non-Digital**
Employment	Large	17.0%	17.9%
	Small	16.3%	00.0%
Digitization	Large	19.9%	19.4%
	Small	16.3%	00.0%

Note: This table shows the combined effects of a firm being "large" or "digital" relative to being a small non-digital firm on employment and digitization during the pandemic. The large digital cell corresponds to $\beta_1 + \beta_2 + \beta_3$ in equation 1. The large non-digital cell corresponds to β_1. The small digital cell corresponds to β_2. All coefficients are statistically significant at the 0.1% level.

For small firms, we find that being more digital pre-pandemic helps retain 16.3% ($p < 0.001$) more of employment and further increases digitization by 16.3% ($p < 0.001$) more through the pandemic. For non-digital firms, we find that being larger helps retain 17.9% ($p < 0.001$) more of employment and catches up in digitization by 19.4% ($p < 0.001$) more through the pandemic. These effects are not significant for large firms or more-digital firms prior to the pandemic. The only exception is that, among digital firms, large firms grow in digitization by 3.6% ($p = 0.002$) more than small firms.

5 Discussion

In this paper, we document how the COVID-19 pandemic impacts firms through their employment and digitization using a novel dataset that contains almost 70 billion emails from Microsoft Outlook between 2019 and 2020. We first use pre-pandemic data to predict trends in the outcomes of interest in the post-pandemic period. We then evaluate the impact of the pandemic by comparing these predicted trends with the realized post-pandemic trends, which allows us to identify systematic differences in the outcomes of interest which we interpret as the causal impact of the pandemic. In this concluding discussion, we highlight the findings and their implications for firm-level digitization and individual employee-level productivity and welfare.

5.1 Firm-Level Implications

We find that, across industries, firms experience a substantial decrease in employment immediately after the March 2020 lockdown. While the reduction in employment slowly recovers toward the end of 2020, the overall employment level is still lower than the expected trend from pre-pandemic level. On the other hand, we found that firms experienced a roughly 20% increase in digitization on average after the lockdown, which is sustained throughout the pandemic period. An important contribution of our work is to formally document the size and scope of the pandemic on the digital economy. Furthermore, we also demonstrate (i) the relative usefulness of digital administrative data in tracking firm employment compared to government survey data, and (ii) the importance of studying societal-level factors that affect IT adoption and firm performance (*2*).

5.1.1 Slack Resources and IT Adoption

The pandemic provides a unique experiment that can be used to identify the importance of slack resources on both firms' IT adoption and performance. We find substantial heterogeneity in these outcomes depending on the pre-pandemic level of firm size and digitization level. While the large decrease in employment virtually disappears for large firms by the end of 2020, persistent gaps in employment of approximately –5% can be found in small firms at the end of our sample. Similarly, we found that large firms also increased digitization more heavily than small firms during the pandemic. One interpretation of these results is that large firms may have the financial and other resources necessary to maintain production and adapt to the unique circumstances of the pandemic, resources that may be unavailable to smaller firms. This interpretation is consistent with the intuition for the importance of slack resources in the literature on IT adoption by firms (*15, 16*).

5.1.2 Absorptive Capacity and IT Adoption

The results of our analysis are also consistent with a literature on the absorptive capacity mechanism of technology adoption (*13, 14*). We find that the increase in digitization was substantially larger for firms with a high level of pre-pandemic digitization. Moreover, firms with a high level of pre-pandemic digitization also retain more of their workforce during the pandemic, demonstrating the value of IT adoption on firm employment in a time of crisis. Taken together, our results indicate that the pandemic widens gaps in pre-pandemic levels of employment size and digitization level. This is consistent with an absorptive capacity mechanism where prior investments in digital technology cause the firm to incorporate new digital technologies more easily in the future.

5.2 Employee-Level Implications

In addition to firm-level trends, the results of our analysis also document several important trends for the digital productivity of employees. The pandemic brings about large, potentially persistent changes in how employees use digital technology. Long-run changes in technology use appear to be on the internal margin rather than the external margin, where workers are writing substantially more emails, but the number of workers per firm has not increased. In fact, the number of email accounts for large firms returned to pre-pandemic trends by the end of 2021. Furthermore, their digital productivity remained substantially above pre-pandemic trend levels, even as workers have begun to return to in-person work. Email writing is a crude measure of worker productivity, and so these results should be interpreted with caution. Still, to our knowledge, this is the first measure of individual employee output that can in some coherent sense be compared across firms and industries.

An essential question that remains is whether digital effort substitutes for previously non-digital work or whether workers in fact produce more digital content during the pandemic. We can think of these mechanisms as analogous to a substitution effect and an income effect that result from the change in the price of a good in economic consumer theory. Here we describe how these mechanisms potentially manifest in our digital technology adoption setting, but we leave the further study of these mechanisms to future work.

5.2.1 Substitution Effect

In our digital technology adoption setting, a substitution effect could result from the fact that during a pandemic in-person interactions are costly and workers substitute email writing for a substantial amount of what would otherwise be undocumented communication. One interpretation of our results is that the pandemic pushed firms to invest in digital infrastructure that is maintained even after in-person interactions resumed by the end of the year 2020. Under this explanation of the data, it is possible that the pandemic could lead to long-run economic growth insofar as it pushes firms to make investments they would otherwise not make in a non-pandemic world. This is reminiscent of a literature in economic growth suggesting that the origins of the European industrial revolution derive from labor supply shocks resulting from the Black Death (see for instance (*34*) and (*35*)). Such a technology shock may potentially improve welfare for both firms and workers in the long run, although obviously the pandemic has had many negative economic effects in the short run.

5.2.2 Income Effect

An income effect could result from the fact that a pandemic may deteriorate the bargaining position for workers. This could be, for example, because it is relatively costly for employees to find alternative work arrangements or take a financial risk in a pandemic. Conscious of this change, profit-maximizing employers may then lay off some employees and assign additional work to the remaining employees. Such a shock is necessarily associated with lower worker welfare. In contrast to the substitution effect outlined previously, there are no long-run worker productivity gains associated with this shock. Instead, there is a redistribution of welfare from workers to firms.

Ultimately, time will tell the extent to which the pandemic has impacted worker and firm welfare in the long run. We view our paper as an important step in understanding this phenomenon from the perspective of digital technology adoption.

6 Acknowledgments

We thank Sandi Ganguly, Feng Zhu and Miaomiao Zhang.

Funding: Harvard Business School Division of Research and Faulty Development.

Author contributions: E.A., S.P., P.W., A.W. conceived the project. H.L, S.P., P.W. participated in the data collection and analysis. All authors participated in the data interpretation and writing.

Data and materials availability: An anonymized version of the data supporting this study is retained indefinitely for scientific and academic purposes. The data are not publicly available due to employee privacy and other legal restrictions. The data are available from the authors on reasonable request and with permission from Microsoft Corporation.

References and Notes

1. E. Brynjolfsson, L. M. Hitt, *Journal of Economic Perspectives* **14**, 23 (2000).

2. N. Melville, K. Kraemer, V. Gurbaxani, *MIS Quarterly* **28**, 283 (2004).

3. L. Wu, *Information Systems Research* **24**, 30 (2013).

4. R. Adner, P. Puranam, F. Zhu, *Strategy Science* **4**, 253 (2019).

5. L. Wu, G. C. Kane, *Organization Science* **32**, 273 (2021).

6. L. M. Hitt, D. Wu, X. Zhou, *Journal of Management Information Systems* **19**, 71 (2002).

7. N. Greenana, J. Mairesse, *Economics of Innovation and New Technology* **9**, 275 (2000).

8. E. Brynjolfsson, L. Hitt, *Management Science* **42**, 541 (1996).

9. E. Brynjolfsson, L. Hitt, *Economics of Innovation and New technology* **3**, 183 (1995).

10. S. Aral, E. Brynjolfsson, D. Wu, *Proceedings of the 27th International Conference on Information Systems (ICIS)* pp. 1819–1839 (2006).

11. P. Tambe, L. M. Hitt, *Information Systems Research* **23**, 599 (2012).

12. T. F. Bresnahan, E. Brynjolfsson, L. M. Hitt, *Quarterly Journal of Economics* **117**, 339 (2002).

13. P. Huang, M. Ceccagnoli, C. Forman, D. Wu, *Available at SSRN 2243886* (2013).

14. N. Zolas, *et al.*, *National Bureau of Economic Research Working Paper* (2021).

15. C. Forman, *Management Science* **51**, 641 (2005).

16. P. Tambe, L. Hitt, D. Rock, E. Brynjolfsson, *National Bureau of Economic Research Working Paper* (2020).

17. G. F. Retana, C. Forman, S. Narasimhan, M. F. Niculescu, D. Wu, *MIS Quarterly* **42**, 961 (2018).

18. S. Dewan, K. L. Kraemer, *Management Science* **46**, 548 (2000).

19. K. Engemann, H. J. Wall, *Available at SSRN 1490041* (2009).

20. T. Alon, S. Coskun, M. Doepke, D. Koll, M. Tertilt, *National Bureau of Economic Research Working Paper* (2021).

21. J. M. Barrero, N. Bloom, S. J. Davis, *National Bureau of Economic Research Working Paper* (2020).

22. S. Y. T. Lee, M. Park, Y. Shin, *National Bureau of Economic Research Working Paper* (2021).

23. R. Cui, H. Ding, F. Zhu, *Manufacturing & Service Operations Management* (2021).

24. F. Almeida, J. D. Santos, J. A. Monteiro, *IEEE Engineering Management Review* **48**, 97 (2020).

25. J. J. Bai, E. Brynjolfsson, W. Jin, S. Steffen, C. Wan, *National Bureau of Economic Research Working Paper* (2021).

26. D. Papanikolaou, L. D. Schmidt, *National Bureau of Economic Research Working Paper* (2020).

27. E. Forsythe, L. B. Kahn, F. Lange, D. Wiczer, *Journal of Public Economics* **189**, 104238 (2020).

28. N. Bloom, R. S. Fletcher, E. Yeh, *National Bureau of Economic Research Working Paper* (2021).

29. W. Ding, R. Levine, C. Lin, W. Xie, *Journal of Financial Economics* **141**, 802 (2021).

30. R. Fairlie, *Journal of Economics & Management Strategy* **29**, 727 (2020).

31. S. Aral, E. Brynjolfsson, M. Van Alstyne, *Information Systems Research* **23**, 849 (2012).

32. S. Aral, M. Van Alstyne, *American Journal of Sociology* **117**, 90 (2011).

33. J. L. Bernal, S. Cummins, A. Gasparrini, *International Journal of Epidemiology* **46**, 348 (2017).

34. N. Voigtländer, H.-J. Voth, *Review of Economic Studies* **80**, 774 (2013).

35. G. Clark, *Journal of Demographic Economics* **82**, 139 (2016).

36. G. W. Imbens, D. B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences* (Cambridge University Press, 2015). New York, NY, USA.

7 Appendix

7.1 Methodology details

Formally, let i index the firms in our sample and t index the time periods as measured in months from March 2019 to December 2020. Let Y_{it} describe one of two firm outcomes: employee headcount and average emails per employee. The binary variable W_t is an indicator for "treatment". It denotes whether the COVID-19 virus has spread to the United States by time period t. Specifically, W_t takes the value of 1 if t is after March 2020 (the pandemic period) and a value of 0 if t is before March 2020 (the pre-pandemic period). We take March 2020 to be the start of the pandemic period. This treatment variable is the same for every firm in our study.

We use a potential outcomes framework to describe the approach (see generally (*36*)). Intuitively, we use potential outcomes to describe two different "worlds" experienced by every firm in our setting. The first world is the pre-pandemic world which we designate as having ended in March 2020. The second world is a pandemic world that we designate as having begun in March 2020. The potential outcomes modeling exercise imagines as if these two worlds are separate and exist in parallel for the full duration of March 2019 to December 2020. That is, there is both a non-pandemic world over these two years and a pandemic world over these two years. In each of these two worlds there exist sequences of outcomes for every firm: $Y_{it}(1)$ and $Y_{it}(0)$. The sequence $Y_{it}(1)$ describes the outcome for firm i in time period t in the pandemic world and $Y_{it}(0)$ is the corresponding sequence in the non-pandemic world. Importantly, $Y_{it}(1)$ exists in the time periods before (and after) March 2019 and $Y_{it}(0)$ exists in the time periods after (and before) March 2019.

The potential outcomes framework supposes that $Y_{it}(0)$ coincides with our observed outcome Y_{it} before March 2020 (i.e. when $W_t = 0$) but is otherwise unknown for the later time periods (i.e. when $W_t = 1$). Similarly, $Y_{it}(1)$ coincides with the observed outcome Y_{it} after March 2020 (i.e. when W_t = 1) but is otherwise unknown for the earlier time periods (i.e. when $W_t = 0$). If we observed both sequences of potential outcomes, we could directly characterize the causal impact of the pandemic on the outcome. One way to do this is to compute the percent change in the outcome due to the pandemic

$$\Delta_{it} = (Y_{it}(1) - Y_{it}(0))/Y_{it}(0). \tag{2}$$

This causal parameter Δ_{it} measures the relative difference in the non-pandemic and pandemic potential outcomes for any time period t and firm i. It cannot however be directly computed from the data because for every firm i and time period t we only observe either the potential outcomes $Y_{it}(1)$ or $Y_{it}(0)$ but never both simultaneously. The premise of the interrupted time series design is to address this problem by imputing the missing potential outcomes using a dynamic linear regression model. Formally, we model the potential outcome in the non-pandemic world as

$$Y_{it}(0) = \alpha_i + \beta_i t + u_{it} \tag{3}$$

where α_i is the expected non-pandemic potential outcome for firm i at time period $t = 0$, β_i is the change in the expected non-pandemic outcome over the course of one month for firm i, and

u_{it} is an idiosyncratic mean-zero error. We estimate the parameters α_i and β_i using data on the pre-pandemic outcomes via ordinary least squares and then impute the non-pandemic potential outcomes in the pandemic time periods using the formula

$$\hat{Y}_{it}(0) = (\hat{\alpha}_i + \hat{\beta}_i t)_+ \tag{4}$$

where $\hat{\alpha}_i$ and $\hat{\beta}_i$ are the usual ordinary least squares estimators and the function $(x)_+$ refers to the positive part of x. That is, if our model predicts a negative outcome for a given it-pair then we replace that prediction with a zero. Our estimate for the causal parameter Δ_{it} is then the plug-in

$$\hat{\Delta}_{it} = (Y_{it}(1) - \hat{Y}_{it}(0))/\hat{Y}_{it}(0). \tag{5}$$

In the results section below we report average values of $\hat{\Delta}_{it}$ for various types of firms in each of the time period. We show how the impact of the pandemic on firm employment and technology adoption depends on both firm size and the pre-pandemic level of digital technology adoption.