Internet Appendix for ”Stress Testing Banks’ Digital Capabilities: Evidence From the COVID-19 Pandemic”

Alan Kwan, Chen Lin, Vesa Pursiainen, and Mingzhu Tai

IA.I. Measuring Mobility Restrictions

In this section, we present the analysis underlying our specification for the SEVERE_ RESTRICTIONS variable used in the main analyses. As discussed in Section III.E, we start by collecting data on various types of mobility restrictions at the county level. We then perform an analysis to select the effective restriction categories in reducing county-level bank branch visits. Lastly, we validate our measure of restrictions by showing that this index strongly correlates with the Google Mobility Report.

Our first goal is to identify the mobility restrictions which might have an impact on branch visits. One challenge is that these mobility restrictions are not of equal importance to branch banking services, and are likely to cluster in time. For example, "religious gatherings being banned" is obviously unlikely to be relevant to our setting, since banks are typically considered as essential business services and it is not clear whether restrictions on non-essential services are relevant. To identify the ones which should be used to form a restrictions index, we run a variety of regression estimators to get a sense of which restrictions drive the most variation in branch visits. Specifically, we form a panel of the inverse hyperbolic sine of branch visits (de-meaned by branch and time fixed effects to be analogous to our main empirical specification). We use the LASSO model to select the restrictions which best explain branch visits. The LASSO model is appropriate because it adds a penalty term for large covariates, and thus will eliminate covariates that do not add incremental explanatory power. To pick the proper penalty parameter (λ), we run a cross-validation exercise and select the version of the model that produces the best cross-validated performance.[1] This exercise reveals three restrictions which matter: *retail*, *shelter-in-place* and *social-distancing*.

To help us build confidence that these are the most relevant mobility restrictions, we do two things. First, we use alternative regression methodologies and find similar results. Both OLS and cross-validated ridge regressions estimates point to these mobility restrictions as being the most

[1] Our sampling uses an arbitrary seed of 123 in this analysis. We also perform 10-fold cross-validation per statistical package defaults. Finally, we pick the value of λ that produces the minimum mean squared error plus one standard deviation, which is a rule-of-thumb advocated to produce a well-regularized model that is as simple as possible (Krstajic, Buturovic, Leahy, and Thomas (2014)).

important – although these methods, by design, do not eliminate covariates, this gives us assurance that our estimates are not sensitive to the peculiarities of the LASSO model. In particular, the OLS estimates affirm that only the three selected measures are *significantly* negatively correlated with branch visits while other measures are not. Finally, amalgamating these three methods, the same three measures are retrieved if we use the adaptive LASSO, initializing guesses for the β coefficients based on a preliminary OLS or ridge regression estimate. The adaptive LASSO proposed by Zou (2006) has been used in recent papers in the finance literature to aid with variable selection (e.g., Reeb and Zhao (2018)) and is considered superior to the LASSO on the dimension of bias.

Second, we validate our analysis using the Google Mobility restrictions at the county-week level (instead of the branch-level data at the branch-week level). We regress each component of the Google Community Mobility Reports on SEVERE_RESTRICTIONS. The Google Community Mobility Report (GCMR) uses Google-sourced location data pooling from Google's database based on Google Maps, Android phones and other data sources owned by Google. In the U.S., these data are available at the county level. The GCMR reports metrics of visits to different types of locations, including Workplaces, Residential, Retail & Recreation, Transit Stations, Parks, and Grocery & Pharmacy. We would expect the relevant restriction measures to be negatively correlated with visits to business establishments, parks and other outdoor locations and positively correlated with time spent at home. To obtain a single mobility measure, we also create a variable which aggregates all of the GCMR data series as their first principal component.

The results are reported in Table IA.2. Each column represents a different pillar of the Google Community Mobility report (except the first column, which is the first principal component of all the measures). We see that visits to public places and workplaces decrease significantly when mobility restrictions are introduced, while time spent at residences increases. Our PCA-based aggregated mobility measure is also significantly negatively associated with restrictions.[2]

[2]All of the measures generally contribute positively to the PCA, except for *Residential*, which contributes negatively. This is sensible as staying at home implies not visiting workplaces, outdoors, or retail establishments.

Table IA.1
Restriction Components and County-Level Mobility Patterns

This table reports the selected coefficients regressing branch visits on mobility restriction categories using the LASSO regression. The dependent variable, ihs(BRANCH_VISITS), is the inverse hyperbolic sine of bank branch visits, de-meaned by branch and week.

	ihs(BRANCH_VISITS)
Retail	-0.0280
Social distancing	-0.0076
Shelter in place	-0.0064
Business	Not selected by the model
Closing of public venues	Not selected by the model
Emergence	Not selected by the model
Gathering	Not selected by the model
Lockdown	Not selected by the model
Nonessential services closure	Not selected by the model
Religious gatherings banned	Not selected by the model
School closure	Not selected by the model
Stay at home	Not selected by the model

Table IA.2
Components of the Google Mobility Index

This table presents a county-week analysis relating our restrictions measure to the various metrics from the Google Community Mobility Report. The dependent variable is shown above each column. SEVERE_RESTRICTIONS is a dummy variable indicating that two of the three most important mobility restrictions (retail, social distancing, and shelter-in-place) are in place in the county. Heteroskedasticity-robust standard errors clustered by county are reported in parentheses.

	$PCA^{GoogleMobilityIndex}$	Retail/Recreation	Grocery/Pharmacy	Parks	Transit Stations	Workplace	Residential
	1	2	3	4	5	6	7
SEVERE_RESTRICTIONS	-0.5119***	-6.0587***	-4.6705***	-8.1282***	-6.4827***	-4.3391***	1.8330***
	(0.0346)	(0.2972)	(0.2718)	(1.1926)	(0.6328)	(0.1716)	(0.1021)
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,088	22,096	21,117	8,393	10,749	27,434	14,104
R^2	0.9499	0.9205	0.8043	0.6476	0.8479	0.9405	0.9539

IA.II. Mobility Restrictions and Bank IT

In this section, we perform additional dynamic event-study type analyses supporting our main results in Sections IV.A, IV.B, and IV.D. These analyses regress different outcome variables, including branch visits, website traffic, and deposits, on the interaction of IT_INDEX and SEVERE_RESTRICTIONS. These analyses imply a parallel trends assumption before restrictions take place – i.e., that banks with different IT_INDEX values have similar outcomes in the period before restrictions. Below, we present evidence to support this assumption and to validate the difference-in-differences design whereby restrictions affect our measures of branch visits and web traffic starting with the onset of the pandemic.

IA.II.A. Mobility Restrictions and Branch Visits

In Figure IA.1, we present our results on physical branch visits. Panel A estimates the overall effect of restrictions directly on branch visits. We use the estimator proposed by Borusyak, Jaravel, and Spiess (2021) to account for recent criticisms levied against standard staggered difference-in-difference designs in OLS. This supplements the pure time-series evidence that we presented in Figure 1, which does not exploit or rely on staggered differences. Panel B shows a similar analysis but compares the difference between high- versus low-IT banks. This can be viewed not just as a validation of the restriction index, but also as evidence that high-IT banks diverge from low-IT banks particularly as mobility restrictions take place.

Figure IA.1: Event-Time Analysis of Restrictions on Branch Visits

The dependent variable is *ihs(Branch visits)*. In Panel A, we present an event time plot of the average relationship between restrictions and bank branch visits as measured in the SafeGraph dataset relative to the imposition of mobility restrictions. The treated group consists of those bank branches for whom the county-level SEVERE_RESTRICTIONS are in place, defined as at least two of the three most important restrictions being in force. In Panel B, the treated group is the "high-IT" banks with SEVERE_RESTRICTIONS in place, where "high-IT" is defined as an IT_INDEX above the sample median. In both panels, point estimates are attained using the estimator proposed by Borusyak et al. (2021), controlling for branch linear trends. We plot 95% confidence intervals.

Panel A: Standalone Analysis of Restrictions on Bank Branch Visits

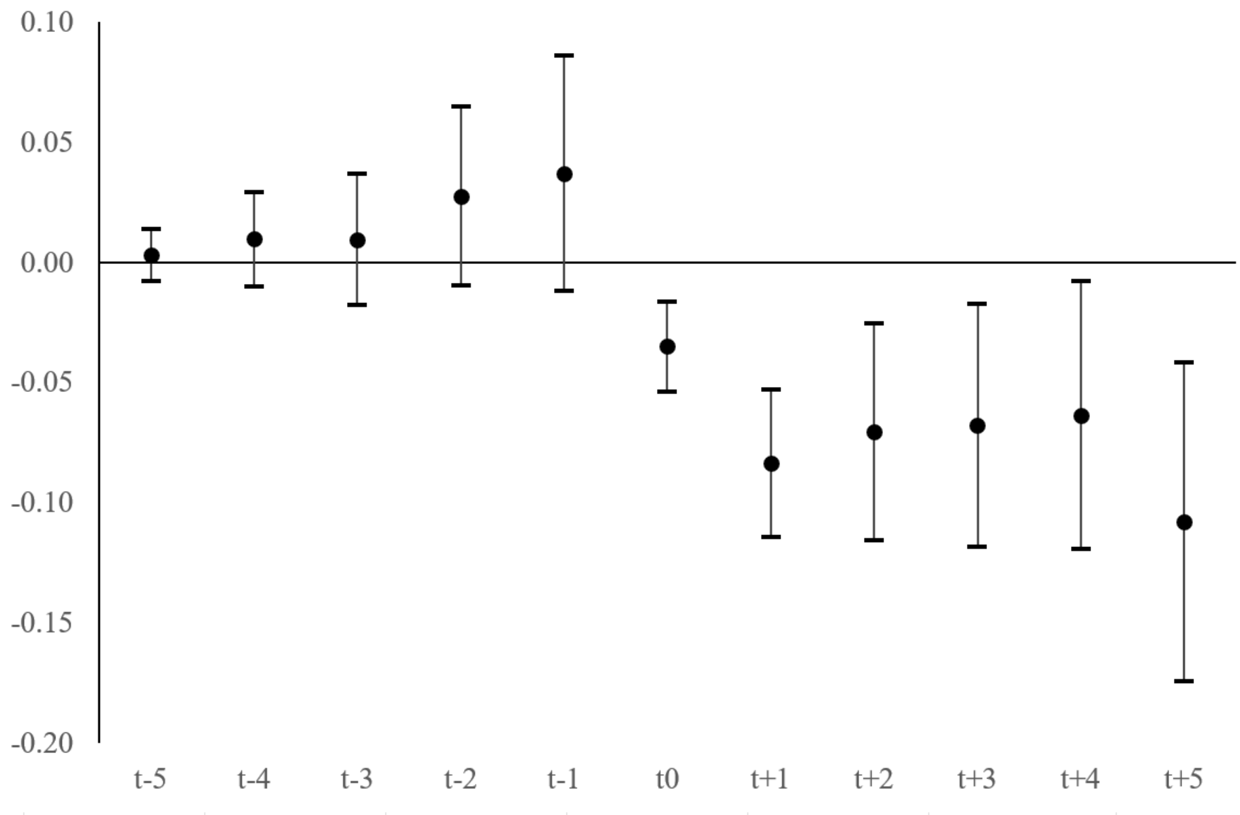

Panel B: High Versus Low IT Banks

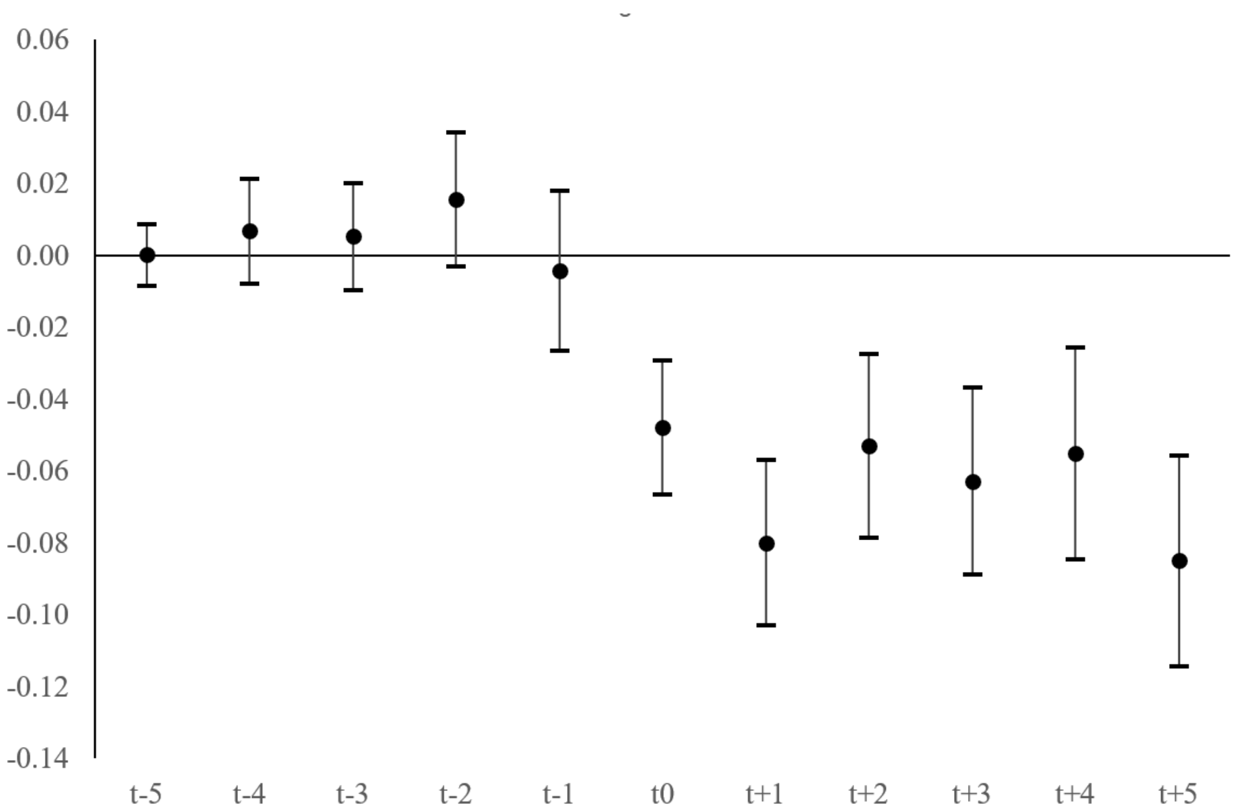

IA.II.B. Mobility Restrictions and Website Traffic

In Figure IA.2, we present our results on website traffic. Here we do not consider the analysis on the average effect of restrictions, as our measure is a *relative* rank index. We present the analysis comparing high- versus low-IT banks. What we find again mirrors the finding for branch visits: in event-time, restrictions generally coincide more with changes in the web traffic of high-IT banks, compared to low-IT banks. This effect intensifies in the weeks following restrictions but does not occur before the restrictions take place.

Figure IA.2: Event-Time Analysis of Restrictions on Website Traffic

In this figure, we present an event time plot of the relationship between restrictions and the difference between high-IT and low-IT banks' website traffic, where "high-IT" is defined as an IT_INDEX above the sample median. The dependent variable is a dummy taking the value of one if the bank's median AlexaRank during the week is below 100,000. The treated group consists of those high-IT banks for which at least half of branches are in counties under SEVERE_RESTRICTIONS. Point estimates are attained using the estimator proposed by Borusyak et al. (2021), controlling for bank linear trends. We plot 95% confidence intervals.

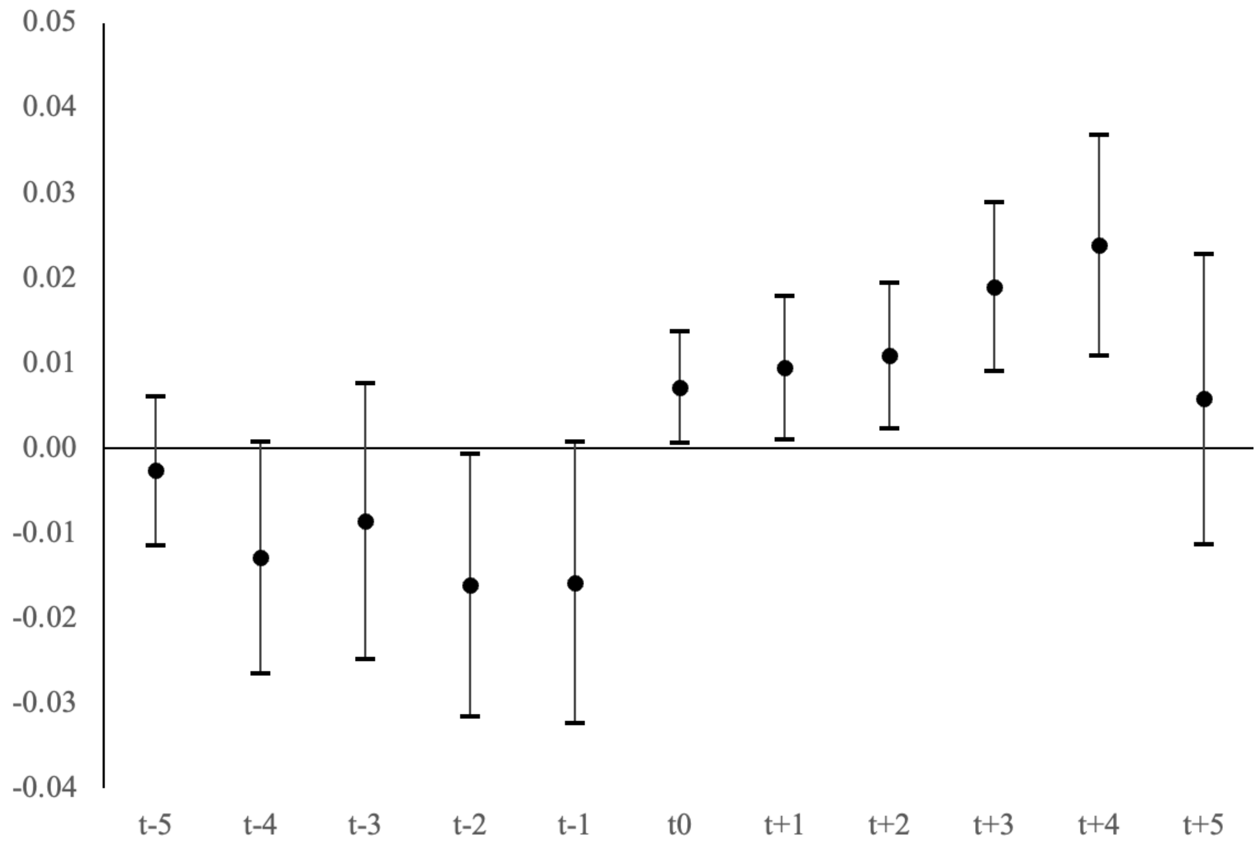

IA.II.C. Mobility Restrictions and Deposits

We next present results for quarterly bank-level deposits. However, given that we only observe quarterly data, this is effectively a one-shot treatment in 2020 Q1. Hence, we estimate the dynamic effect of bank IT using OLS regression, which is unbiased under a difference-in-differences setting as it is a one-time event instead of a staggered event. Figure IA.3 presents our results. The estimated effect of IT sharply increases during Q1 of 2020 and persists in Q2. This affirms our econometric results showing that the timing of the increase in the effect of IT matches the onset of the pandemic.

Figure IA.3: Event-time Analysis of Bank Deposits

In this table, we present the quarterly coefficient estimates for IT_INDEX from a regression where the dependent variable is ihs(DEPOSITS). We control for bank fixed effects, quarter fixed effects, and headquarter-county-time linear trends. We plot 95% confidence intervals.

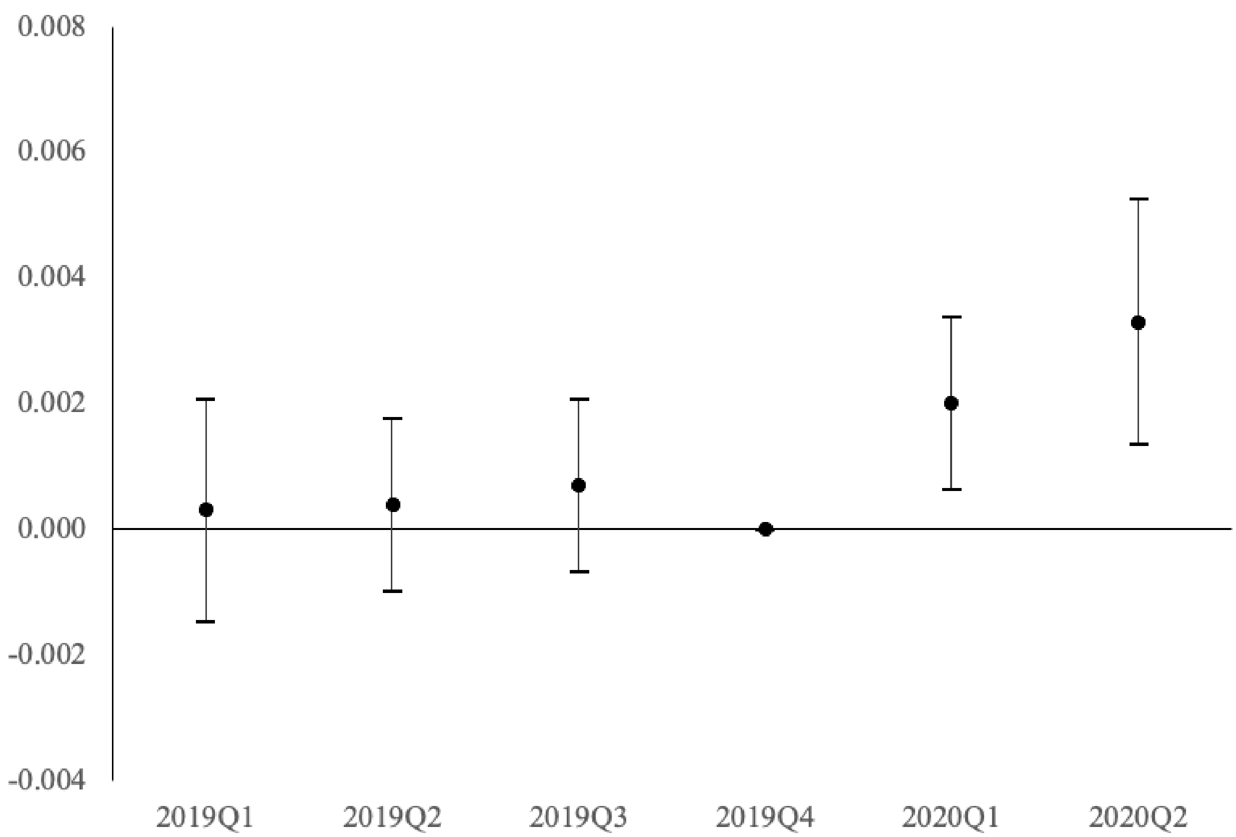

IA.III. The Role of Remote Work at Banks

In this section, we investigate to what extent the decline in branch visits we observe is explained by employees' remote work behavior. While remote work might be relevant for branch visits, it is less relevant for web traffic analysis as websites are customer-facing. To quantify remote work, we use the measure from Kwan and Matthies (2022). Kwan and Matthies (2022) partners with a data analytics company with visibility into thousands of internet publishers. The company specializes in identifying audiences and classifying internet content read. The classification operates by taking a sample of "truth-set" addresses from four categories (VPN, Residential, Business, and Mobile) and classifies them using a machine learning model, which is then combined with a rule-of-thumb IP rule set to form a classification of over 760 million IP addresses. For example, for residential IPs, residential proxies (businesses that are used to serve residential VPN traffic) have publicly known IPs. Prior computer science research suggests that all IPs in the same IP address block as such IP addresses are likely residential. For work IPs, IP addresses which displayed pre-pandemic work hour behavior are likely to be identified as work IPs. Mobile IPs are often those which have hundreds of users on them, most of which are mobile phones. About two-thirds of internet traffic observed in the dataset can be classified using a rule-of-thumb approach, and the remainder is classified with the model. Kwan and Matthies (2022) validates this measure.[3]

We define the variable REMOTE_WORK as the fraction of bank traffic during work hours (Monday through Friday, 8 a.m. to 5 p.m.) which comes from an IP address that Kwan and Matthies (2022) classifies as likely not a business IP address. Using this measure, we can classify the remote work traffic of our sample banks and hence quantify the extent of remote work performed at each of the banks. We also use the same measure for banks' SME customers in our analysis in Panel A of Table 9.

First, we study the relationship between mobility restrictions and bank remote work. The

[3]For example, this measure of remote work rises in March 2020 and shows a strong relationship with mobile phone mobility data. When people physically go to work, a percentage increase in workplace attendance leads to a 0.75% drop in remote work during daytime hours in the same county. At the firm level, remote work is more common in knowledge intensive industries in urban areas with high COVID-19 cases.

results are shown in Panel A of Table IA.3. We see that banks tend to shift toward remote work if they have better IT at the onset of restrictions. This validates the assumption that the remote work measure could partly react to restrictions, as expected.

However, as shown in Panel B where we repeat the analysis from Table 2 but control for remote work, remote work at the bank level does not significantly explain branch visits. This is intuitive, as in our sample, no more than 15% of our branch visitors could possibly be employees based on their common daytime location.[4] In addition, controlling for bank employee remote work, the effect of local mobility restrictions on branch footfall remains statistically and economically significant. Overall, these additional findings are consistent with the notion that both demand and supply effects play a role in shifting banking service from in-person to online.

[4]As mentioned in Section IV.A, we conduct diagnostic tests. If we assume employees of a bank in the data to have the same common daytime location as the bank branch's census block group (CBG), the fraction of visitors with the same CBG is around 15%.

Table IA.3
Remote Work and Branch Visits

In Panel A, we conduct a bank-week analysis, where the dependent variable is REMOTE_WORK, the measure of employee remote work constructed by Kwan and Matthies (2022). In Panel B, we repeat the analysis from Table 2, with ihs(BRANCH_VISITS) as the dependent variable, but controlling for REMOTE_WORK and its interaction with the bank's IT_INDEX. Heteroskedasticity-robust standard errors clustered by bank (in Panel A) or county and bank (Panel B) are reported in parentheses.

Panel A: Validation Test: Restrictions and Remote Work

	REMOTE_WORK	
	1	2
IT_INDEX × SEVERE_RESTR.		0.0038**
		(0.0018)
SEVERE_RESTRICTIONS	0.0063	-0.0012
	(0.0053)	(0.0064)
Bank FE	Yes	Yes
Week FE	Yes	Yes
Observations	67,527	67,527
R^2	0.6516	0.6517

Panel B: Branch Visits Controlling for Bank Remote Work

	ihs(BRANCH_VISITS)	
	1	2
IT_INDEX × SEVERE_RESTR.	-0.0226***	-0.0225***
	(0.0033)	(0.0038)
SEVERE_RESTRICTIONS	-0.0761***	-0.0761***
	(0.0120)	(0.0120)
IT_INDEX × REMOTE_WORK		-0.0004
		(0.0094)
REMOTE_WORK	0.0268	0.0283
	(0.0252)	(0.0226)
Branch FE	Yes	Yes
Week FE	Yes	Yes
Observations	504,157	504,157
R^2	0.8568	0.8568

IA.IV. Additional Analysis on Web Traffic

In this section, we perform an additional analysis of bank web traffic by using website traffic data from another provider, SimilarWeb, which combines a consumer panel with verified data directly gathered from websites and other data sources to estimate web traffic. We obtain data at the monthly level. Like AlexaRank, SimilarWeb aims to correct skews in the data to make their data more representative, and thus the use of complementary data will make our results more reliable.

In Table IA.4, we show results using the monthly data from SimilarWeb. The dependent variable is ihs(WEBSITE_VISITS), measuring the inverse hyperbolic sine of the number of monthly visits to each bank website. The results are qualitatively similar to the weekly results we present in Table 3. This shows that our findings are not sensitive to the rank-based construction methodology of the AlexaRank variable.

Table IA.4
Alternative Web Traffic Measure: SimilarWeb Sample

In this table, we present an analysis at the bank-month level. The outcome variable is ihs(WEBSITE_VISITS), the inverse hyperbolic sine of the number of website visits provided by SimilarWeb. Standard errors, reported in parentheses, are clustered by bank.

	ihs(WEBSITE_VISITS)	
	1	2
IT_INDEX × SEVERE_RESTR.		0.0176***
		(0.0063)
SEVERE_RESTRICTIONS	0.0013	0.0004
	(0.0317)	(0.0317)
Bank FE	Yes	Yes
Month FE	Yes	Yes
Observations	14,322	14,322
R^2	0.9475	0.9475

IA.V. Alternative IT Measures

We next perform our analysis using three alternative IT measures: IT_BUDGET, IT_STAFF, and IT_INDEX_OTHER which consists of the remainder of the full 64 technologies that we did not include in our main IT_INDEX. The basic idea of this is to reduce measurement error deriving from simply aggregating the number of technologies. For IT staff, for a given establishment, we take the midpoint of the listed range and normalize the number by total employment. For IT budget, we take the inverse hyperbolic sine of the IT budget per employee.

As mentioned before, despite the fact that our main results generally work for various measures of IT, we emphasize our main IT_INDEX measure for two reasons. First, the IT_INDEX measure is directly verified from company websites, employee resumes and job postings, whereas IT_STAFF and *IT Budget* are estimated. Second, the IT_INDEX allows us to isolate the specific technologies of interest. In our tests, we argue that the 14 IT components that we isolate outperform the remainder 50 technologies that we do not include in explaining our various outcome variables.

In what follows, we report our main analyses using the alternative IT measures. In Table IA.5, we present our results on branch visits. In Table IA.6, we present our results on web traffic. In Table IA.8, we present our results on PPP lending. In Table IA.9, we present our results on deposits.

Broadly, our analyses corroborate two general sets of findings. First, the variable IT_INDEX_OTHER produces qualitatively similar, but generally weaker results than using the IT_INDEX. In specifications where we horse race the two variables, our main effect is driven entirely by the main IT_INDEX consisting of 14 variables that we include. This implies that the 14 technologies we chose tend to be most relevant for driving the outcome variables of interest. We chose these technologies to be particularly relevant for remote work, digital operations and customer communications, and so this result broadly affirms our interpretation that we are capturing a technology-driven shift, given that less relevant technologies have less explanatory power. Second, using alternative measures of IT (budget and staff) produces qualitatively similar inferences as if we use our IT_INDEX. This reduces the concern that our measure of IT is confounded with measurement error.

Table IA.5
Alternative IT measures: Branch Visits during Mobility Restrictions

This table repeats the analysis in Table 2 by using alternative measures for the IT capability of each bank. The alternative measures include a similar IT_INDEX measured based on other less relevant technologies, the number of bank IT staff (per employee), and the inverse hyperbolic sine of the amount of IT budget normalized by bank employee number. The dependent variable, ihs(BRANCH_VISITS) is the inverse hyperbolic sine of the number of visits recorded in Safegraph's Places of Interest file. The SEVERE_RESTRICTIONS variable is a dummy variable indicating the county is in a state of "severe restrictions". The sample period is from January to April, 2020. Heteroskedasticity-robust standard errors clustered by county and bank are reported in parentheses.

	ihs(BRANCH_VISITS)			
	1	2	3	4
IT_INDEX_OTHER × SEVERE_RESTR.	-0.0090***	-0.0017		
	(0.0013)	(0.0036)		
IT_INDEX × SEVERE_RESTR.		-0.0185**		
		(0.0088)		
IT_STAFF × SEVERE_RESTR.			-0.3342*	
			(0.1790)	
ihs(IT_BUDGET) × SEVERE_RESTR.				-0.0305***
				(0.0088)
SEVERE_RESTRICTIONS	-0.0760***	-0.0760***	-0.0794***	-0.0795***
	(0.0126)	(0.0119)	(0.0094)	(0.0097)
Branch FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
Observations	522,370	522,370	522,340	522,370
R^2	0.8567	0.8568	0.8561	0.8562

Table IA.6
Alternative IT Measures: Website Traffic during Mobility Restrictions

In this table, we repeat the analysis from Table 4 using alternative IT measures. The alternative measures include a similar IT_INDEX measured based on other less relevant technologies, the number of bank IT staff (per employee), and the inverse hyperbolic sine of the IT budget normalized by bank employee number. The dependent variable is a dummy indicator for whether the median rank for the bank-week is in the AlexaRank top 100,000. The SEVERE_RESTRICTIONS variable is a dummy variable indicating the county is in a state of "severe restrictions" for at least half of bank branches during that bank-week. Heteroskedasticity-robust standard errors clustered by bank are reported in parentheses.

	MEDIAN_RANK≤100k			
	1	2	3	4
IT_INDEX_OTHER × SEVERE_RESTR.	0.0016**	-0.0008		
	(0.0008)	(0.0009)		
IT_INDEX × SEVERE_RESTR.		0.0063***		
		(0.0019)		
IT_STAFF × SEVERE_RESTR.			0.0425*	
			(0.0228)	
ihs(IT_BUDGET) × SEVERE_RESTR.				0.0033**
				(0.0016)
SEVERE_RESTRICTIONS	0.0042*	0.0039*	0.0042*	0.0041*
	(0.0022)	(0.0022)	(0.0023)	(0.0022)
Bank FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
Observations	29,946	29,946	29,946	29,946
R^2	0.8881	0.8882	0.8880	0.8880

Table IA.7
Alternative IT Measures: SimilarWeb Sample

In this table, we repeat the analysis from Table IA.4 using alternative IT measures. The alternative measures include a similar IT_INDEX measured based on other less relevant technologies, the number of bank IT staff (per employee), and the inverse hyperbolic sine of IT budget normalized by bank employee number. The outcome variable is the inverse hyperbolic sine of the number of website visits provided by SimilarWeb. Heteroskedasticity-robust standard errors are reported in parentheses are clustered by bank.

	ihs(WEBSITE_VISITS)			
	1	2	3	4
IT_INDEX_OTHER × SEVERE_RESTR.	0.0058*	-0.0014		
	(0.0032)	(0.0057)		
IT_INDEX × SEVERE_RESTR.		0.0198*		
		(0.0115)		
IT_STAFF × SEVERE_RESTR.			0.8890***	
			(0.1706)	
ihs(IT_BUDGET) × SEVERE_RESTR.				0.0464***
				(0.0085)
SEVERE_RESTRICTIONS	0.0007	0.0004	-0.0015	-0.0019
	(0.0317)	(0.0317)	(0.0317)	(0.0317)
Bank FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	14,322	14,322	14,322	14,322
R^2	0.9475	0.9475	0.9476	0.9477

Table IA.8
Alternative IT Measures: PPP Lending

This table repeats the analysis in Table 5 by using alternative measures for the IT capability of each bank. The alternative measures include a similar IT_INDEX measured based on other less relevant technologies, the number of bank IT staff (per employee), and the inverse hyperbolic sine of IT budget normalized by bank employee number. The dependent variable is *ihs(PPP)*, the inverse hyperbolic sine of the amount of PPP loans originated by a bank in a county. Heteroskedasticity-robust standard errors clustered by county and bank are reported in parentheses.

	ihs(PPP)			
	1	2	3	4
IT_INDEX_OTHER	0.0207	-0.0308		
	(0.0164)	(0.0271)		
IT_INDEX		0.1488***		
		(0.0516)		
IT_STAFF			2.5054**	
			(1.1669)	
ihs(IT_BUDGET)				0.2482***
				(0.0898)
Bank-County Controls	Yes	Yes	Yes	Yes
Bank Controls	Yes	Yes	Yes	Yes
HQ State FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes
Observations	30,166	30,166	30,166	30,166
R^2	0.6496	0.6527	0.6503	0.6525

Table IA.9
Alternative IT Measures: Deposit Analysis

This table repeats the analysis in Table 7 by using alternative measures for the IT capability of each bank. The alternative measures include a similar IT_INDEX measured based on other less relevant technologies, the number of bank IT staff (per employee), and the inverse hyperbolic sine of IT budget normalized by bank employee number. The dependent variable is ihs(DEPOSITS), the inverse hyperbolic sine of the quarterly bank-level deposits. The sample period is 2019 Q1 to 2020 Q2. Q1_2020_ONWARD is a dummy variable that equals one for the first two quarters of 2020. *Bank Controls* include size (measured by the inverse hyperbolic sine of total assets), capitalization (measured by both equity/asset ratio and tier-1 capital ratio), profitability (measured by return on equity and cost/income ratio), funding cost, and personnel costs. Heteroskedasticity-robust standard errors clustered by bank are reported in parentheses.

	ihs(DEPOSITS)			
	1	2	3	4
IT_INDEX_OTHER × Q1_2020_ONWARD	0.0013***	0.0003		
	(0.0003)	(0.0004)		
IT_INDEX × Q1_2020_ONWARD		0.0028***		
		(0.0008)		
IT_STAFF × Q1_2020_ONWARD			0.0702***	
			(0.0122)	
ihs(IT_BUDGET) × Q1_2020_ONWARD				0.0035***
				(0.0006)
Bank Controls	Yes	Yes	Yes	Yes
Bank FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Observations	27,472	27,472	27,469	27,472
R^2	0.9993	0.9993	0.9993	0.9993

IA.VI. Alternative Formulations of Restrictions Measure

We next examine alternative formulations of our restrictions index. Instead of a dummy variable indicating severe restrictions, as an alternative specification, we look at counting the three most important restrictions we identified as well as all other available restrictions. This exercise serves two purposes. First, it alleviates the concern that measurement error in how we define restrictions impacts our results. Second, we can compare the coefficients across using two different measures of restrictions – one that only considers the three chosen restrictions versus one that considers all.

In Table IA.10 we repeat the analysis from Table 2 but count the number of restrictions in force instead of defining a strict threshold index. Likewise, Table IA.11 repeats the analysis from Table 4 but counting the number of restrictions in force instead of defining a strict threshold index. Together, the analyses generally indicate that there is a significant relationship between restrictions and the outcome variables similar to when we use the threshold indicator in our main results.

Table IA.10
Count of Restrictions in Force and Branch Visits

In this table, we redo the analysis from Table 2 but counting the number of restrictions in force instead of defining a strict threshold index. The first three columns count the number of the three main restriction categories we use in the main analyses. The last three columns count the total number of restrictions. Standard errors are reported in parentheses and clustered by bank and county.

	ihs(BRANCH_VISITS)					
	1	2	3	4	5	6
IT_INDEX × RESTRICTIONS		-0.0081***	-0.0024**		-0.0023***	-0.0007**
		(0.0012)	(0.0011)		(0.0004)	(0.0003)
RESTRICTIONS	-0.0555***	-0.0464***		-0.0123***	-0.0098***	
	(0.0062)	(0.0067)		(0.0025)	(0.0026)	
Restrictions Measure		Main 3			All Restr.	
Branch FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
County-Week FE	No	No	Yes	No	No	Yes
Observations	522,370	522,370	515,256	522,370	522,370	515,256
R^2	0.8563	0.8569	0.8685	0.8561	0.8568	0.8685

Table IA.11
Count of Restrictions in Force and Web Traffic

In this table, we redo the analysis from Table 4 but counting the number of restrictions in force instead of defining a strict threshold index. The first two columns count the number of the three main restriction categories we use in the main analyses. The last two columns count the total number of restrictions. Standard errors are reported in parentheses and clustered by bank.

	MEDIAN_RANK≤100k			
	1	2	1	2
IT_INDEX × RESTRICTIONS		0.0016***		0.0004***
		(0.0006)		(0.0001)
RESTRICTIONS	0.0009	0.0005	0.0005	0.0004
	(0.0010)	(0.0009)	(0.0004)	(0.0004)
Restrictions Measure	Main 3		All Restr.	
Bank FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
Observations	29,946	29,946	29,946	29,946
R^2	0.8880	0.8882	0.8880	0.8882

IA.VII. Additional Information: Customer Review Analysis

We collect data from the Android app store and Apple iOS app store. These data allow us to find the version of an app of a bank, if it exists. To find the appropriate app, we obtain a public data set from Github, which is a scrape of the entire Google and Apple app stores as of the middle of 2020.[5] We build a scraper for each platform.

One limitation worth noting is that the Android app store does not let us isolate reviews made in the U.S. Thus, some of our data likely come from bank subsidiaries in other countries. For example, HSBC operates in the U.S. but has a significant presence in Hong Kong and the United Kingdom. Citibank has ATMs, and thus likely a mobile site, in 20 different countries. The Apple app store allows us to geographically target reviews, so we target the U.S. app store directly. Another limitation worth noting is that it is likely both app stores only provide ratings with text, as the app store tends to display non-empty ratings. Thus, we only have a subset of the total ratings where the mobile app user took the effort to write something, however short. One could argue this is the sample of interest in that such ratings are likelier to be more meaningful. Nevertheless we have a significant, but incomplete subset of ratings.

The reviews allow us to extract the rating (a numeric scale from 1–5), time the review was issued (which may not be the time the app was downloaded), and the texts. The text of the review permits us to conduct textual analysis to ascertain whether the score is associated with a negative or a positive reaction. To conduct our textual analysis, we apply a state-of-the-art natural language processing algorithm recently open-sourced by Facebook called the BART Natural Language Inference model. It is a "zero shot classifier" meaning that it is designed to be useful for extrapolating to an arbitrary label, without being able to see the labels beforehand.

Standard classification problems typically require one to manually train an algorithm to do a particular classification, labeling a "training set" of data by hand, and calibrating a machine learning model until it achieves acceptable performance. The benefit of a zero-shot approach is

[5]We thank Gautham Prakash for making his data available via Github.

that it sidesteps the requirement of having to tailor an algorithm, as the algorithm is to designed to be extrapolated beyond its original training set and instead be amenable to an arbitrary label. In other words, one could theoretically pass over a phrase like ”courteous service,” even though that was not originally label that the model was trained on, because ”courteous service” is likely related to words that were in the original model such as ”courteous” and ”service.”[6]

The second piece of this methodology is “BART,” called “Bidirectional Auto-regressive Transformers” released by Facebook which builds upon Google’s BERT (Bidirectional Encoder Representations from Transformers). Essentially, instead of keeping specific words or n-grams, BERT and BART use a variety of machine learning methodologies to encode sequences of text into a contextual representation around a word. For example, suppose we are given two sentences: 1) “The man was accused of robbing a bank.” and 2) “The man went fishing by the bank of the river.” Simply counting word occurrences would produce the same word embedded for the word “bank” in both sentences, while under BERT the word embedding for “bank” would be different for each sentence. BART generalizes BERT and is widely considered to be the next advancement of BERT.

We demonstrate this with examples. To capture the quality of “service,” we created eight labels, namely: easy to use, effective, unintuitive, reliable, lacks features, doesn’t work, slow service, and aesthetically pleasing. These are four positive and four negative labels. Unintuitive and aesthetically pleasing refer to appearance. These labels are motivated by reading a few hundred reviews and getting a general sense of the common types of complaints or commendations given to reviews.

We apply these tools to all 2.5 million reviews that we collect. Training this model is computationally expensive and requires substantial graphics processing unit resources, given that the model we use has 400 million hyperparameters. First, in Table IA.12 we document the most common key words derived using this methodology. Second, in Table IA.13 we show example customer reviews and the top three labels. After reading several dozens of these examples, we believe that the output it provides is intuitive and reasonable, and it permits us a succinct economic interpretation.

[6]The package we use is hosted at this repository.

Table IA.12
Frequency of Label in the Top Position

Label	N
effective	20.805
aesthetically pleasing	20.662
doesnt work	15.521
easy to use	15.512
reliable	14.676
lacks features	4.628
slow service	4.385
unintuitive	3.811

Table IA.13
Examples of Customer Review Textual Analysis Using the BART MLNI Algorithm

We show 10 reviews randomly drawn. We created 10 labels: *"easy to use, effective, unintuitive, reliable, lacks features, doesnt work, slow service, aesthetically pleasing."* We list the top three topics sorted by label probability/relevancy. The relevancy score is shown in parentheses. The probabilities of all the eight labels sum to 1.

Review	1	2	3
Hopeless App Unable to reset Login Pin, After i enter debit card details and enter new Pin in required format, unable to submit (Submit button is inactive	doesnt work(0.598)	lacks features(0.299)	unintuitive(0.073)
wells fargo is the best bank and has the best fraud protection around. the app is very easy to navigate and helps me stay on top of my finances.	easy to use(0.469)	reliable(0.345)	effective(0.17)
Great bank I recommend them to eanyone	reliable(0.6)	effective(0.25)	aesthetically pleasing(0.12)
Banking is great when you can see what's going on with your money at all times.	effective(0.544)	reliable(0.264)	aesthetically pleasing(0.082)
Only worked half the time, and when it did worl, it was super slow. Now it just does not work at all.	slow service(0.582)	doesnt work(0.357)	unintuitive(0.025)
I live the ability to be able to access my accounts with no issues except for the finger ID was removed other than that great tool to keep up with everyday expenses.	effective(0.57)	reliable(0.125)	unintuitive(0.082)
Awesone way to make deposits. It's very useful	effective(0.653)	reliable(0.197)	easy to use(0.077)
Every 3 out of 5 times I open the app I get a modal popup saying, "Banking is unavailable for maintenance, please try again later." I work at a software company and having an application in production being down this often is absolutely unacceptable, especially something as critical as a customer's way to access their personal finances. When I get the time I will be switching banks.	doesnt work(0.365)	slow service(0.344)	effective(0.148)
works great never a problem and very easy to access and use	reliable(0.508)	effective(0.271)	easy to use(0.203)
Thought the update would fix this, but no matter what I do the pictures of my checks aren't accepted. The only reason I have this app and it no longer works.	doesnt work(0.776)	effective(0.066)	unintuitive(0.056)
I hate that I have to see the daily balance after each transaction on my main account. if I do 15 transactions I have to see the daily bslsnce 15 times!!! I don't want to scroll for 20 pages to see my last 20 transactions. it wasn't like that a couple of months ago. it was much much much much better	unintuitive(0.365)	doesnt work(0.187)	effective(0.138)

References

Borusyak, K.; X. Jaravel; and J. Spiess. "Revisiting event study designs: Robust and efficient estimation." arXiv preprint arXiv:2108.12419 (2021).

Krstajic, D.; L. J. Buturovic; D. E. Leahy; and S. Thomas. "Cross-validation pitfalls when selecting and assessing regression and classification models." *Journal of Cheminformatics*, 6 (2014), 1–15.

Kwan, A., and B. Matthies. "Measuring Remote Knowledge Work Using Big Data." Working Paper (2022).

Reeb, D. M., and W. Zhao. "Dissecting innovation." Working paper (2018).

Zou, H. "The adaptive lasso and its oracle properties." *Journal of the American Statistical Association*, 101 (2006), 1418–1429.