

The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data^{*}

Raj Chetty, John N. Friedman, Nathaniel Hendren, Michael Stepner,
and the Opportunity Insights Team[†]

October 2020

Abstract

We build a publicly available database that tracks economic activity at a granular level in real time using anonymized data from private companies. We report daily statistics on consumer spending, business revenues, employment rates, and other key indicators disaggregated by ZIP code, industry, income group, and business size. Using these data, we study the mechanisms through which COVID-19 affected the economy by analyzing heterogeneity in its impacts. We first show that high-income individuals reduced spending sharply in mid-March 2020, particularly in areas with high rates of COVID-19 infection and in sectors that require in-person interaction. This reduction in spending greatly reduced the revenues of businesses that cater to high-income households in person, notably small businesses in affluent ZIP codes. These businesses laid off many of their employees, leading to widespread job losses especially among low-wage workers in affluent areas. High-wage workers experienced a “V-shaped” recession that lasted a few weeks in terms of employment loss, whereas low-wage workers experienced much larger job losses that persisted for several months. Building on this diagnostic analysis, we use event study designs to estimate the causal effects of policies aimed at mitigating the adverse impacts of COVID-19. State-ordered reopenings of economies have small impacts on spending and employment. Stimulus payments to low-income households increased consumer spending sharply, but little of this increased spending flowed to businesses most affected by the COVID-19 shock, dampening its impacts on employment. Paycheck Protection Program loans increased employment at small businesses by only 2%, implying a cost of \$450,000 per job saved. These results suggest that traditional macroeconomic tools – stimulating aggregate demand or providing liquidity to businesses – have diminished capacity to restore employment when consumer spending is constrained by health concerns. During a pandemic, it may be more fruitful to mitigate economic hardship through social insurance. More broadly, this analysis shows how public statistics constructed from private sector data can support many research and policy analyses without compromising privacy, providing a new tool for real time empirical macroeconomics.

^{*}An earlier draft of this paper was circulated under the title “How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data.” We thank Gabriel Chodorow-Reich, Jason Furman, Xavier Jaravel, Erik Hurst, Lawrence Katz, Emmanuel Saez, Ludwig Straub, Danny Yagan, and numerous seminar participants for helpful comments. We also thank the corporate partners who provided the underlying data used in the Economic Tracker: Affinity Solutions (especially Atul Chadha and Arun Rajagopal), Burning Glass (Anton Libsch and Bledi Taska), CoinOut (Jeff Witten), Earnin (Arun Natesan and Ram Palaniappan), Homebase (Ray Sandza and Andrew Vogeley), Intuit (Christina Foo and Krithika Swaminathan), Kronos (David Gilbertson), Paychex (Mike Nichols and Shadi Sifain), Womply (Derek Doel and Ryan Thorpe), and Zearn (Billy McRae and Shaline Sharma). We are very grateful to Ryan Rippel of the Gates Foundation for his support in launching this project and to Gregory Bruich for early conversations that helped spark this work. The work was funded by the Chan-Zuckerberg Initiative, Bill & Melinda Gates Foundation, Overdeck Family Foundation, and Andrew and Melora Balson. The project was approved under Harvard University IRB 20-0586.

[†]The Opportunity Insights Economic Tracker Team consists of Camille Baker, Harvey Barnhard, Matthew Bell, Gregory Bruich, Tina Chelidze, Lucas Chu, Westley Cineus, Sebi Devlin-Foltz, Michael Droste, Shannon Felton Spence, Dhruv Gaur, Federico Gonzalez, Rayshauna Gray, Abby Hiller, Matthew Jacob, Tyler Jacobson, Margaret Kallus, Laura Kincaide, Caitlin Kupsc, Sarah LaBauve, Maddie Marino, Kai Matheson, Christian Mott, Kate Musen, Danny Onorato, Sarah Oppenheimer, Trina Ott, Lynn Overmann, Max Pienkny, Jeremiah Prince, Sebastian Puerta, Daniel Reuter, Peter Ruhm, Tom Rutter, Emanuel Schertz, Kamelia Stavreva, James Stratton, Clare Suter, Elizabeth Thach, Nicolaj Thor, Amanda Wahlers, Kristen Watkins, Alanna Williams, David Williams, Chase Williamson, Shady Yassin, and Ruby Zhang.

I Introduction

Since the pioneering work of Kuznets (1941), macroeconomic policy decisions have been made on the basis of publicly available statistics constructed from recurring surveys of households and businesses conducted by the federal government. Although such statistics have great value for understanding total economic activity, they have two limitations. First, survey-based data are typically available only at low frequencies, often with a significant time lag. For example, disaggregated data on consumer spending are only available at a quarterly frequency with a one year lag in the Consumer Expenditure Survey (CEX). Second, such statistics typically cannot be used to assess granular variation across geographies or subgroups; due to relatively small sample sizes, most statistics are typically reported only at the national or state level and breakdowns for demographic subgroups or sectors are generally unavailable. Because of these two limitations, existing macroeconomic statistics generally do not permit precise diagnostic analyses of the sources of economic fluctuations or the impacts of macroeconomic policies in a timely manner.

In this paper, we address these challenges by (1) building a publicly accessible [platform](#) that measures spending, employment, and other outcomes at a high-frequency, granular level using anonymized data from private companies and (2) demonstrating how these new data can be used to obtain insights into the effects of the coronavirus pandemic (COVID-19) and subsequent stabilization policies in near-real-time – within three weeks of the shock or policy change of interest.¹

We organize the paper into three parts. The first part describes how we construct statistics on consumer spending, business revenues, employment rates, job postings, and other key indicators – disaggregated by area (ZIP code or county), industry, income level, and business size – by combining data from credit card processors, payroll firms, and financial services firms. The main challenge in using private sector data sources to measure economic activity is a tension between research value and privacy protection. For research, it is beneficial to use raw, disaggregated data – ideally down to the individual consumer or business level – to maximize precision and flexibility of research designs. But from a privacy perspective, it is preferable to aggregate and mask data to reduce the risk of disclosure of information about businesses or their clients. To balance these conflicting interests, one must construct statistics that are sufficiently aggregated and masked to mitigate privacy concerns yet sufficiently granular to support research and policy analysis.²

1. Our approach builds on a recent literature that analyzes the utility of private sector data sources to measure economic activity, notably a set of studies collected in Abraham et al. (2019) that we discuss further below. To our knowledge, the present study constructs the first *public* platform from such private data sources.

2. An alternative approach – pursued by several recent studies summarized at the end of this section – is to use

We navigate this tradeoff by combining several statistical methods: reporting only changes since January 2020 (rather than raw levels), masking certain cells, and pooling data from multiple companies to comply with regulations governing the disclosure of material non-public information. We then clean the raw transaction data by removing data artifacts (e.g., breaks that arise from changes in platforms) and smoothing seasonal fluctuations. Finally, we address the challenge that these statistics reflect the behavior of the clients of the firms from which we obtain data and hence may not be representative of the broader population. To minimize potential selection biases, we start by obtaining data from companies that have large samples (e.g., at least one million individuals) and span well-defined sectors or subgroups (e.g., small businesses, bottom-wage-quartile workers). We then compare each time series to publicly available benchmarks based on representative surveys and proceed to use only the series that track publicly available data closely.³ After establishing these protocols, we report the final statistics using an automated pipeline that ingests data from businesses and reports statistics shortly after the relevant transactions occur (typically within one week).

In the second part of the paper, we use these new public statistics to analyze the economic impacts of COVID-19. National accounts data reveal that most of the initial reduction in GDP following the COVID-19 shock came from a reduction in consumer spending (rather than business investment, government purchases, or exports). We therefore begin our analysis by examining the drivers of changes in consumer spending, focusing primarily on credit and debit card spending. We first establish that card spending closely tracks historical benchmarks on retail spending and services, which together constitute a large fraction of the reduction in total spending in the national accounts. We then show that the vast majority of the reduction in consumer spending in the U.S. came from reduced spending by *high-income* households. High-income households cut spending more heavily when the COVID shock first hit in mid-March and increased spending more gradually

confidential private data sources for research analyses under non-disclosure agreements with companies. Although a valuable complement to the public data approach we pursue here, that approach has limits in terms of scale and timeliness. The need to write contracts and acquire data from each company separately typically leads most studies to use one or two datasets and limits the number of researchers and policymakers who can analyze data most relevant to their region or sector of interest in a timely manner. Our goal is to assess whether one can eliminate the need to write such contracts by producing aggregated public statistics that can deliver analogous insights.

3. This benchmarking work proves to be quite important in constructing representative series. For example, many studies have used data from Homebase, a company that helps small businesses track their employees hours (e.g., Bartik et al. 2020, Bartlett and Morse 2020, Granja et al. 2020, Altonji et al. 2020), to study employment in the COVID pandemic. As noted by Bartik et al. (2020), the time series patterns in the Homebase data differ significantly from representative statistics on small business employment (although the patterns are generally similar for the sectors it covers). We therefore turn to other sources of employment data to produce publicly available series that track representative benchmarks more closely and are now available for future research.

since that point. As of July 31st, 47% of the reduction in spending since January came from households in the top quartile of the income distribution; only 7% had come from households in the bottom income quartile.⁴ This is both because the rich account for a larger share of total spending to begin with, and because high-income households were spending 16% less in July than they were in January, whereas low-income households were spending only 5% less.

Most of the reduction in spending is accounted for by reduced spending on goods or services that require in-person physical interaction and thereby carry a risk of COVID infection, such as hotels, transportation, and food services, consistent with the findings of Alexander and Karger (2020). The composition of spending cuts – with a large reduction in services – differs sharply from that in prior recessions, where service spending was essentially unchanged and durable goods spending fell sharply. Zooming into specific subcategories, we find that spending on luxury goods that do not require physical contact – such as landscaping services or home swimming pools – did not fall, while spending at salons and restaurants plummeted. Businesses that offer fewer in person services, such as financial and professional services firms, also experienced much smaller losses. The fact that spending fell in proportion to the degree of in-person exposure required across sectors suggests that the reduction in spending by the rich was driven primarily by health concerns rather than a reduction in income or wealth. Indeed, we find that the incomes of the rich fell relatively little in the COVID recession. Consistent with the centrality of health concerns, we find that the reductions in spending and time spent outside home were larger in high-income, high-density areas with higher rates of COVID infection, perhaps because high-income individuals can self-isolate more easily (e.g., by substituting to remote work). Together, these results suggest that consumer spending in the pandemic fell because of changes in firms’ ability to supply certain goods (e.g., restaurant meals that carry no health risk) rather than because of a reduction in purchasing power.⁵

Next, we turn to the impacts of the consumer spending shock on businesses. To do so, we exploit the fact that many of the sectors in which spending fell most are non-tradable goods produced by

4. We impute income as the median household income (based on Census data) in the cardholder’s ZIP code. We verify the quality of this imputation procedure by showing that our estimates of the gap in spending reductions by income group are aligned with those of Cox et al. (2020), who observe income directly for JPMorgan Chase clients, as of mid-April 2020, the last date available in their series. We find that spending levels of low-income households increased much more sharply than those of high-income households since mid-April largely as a result of stimulus payments.

5. This explanation may appear to be inconsistent with the fact that the Consumer Price Index (CPI) shows little increase in inflation, given that one would expect a supply shock to increase prices. However, the CPI likely understates inflation in the current crisis because it does not capture the extreme shifts in the consumption bundle that have occurred as a result of the COVID crisis (Cavallo 2020).

small local businesses (e.g., restaurants) which serve customers in their local area. Building on the results on the heterogeneity of the spending shock, we use differences in average incomes and rents across ZIP codes as a source of variation in the spending shock that businesses face. This geographic analysis is useful both from the perspective of understanding mechanisms and because prior work shows that geography plays a central role in the impacts of economic shocks due to low rates of migration that can lead to hysteresis in local labor markets (Austin, Glaeser, and Summers 2018, Yagan 2019).

Small business revenues in the most affluent ZIP codes in large cities fell by more than 70% between March and late April, as compared with 30% in the least affluent ZIP codes. These reductions in revenue resulted in a much higher rate of small business closure in high-rent, high-income areas within a given county than in less affluent areas. This is particularly the case for non-tradable goods that require physical interaction – e.g., restaurants and accommodation services – where revenues fell by more than 80% in the most affluent neighborhoods in the country, such as the Upper East Side of Manhattan or Palo Alto, California. Small businesses that provide fewer in-person services – such as financial or professional services firms – experience much smaller losses in revenue even in affluent areas.

As businesses lost revenue, they passed the incidence of the shock on to their employees, particularly low-wage workers. Using data from payroll firms, we find that employment rates fell by 34% at the trough of the COVID recession (April 15, 2020) for workers with wages rates in the bottom quartile of the pre-COVID wage distribution. By contrast, employment rates fell by 10% for those in the top wage quartile. Employment for high-wage workers also rebounded much more quickly: employment levels for workers in the top wage quartile were back to pre-COVID levels by the end of May, but remained 17% below baseline for low-wage workers even as of September 2020. The greater persistence of job losses for low-wage workers is not explained purely by sectoral differences; even in sectors where spending rebounded to baseline levels, such as retail trade, employment of low-wage workers remained far below baseline levels, suggesting that firms may have shifted their modes of production to use less low-wage labor (Jaimovich and Siu 2020).

Low-wage individuals working at small businesses in affluent areas were especially likely to lose their jobs. In the highest-rent ZIP codes, more than 65% of workers at small businesses were laid off within two weeks after the COVID crisis began; by contrast, in the lowest-rent ZIP codes, fewer than 30% lost their jobs. Workers at larger firms and in tradable sectors (e.g., manufacturing) were much less likely to lose their jobs than those working in small businesses producing non-tradable

goods, irrespective of their geographic location. Job postings also fell much more sharply in more affluent areas, particularly for positions requiring less education. As a result of these changes in the labor market, unemployment claims surged even in affluent counties, which have generally had relatively low unemployment rates in prior recessions. For example, more than 15% of residents of Santa Clara county – the richest county in the United States, located in Silicon Valley – filed for unemployment benefits by May 2. Perhaps because they face higher rates of job loss and worse future employment prospects, low-income individuals working in more affluent areas cut their *own* spending much more than low-income individuals working in less affluent areas – showing that some workers were not fully insured against job loss despite the substantial expansion of safety net programs such as unemployment insurance.

In summary, the impacts of COVID-19 on economic activity in the first three months after the shock appear to be largely driven by a reduction in spending by higher-income individuals due to health concerns, which in turn affected businesses that cater to the rich – e.g., small businesses in affluent areas – and ultimately reduced the incomes and expenditure levels of low-wage employees of those businesses. In the third and final part of the paper, we analyze the impacts of three sets of policies that were enacted shortly after the crisis began in an effort to break this chain of events and mitigate economic losses: state-ordered shutdowns and reopenings, stimulus payments to households, and loans to small businesses.⁶

State-ordered shutdowns and reopenings of economies had modest impacts on economic activity. Spending and employment remained well below baseline levels even after reopenings, and trended similarly in states that reopened earlier relative to comparable states that reopened later. Spending and employment also fell well *before* state-level shutdowns were implemented, consistent with other recent work examining data on hours of work and movement patterns (Bartik et al. 2020, Villas-Boas et al. 2020). As a result, relatively little of the cross-state variation in spending and employment patterns is explained by the timing of shutdown orders or re-openings, consistent with the findings of Goolsbee and Syverson (2020) from cell phone data on locations.

Stimulus payments made to households in mid-April 2020 increased spending among low-income households sharply, nearly restoring their spending to pre-COVID levels by late April, consistent with evidence from Baker et al. (2020) and recent models of consumption that generate excess

6. Of course, this set of policies is by no means exhaustive: a vast set of other policy efforts ranging from changes in monetary policy to various state-level programs were also undertaken in response to the crisis. We focus on these three policies because they illustrate the ways in which the new high-frequency data we have assembled can be used for real-time policy analysis, and we hope that future work will use these data to analyze other policies.

sensitivity via frictions (e.g., Kaplan and Violante 2014). However, most of this increase in spending was in sectors that require limited physical interaction. Purchases of durable goods surged, while consumption of in-person services (e.g., restaurants) increased by much less. As a result, very little of the increased spending flowed to businesses most affected by the COVID-19 shock, such as small businesses in affluent areas – potentially limiting the capacity of the stimulus to increase economic activity and employment in the communities and sectors where job losses were largest because of diminished secondary multiplier effects (a broken Keynesian cross), as discussed in Guerrieri et al. (2020).

Loans to small businesses as part of the Paycheck Protection Program (PPP) also had small impacts on employment rates at small businesses. Employment rates at firms with fewer than 500 employees (who were eligible for PPP assistance) increased only slightly – by about 2 percentage points – relative to larger firms that were ineligible for PPP when the PPP program began. Our point estimates imply that the cost per job saved by the PPP was \$451,000 (\$114,000 at the lower bound of the 95% confidence interval). The PPP was not a cost-effective way to maintain employment likely because the vast majority of PPP loans went to inframarginal firms that were not planning to lay off many workers. These results are consistent with those of contemporaneous studies by Granja et al. (2020), who exploit cross-sectional variation in community bank shares coupled with the public data from our platform to show that the PPP had small effects on employment, and with those of Autor et al. (2020) and **hubbardhas**, who report similar point estimates using data from ADP and Dunn and Bradstreet. Together, these findings suggest that providing liquidity to firms is an expensive way to maintain employment rates in the short run, although it remains possible that the PPP may have long-term benefits by reducing business closures.

Our findings suggest that economic recovery from a pandemic ultimately requires restoring consumer confidence by addressing the root health concerns themselves (e.g., Allen et al. 2020, Romer 2020). Traditional macroeconomic tools – stimulating aggregate demand or providing liquidity to businesses – may have diminished short-run impacts when consumer spending is fundamentally constrained by health concerns. In such a setting, it may be more fruitful to provide social insurance to reduce hardship for those who have lost their jobs (e.g., via unemployment benefit extensions), consistent with the normative predictions of Guerrieri et al. (2020). In addition, the disparate impacts of the shock across areas suggests it may be useful to target employment assistance to places that have suffered the largest job losses (such as affluent, urban areas), since geographic disparities in unemployment persist for many years (Blanchard and Katz 1992, Yagan 2019).

Our work builds on two literatures: a longstanding literature on the measurement of economic activity and a nascent literature on the economics of pandemics. In the macroeconomic measurement literature, our work is most closely related to recent studies showing that private sector data sources can be used to forecast government statistics (e.g., Aladangady et al. 2019, Ehrlich et al. 2019, Cajner et al. 2019, Gindelsky, Moulton, and Wentland 2019, Dunn, Hood, and Driessen 2020). In the COVID-19 pandemic literature, several recent papers have used confidential private sector data similar to those we use to construct our platform to document related results on consumer spending (e.g., Baker et al. 2020, Chen, Qian, and Wen 2020, Cox et al. 2020), business revenues (e.g., Alexander and Karger 2020), labor market trends (e.g., Bartik et al. 2020, Cajner et al. 2020, Kurmann, Lalé, and Ta 2020, Kahn, Lange, and Wiczer 2020), and social distancing (e.g., Allcott et al. 2020, Chiou and Tucker 2020, Goldfarb and Tucker 2020, Mongey, Pilossoph, and Weinberg 2020).

Our paper makes two main contributions to these literatures. First, our analysis sheds light on the mechanisms through which pandemics affect economic activity. Other contemporaneous studies of the COVID-19 pandemic have focused on a specific subset of outcomes (e.g., spending or employment or job postings) at broad geographies. By combining data sources on multiple outcomes at the ZIP code level, we provide an integrated picture of how COVID-19 affected the macroeconomy – from changes in consumer spending to in-person business revenue losses to employment changes. In addition, analyzing a suite of outcomes allows us to characterize the impacts of major stabilization policies more fully, from changes in consumer behavior to impacts on businesses’ employment and hiring decisions. These findings also suggest new directions for future research. The sharp heterogeneity in impacts we document across ZIP codes provide a novel source of local variation to understand macroeconomic dynamics, similar to the geographic variation widely exploited to understand the Great Recession (Mian and Sufi 2009). More generally, our reduced-form results showing how correlated behavioral responses among consumers lead to downstream changes in product and labor markets in equilibrium provide estimates that can be used to calibrate macroeconomic models.

Second, and more broadly, this study opens new approaches to empirical macroeconomics by demonstrating that it is feasible to construct public statistics that are sufficiently disaggregated to answer many research and policy questions yet are sufficiently aggregated to protect privacy. Unlike the aforementioned studies of COVID-19 – each of which draws upon confidential data sources

– all the results reported here are produced from what are now *publicly available* data.⁷ This ability to analyze impacts in a timely, publicly verifiable manner creates new paths for evidence-based macroeconomic policy and research. For instance, one could potentially adjust policies as one observes their impacts on the economy (as was done repeatedly in the Paycheck Protection Program, but without the benefit of evidence on its ongoing impacts). Moreover, these impacts can be analyzed heterogeneously across areas, permitting tailored responses by local governments and analyses of disaggregated data by a much larger set of researchers. In this sense, the data assembled here provide a prototype for a system of real-time, granular national accounts that we hope will be refined in future work, much as Kuznets (1941) and Summers and Heston (1991) developed prototypes for national accounts within and across countries that were refined in subsequent work (e.g., Feenstra, Inklaar, and Timmer 2015).

The paper is organized as follows. The next section describes the data we use to construct the economic tracker. In Section 3, we analyze the effects of COVID-19 on spending, revenue, and employment. Section 4 analyzes the impacts of policies enacted to mitigate COVID’s impacts. Section 5 concludes. Technical details are available in an online appendix, and the data used to produce the results (along with replication code) can be downloaded from this [Github repository](#).

II Data and Methods

We use anonymized data from several private companies to construct indices of spending, employment, and other outcomes. To systematize our approach and facilitate comparisons between series, we adopt the following set of principles when constructing each series (wherever feasible given data availability constraints).

First, we remove artifacts in transaction data that arise from changes in data providers’ coverage or systems. For instance, firms’ clients often change discretely, sometimes leading to discontinuous jumps in series, particularly in small cells. We systematically search for large jumps in series (e.g., >80%), study their root causes by consulting with the data provider, and address such discontinuities by imposing continuity using series-specific methods described below.

7. In particular, the public platform eliminates the need for researchers or local policymakers to obtain permissions to use confidential data from each company, permitting analyses that make use of a much broader set of data. Indeed, the combination of these datasets is precisely what enables us to trace the macroeconomic impacts of the COVID shock from consumer spending to businesses to labor markets, unlike other studies. The cost of this approach is that one loses potential access to individual or business-level data, but our analysis demonstrates that one can answer many questions studied by others with the aggregate statistics that we construct.

Second, we address low- and high-frequency seasonal fluctuations in the data. We address high-frequency fluctuations through aggregation, e.g. by reporting 7-day moving averages to smooth fluctuations across days of the week. Certain series – most notably consumer spending and business revenue – exhibit strong lower-frequency fluctuations that are autocorrelated across years (e.g., a surge in spending around the holiday season). Where feasible, we de-seasonalize such series by normalizing each week’s value in 2020 relative to corresponding values for the same week in 2019, but we also report raw values for 2020 for researchers who prefer to make alternative seasonal adjustments.

Third, we take a series of steps to protect the confidentiality of businesses and their clients. Instead of reporting levels of each series, we report indexed values that show percentage changes relative to mean values in January 2020.⁸ We suppress small cells and exclude outliers to further protect privacy, with thresholds that vary across datasets as described below. For data obtained from publicly traded firms – whose ability to disclose data is restricted by Securities and Exchange Commission regulations governing the disclosure of material non-public information – we combine data from multiple firms so that the statistics we report do not reveal information about any single company’s activities.⁹

Finally, we address the challenge that our data sources capture information about the customers each company serves rather than the general population. Instead of attempting to adjust for this non-representative sampling, we characterize the portion of the economy that each series represents by comparing each sample we use to national benchmarks. We explicitly label the sector and population subgroup that each series represents and exclude data sources that do not track established benchmarks for that sector/subgroup closely. We examined several sources of spending, revenue, and employment data in addition to those used in the final analysis below and excluded them because they failed benchmarking tests.¹⁰

We release each data series at the highest available frequency using an automated pipeline that

8. We always norm after summing to a given cell (e.g. geographic unit, industry, etc.) rather than at the firm or individual level. This dollar-weighted approach overweights bigger firms and higher-income individuals, but leads to smoother series and is arguably more relevant for certain macroeconomic policy questions (e.g., changes in aggregate spending).

9. For publicly traded firms, a key contribution of our platform is that it serves as a coordination device that allows multiple firms to pool and release their data in an environment where each firm faces restrictions that limit its capacity to share its own data publicly.

10. An alternative approach is to reweight samples based on observable characteristics – e.g., industry – to match national benchmarks. We do not pursue such an approach here because the data sources we choose to work with largely track relevant national benchmarks – at least for the scale of shocks induced by the COVID crisis – without such reweighting. However, the disaggregated data we report by industry and county can be easily reweighted as desired in future applications.

ingests data from data providers, constructs the relevant statistics and conducts quality control tests, and outputs the series publicly (see Appendix A for details on the engineering of this pipeline). To limit revisions, we allow for a lag to account for reporting delays (typically one week). We disaggregate each series by two-digit NAICS industry code; by county, metro area, and state; and by income quartile where feasible.

In the rest of this section, we describe each of the series in turn, discussing the raw data sources, construction of key variables, and cross-sectional comparisons to publicly available benchmarks.¹¹ All of the data series described below can be freely downloaded from the Economic Tracker website: www.tracktherecovery.org.

II.A Consumer Spending: Affinity Solutions and CoinOut

We measure consumer spending primarily using aggregated and anonymized consumer purchase data collected by Affinity Solutions.

Affinity Solutions. [Affinity Solutions Inc](#) is a company that aggregates consumer credit and debit card spending information to support a variety of financial service products. Affinity Solutions captures nearly 10% of debit and credit card spending in the U.S. We obtain raw data from Affinity Solutions at the county-by-ZIP code income quartile-by-industry-by-day level starting from January 1, 2019. Cells with fewer than five unique card transactions are masked to protect privacy.

The raw Affinity data have discontinuous breaks caused by entry or exit of card providers from the sample. We identify these breaks using data on the total number of active cards in the cell. When we observe a discontinuity in the total number of cards on a given day, we impute the change in spending for that day using the change from a higher level of geography (see Appendix B for details). For instance, if there is a discontinuous jump in the number of cards at the county level on day t , we replace spending on day t in the relevant county with the spending in day $t - 1$ multiplied by one plus the growth rate in spending from $t - 1$ to t from the corresponding state-level series.

After cleaning the raw data in this manner, we construct daily values of the consumer spending series using a seven-day moving average of the current day and previous six days of spending. We then seasonally adjust the series by dividing each calendar date’s 2020 value by its corresponding value from 2019.¹² Finally, we index the seasonally-adjusted series relative to pre-COVID-19

11. We benchmark trends in each series over time to publicly-available data in the context of our analysis in the next section.

12. We divide the daily value for February 29, 2020 by the average value between the February 28, 2019 and March 1, 2019.

spending by dividing each value by the mean of the seasonally-adjusted average spending level in the first four complete weeks of 2020.

CoinOut. A concern with card-based spending measures is that they omit cash transactions, which account for 6.3% of consumer spending in the United States (Greene and Stavins 2020) and could potentially respond differently to the COVID shock and subsequent policies. To address this concern, we measure cash spending using transaction data from CoinOut, a company that allows individuals to receive rewards by uploading photos of their receipts to a mobile app. We obtain anonymized data from CoinOut starting from January 1, 2018 on the date and amount of each transaction; the user’s ZIP code; and the date and time the receipt was uploaded. We identify cash transactions by searching for the string “cash” in the text of each receipt and construct series on the total number and amount of cash transactions by day. The CoinOut data are not disaggregated by industry; however, since cash is almost always used in person, we view this series as representing spending on in-person goods (e.g., at grocery stores).

Comparison to QSS and MRTS. Total debit and credit card spending in the U.S. was \$7.08 trillion in 2018 (Board of Governors of the Federal Reserve System 2019), approximately 50% of total personal consumption expenditures recorded in national accounts. Appendix Figure 1 compares the spending distributions across sectors in our card spending series to spending captured in the nationally representative Quarterly Services Survey (QSS) and Monthly Retail Trade Survey (MRTS). Our card spending series has broad coverage across industries. However, as expected, it over-represents categories in which credit and debit cards are used for purchases. In particular, accommodation and food services and clothing are a greater share of the the card spending data than financial services and motor vehicles. We therefore view our series as providing statistics that are representative of total card spending (but not total consumer spending). We assess whether our Affinity Solutions series accurately captures changes in total card spending around the crisis in Section 3.1 below.

II.B Small Business Revenue: Womply

We obtain data on small business transactions and revenues from [Womply](#), a company that aggregates data from several credit card processors to provide analytical insights to small businesses and other clients. In contrast to the Affinity series on consumer spending, which is a cardholder-based panel covering total spending, Womply is a firm-based panel covering total revenues of small businesses. The key distinction is that location in Womply refers to the location where the business

transaction occurred as opposed to the location where the cardholder lives.

We obtain raw data on small business transactions and revenues from Womply at the ZIP-industry-day level starting from January 1, 2019.¹³ Small businesses are defined as businesses with annual revenue below [Small Business Administration thresholds](#). To reduce the influence of outliers, firms outside twice the interquartile range of firm annual revenue within this sample are excluded and the sample is further limited to firms with 30 or more transactions in a quarter and more than one transaction in 2 out of the 3 months (see Appendix C for details).

We aggregate these raw data to form two publicly available series at the county by industry level: one measuring total small business revenue and another measuring the number of small businesses open. We measure small business revenue as the sum of all credits (generally purchases) minus debits (generally returns). We define small businesses as being open if they have a transaction in the last three days. We exclude counties with a total average revenue of less than \$250,000 during the pre-COVID-19 period (January 4-31).

For each series, we construct daily values in exactly the same way that we constructed the consumer spending series. We first take a seven-day moving average, then seasonally adjust by dividing each calendar date’s 2020 value by its corresponding value from 2019. Finally, we index relative to pre-COVID-19 by dividing the series by its average value over January 4-31.

Comparison to QSS and MRTS. Appendix Figure 1 shows the distribution of revenues observed in Womply across industries in comparison to national benchmarks. Womply revenues are again broadly distributed across sectors, particularly those where card use is common. A larger share of the Womply revenue data come from industries that have a larger share of small businesses, such as food services, professional services, and other services, as one would expect given that the Womply data only cover small businesses.

II.C Employment and Earnings: Paychex, Intuit, Earnin, and Kronos

We combine several data sources to obtain information on employment and earnings: payroll data from Paychex and Intuit, worker-level data on employment and earnings from [Earnin](#), and time sheet data from Kronos. We describe each of these data sources in turn and then discuss how we combine them to construct a weekly series that measures private non-farm employment rates in the U.S.¹⁴ Further details are provided in Appendix D.

13. We crosswalk Womply’s transaction categories to two-digit NAICS codes using an internally generated Womply category-NAICS crosswalk, and then aggregate to NAICS supersectors.

14. The private payroll providers from whom we obtain data have limited coverage of government agencies; we therefore do not attempt to measure government employment here.

Paychex. Paychex provides payroll services to approximately 670,000 small and medium-sized businesses across the United States and pays nearly 10% of U.S. private-sector workers (Paychex 2020). We obtain aggregated weekly data on total employment, hours worked, and payroll earnings for each county by industry by 2019 hourly wage quartile by 2019 firm size by pay frequency. Industries are defined as two-digit NAICS codes. Hourly wage quartiles are based on fixed thresholds of the hourly wage distribution in 2019 ($< \$13.00$, $\$13.00$ - $\$18.18$, $\$18.18$ - $\$29.17$, $> \$29.17$). Salaried employees' wages are translated to hourly wages by dividing weekly pay by 40 hours. Firm size is measured as the average number of workers employed by the firm in 2019, broken into a set of broad groups (e.g., 1-49 employees, 50-99 employees, 100-199 employees, ..., 900-999 employees, > 999 employees). Firms that are new in 2020 are assigned a size of 0 (as they had no employees in 2019). Firm size at multi-establishment firms is calculated by summing establishment-level employment at each establishment within the same Dun & Bradstreet parent firm. Since we seek to measure private sector employment, we exclude workers employed in public administration and those with an unclassified industry (which each represent 0.8% of workers as of January 2020). We restrict the sample to workers with weekly, bi-weekly, semi-monthly or monthly pay frequencies; these workers represent over 99% of employees in the Paychex data.

We obtain raw data from Paychex on checks processed by week. We convert these data into an employment series using methods analogous to those developed by Cajner et al. (2019); see Appendix D for details. We first construct a daily series of pay checks processed using a linear interpolation between weekly values. Next, we distribute paychecks to the last date of the corresponding pay period. To do so, we use data provided by Paychex on the distribution of (date at which payroll is processed – last date in pay period) for weekly and bi-weekly pay frequencies, treating the distribution of (date at which payroll is processed – last date in pay period) as constant across geographies and NAICS codes.¹⁵ Finally, we record a worker as employed for the full duration of the pay cycle up until the last date in their pay period, under the assumption that workers are employed for each day during their pay period.

Intuit. Intuit offers payroll services to businesses as part of its Quickbooks program, covering approximately one million businesses as of January 2020. Businesses that use Quickbooks tend to be very small (fewer than 20 employees) and are concentrated in professional services industries, as

15. For monthly and semi-monthly pay frequencies, where cycles regularly occur on fixed calendar dates (e.g. the 15th and 30th of each month for semi-monthly paycycles) rather than on a weekly basis, we assume that the last date within each pay period is the closest preceding calendar date that is the 15th or the 30th day of the month (semi-monthly paycycles) or the 30th day of the month (monthly paycycles).

discussed in further detail below. Employment is defined in the Intuit data as the total number of workers who were paid a non-zero amount in the prior month. We obtain anonymized, aggregated data on month-on-month and year-on-year changes in total employment and average earnings at the state and county level by month, based on repeated cross-sections. To develop a national series, we take population-weighted averages of state changes in each month.

Earnin. Earnin is a financial management application that provides its members with access to their income as they earn it, in advance of their paychecks. Workers sign up for Earnin individually using a cell phone app, which connects to the bank account in which paychecks are deposited. Earnin complements the firm-based payroll datasets discussed above by providing a *worker-level* sample. This yields insight into employment rates at a much wider spectrum of firms – ranging from the largest firms in the U.S. to small businesses – at the expense of having fewer workers per firm. Since employment and hours decisions are highly correlated across workers within firms at business cycle frequencies, Earnin’s coverage of a large set of firms proves to be a valuable complement to the firm-based payroll datasets for our analysis. Because its users tend to have lower income levels, Earnin primarily provides information on hours for the bottom quartile of the wage distribution; we discuss the characteristics of the workers who use Earnin further below.

We obtain anonymized data from Earnin from January 2020 to present at the paycheck level with information on home ZIP code, workplace ZIP code, unemployment status, earnings, and hours worked. We assign firm sizes and NAICS codes to Earnin employers by matching to external data provided by ReferenceUSA using a custom-built crosswalk constructed by Digital Divide Data; for details, see Appendix D.¹⁶ Starting from this raw data, we measure employment as the total number of active Earnin users on a given day who are in paid employment, restricting to workers paid weekly or bi-weekly (see Appendix D for details). We distribute each individual’s paycheck over their pay cycle by assuming that individuals are employed for all days in their pay period. We assign workers to locations using their workplace ZIP codes to construct ZIP-code and county level measures.

Kronos. Kronos is a workforce management service used by many firms across the U.S. The data we obtain from Kronos cover approximately 30,000 mid-sized firms who together employed about 3.2 million workers pre-COVID. We obtain anonymized and aggregated weekly data on the total number of “punches,” where a punch represents an employee clocking into work on an

16. We match 68% of user-provided employer entries in the Earnin database to the ReferenceUSA data; we have missing information on firm size for the remaining firms.

electronic time sheet. We obtain these data by county, industry, and firm size at the point that the firm entered the Kronos database. The employees in the database are primarily hourly workers who must record their time, and are concentrated in the bottom quartile of the wage distribution: assuming full-time employment, their wage rates translate to average earnings of \$24,000 per year (with a 10th-90th percentile range of \$7,200 to \$45,600).

The Kronos data differ from the other three data sources above in that they measure data from time sheets rather than paychecks. The advantage of time sheets is that they provide very timely information on employment, with a lag of just 2-3 days (whereas payroll data naturally are obtained only once payroll is processed for the prior period, which can result in a lag of up to four weeks). The disadvantage of time sheets is that they do not capture total wage employment (e.g., workers may remain on payroll despite clocking fewer hours) and, naturally, only provide information for the subset of workers who are required to record their time.

Homebase. Homebase provides scheduling tools for small businesses (on average, 8.4 employees) such as restaurants (64% of employees for whom sectoral data are available) and retail stores (15% of employees for whom sectoral data are available). We obtain de-identified individual-level data on hours and total pay for employees and construct employment series at the county and industry level, assigning location based on the ZIP code of establishment. We include Homebase as a point of comparison because it has been widely used in other studies of small business employment in the COVID pandemic, but we do not include it in our primary employment indices because it does not track publicly available benchmarks of overall employment at small businesses as closely as our other data sources (see Section III.C below).

Combined Employment Series. We combine the data sources above to construct our primary employment series. Because Paychex covers all sectors and wage levels fairly comprehensively, we use Paychex data at the industry x wage quartile level for each geography (county, state and national) as the base for the combined employment series. We then use Earnin, Intuit and Kronos to refine the series in cells represented by those datasets. Earnin best represents workers in the bottom wage quartile (see Appendix Table 2 below). We therefore combine Earnin data with Paychex data to construct employment estimates for the bottom wage quartile. Next, we combine Intuit with the Paychex+Earnin data, accounting for the fact that Intuit data are available at a lower frequency and are not broken down by wage level or industry (see Appendix D). We report seven-day moving averages of these series, expressed as a percentage change relative to January 4-31.

The employment series constructed based on payroll data is generally available only with a one month lag because people are paid after completing work over multiple prior weeks. To obtain more current estimates, we use Kronos time sheet data along with Paychex data from firms with weekly paycycles to forecast employment rates (see Appendix D). We regress the combined employment series for each quartile of workers on the Kronos series for the same week, the corresponding quartile of the Paychex weekly series for the same week, as well as the three prior weeks of Paychex weekly data. We then use these regression coefficients combined with the most recent Kronos and Paychex weekly data to create a series of predicted employment rates for workers in each wage quartile.

Comparisons to QCEW and OES. Appendix Table 1 compares industry shares in each of the data sources above to nationally representative statistics from the Quarterly Census of Employment and Wages (QCEW). The Earnin and Paychex samples are broadly representative of the industry mix in the U.S., although high-skilled sectors (such as professional services) are slightly under-represented in Earnin as expected given that it consists primarily of lower-income workers. Intuit is concentrated primarily in professional services, construction, other services, and health care and social assistance. Homebase covers primarily food services.

Next, we assess how these datasets compare to national benchmarks in terms of wage rates by comparing the median wage rates of workers in Paychex, Intuit, and Earnin to nationally representative statistics from the BLS’s Occupational Employment Statistics (Appendix Table 2). Median wage rates in Paychex closely match the OES estimates. Average wages in Intuit closely mirror OES estimates in the industries that Intuit covers. Workers who use the Earnin app have median wages that are at roughly the 10th percentile of the wage distribution within each NAICS code. The one exception is the food and drink industry, where the median wages are close to the population median wages in that industry (reflecting that most workers in food services earn relatively low wages). Homebase exhibits a similar pattern, with lower wage rates compared to industry averages, except in sectors that have low wages, such as food services and retail.

We conclude based on these comparisons that our combined datasets provide a representative picture of private non-farm employment in the United States, and that Earnin provides good coverage of workers at the bottom of the wage distribution, who are a group of particular interest given their volatile employment rates over the business cycle.

II.D Job Postings: Burning Glass

We obtain data on job postings from 2007 to present from [Burning Glass Technologies](#). Burning Glass aggregates nearly all jobs posted online from approximately 40,000 online job boards in the United States. Burning Glass then removes duplicate postings across sites and assigns attributes including geographic locations, required job qualifications, and industry.

We obtain raw data on job postings at the industry-week-job qualification-county level from Burning Glass. Industry is defined using select [NAICS supersectors](#), aggregated from 2-digit NAICS classification codes assigned by a Burning Glass algorithm. Job qualifications are defined using ONET Job Zones. These [job zones](#) are mutually exclusive categories that classify jobs into five groups: needing little or no preparation, some preparation, medium preparation, considerable preparation, or extensive preparation. We also obtain analogous data broken by educational requirements (e.g., high school degree, college, etc.).

Comparison to JOLTS. Burning Glass data have been used extensively in prior research in economics; for instance, see Hershbein and Kahn (2018) and Deming and Kahn (2018). Carnevale, Jayasundera, and Repnikov (2014) compare the Burning Glass data to government statistics on job openings and characterize the sample in detail. In Appendix Figure 3, we compare the distribution of industries in the Burning Glass data to nationally representative statistics from the Bureau of Labor Statistics’ Job Openings and Labor Market Turnover [Survey](#) (JOLTS) in January 2020. In general, Burning Glass is well aligned across industries with JOLTS, with the one exception that it under-covers government jobs. We therefore view Burning Glass as a sample representative of private sector jobs in the U.S.

II.E Education: Zearn

[Zearn](#) partners with schools to provide a math program that combines in-person instruction with digital lessons. Zearn was used by approximately 800,000 students in the U.S. in Spring 2020. Many schools continued to use Zearn as part of their math curriculum after COVID-19 induced schools to shift to remote learning.

We obtain data on the number of students using Zearn Math and the number of lessons they completed at the school-grade-week level. The data we obtain are masked such that any county with fewer than two districts, fewer than three schools, or fewer than 50 students on average using Zearn Math during the pre-period is excluded. We fill in these masked county statistics with the commuting zone mean whenever possible. We winsorize values reflecting an increase of greater than

300% at the school level. We exclude schools who did not use Zearn Math for at least one week from January 6 to February 7 and schools that never have more than five students using Zearn Math during our analysis period. To reduce the effects of school breaks, we replace the value of any week for a given school that reflects a 50% decrease (increase) greater than the week before or after it with the mean value for the three relevant weeks.

We measure online math participation as the number of students using Zearn Math in a given week. We measure student progress in math using the number of lessons completed by students each week. We aggregate to the county, state, and national level, in each case weighting by the average number of students using the platform at each school during the base period of January 6-February 7, and we normalize relative to this base period to construct the indices we report.

Comparison to American Community Survey. In Appendix Table 3, we assess the representativeness of the Zearn data by comparing the demographic characteristics of the schools for which we Zearn data (based on the ZIP codes in which they are located) to the demographic characteristics of K-12 students in the U.S. as a whole. The distribution of income, education, and race and ethnicity of the schools in the Zearn sample is similar to that in the U.S. as a whole, suggesting that Zearn provides a representative picture of online learning for public school students in the U.S.

II.F Public Data Sources: UI Records, COVID-19 Incidence, and Google Mobility Reports

In addition to the new private sector data sources described above, we also collect and use three sets of data from public sources to supplement our analysis: data on unemployment benefit claims obtained from state government agencies; data on COVID-19 cases and deaths obtained from the New York Times; and data on the amount of time people spend at home vs. other locations obtained from Google’s COVID-19 Community Mobility Reports. Further details on these data sources are provided in Appendix E.

III Economic Impacts of COVID-19

In this section, we analyze the economic impacts of COVID-19, both to shed light on the COVID crisis itself and to demonstrate the utility of private sector data sources assembled above as a complement to national accounts data in tracking economic activity.

To structure our analysis, we begin from national accounts data released by the Bureau of

Economic Analysis (2020). GDP fell by \$1.73 trillion (an annualized rate of 31.7%) from the first quarter of 2020 to the second quarter of 2020, shown by the first bar in Figure 1a. GDP fell primarily because of a reduction in personal consumption expenditures (consumer spending), which fell by \$1.35 trillion. Government purchases and net exports did not change significantly, while private investment fell by \$0.47 trillion.¹⁷ We therefore begin our analysis by studying the determinants of this sharp reduction in consumer spending. We then turn to examine downstream impacts of the reduction in consumer spending on business activity and the labor market.

III.A Consumer Spending

We analyze consumer spending using data on aggregate credit and debit card spending. National accounts data show that spending that is well captured on credit and debit cards – essentially all spending excluding housing, healthcare, and motor vehicles – fell by approximately \$138 billion, comprising roughly 60% of the total reduction in personal consumption expenditures.¹⁸

Benchmarking. We begin by assessing whether the our card spending data track patterns in corresponding spending categories in the national accounts. Figure 1b plots spending on retail services (excluding auto-related expenses) and food services in the Affinity Solutions data vs. corresponding series from the Monthly Retail Trade Survey (MRTS), one of the main inputs used to construct the national accounts.¹⁹ All series are indexed to have a value of 1 in January of each calendar year; each point shows the average level of daily spending in a given month divided by spending in January of that year. The credit/debit card spending series tracks the MRTS closely both before and after the COVID crisis. In particular, both series show a rapid drop in food services spending in March and April 2020, while total retail spending fluctuates much less.

Figure 1c plots the change in spending from January to April 2020 in the Affinity Solutions series against the decline in consumer spending as measured in the MRTS. Despite the fact that

17. Most of the reduction in private investment was driven by a reduction in inventories and equipment investment in the transportation and retail sectors, both of which are plausibly a response to reductions in current and anticipated consumer spending. In the first quarter of 2020, consumer spending accounted for an even larger share of the reduction in GDP, further supporting the view that the initial shock to the economy came from a reduction in consumer spending.

18. The rest of the reduction is largely accounted for by healthcare and motor vehicle expenditures; housing expenditures did not change significantly. We view the incorporation of data sources to study these other major components of spending as an important direction for future work; however, we believe that the mechanisms discussed below may apply at least qualitatively to those sectors as well.

19. The series are not perfectly comparable because the category definitions differ slightly across the datasets. For example, we observe food and accommodation services combined together in the card data but only food services in the MRTS. In addition, the MRTS includes corporate card transactions, whereas we exclude them in order to isolate consumer spending. Hence, we would not expect the series to track each other perfectly even if the card spending data provided a perfect representation of national spending patterns.

the MRTS category definitions are not perfectly aligned with those in the card spending data, the relative declines are generally well aligned across sectors, with a correlation of 0.92. Given that credit card spending data tracks the MRTS at the national level quite well, we proceed to use it to disaggregate the national series in several ways to understand why consumer spending fell so sharply.²⁰

Heterogeneity by Income. We begin by examining spending changes by household income. We do not directly observe cardholders’ incomes in our data; instead, we proxy for cardholders’ incomes using the median household income in the ZIP code in which they live (based on data from the 2014-18 American Community Survey). ZIP-codes are strong predictors of income because of the degree of segregation in most American cities; however, they are not a perfect proxy for income and can be prone to bias in certain applications, particularly when studying tail outcomes (Chetty et al. 2020). To evaluate the accuracy of our ZIP code imputation procedure, we compare our estimates to those of Cox et al. (2020), who observe cardholder income directly based on checking account data for clients of JPMorgan Chase. Our estimates are closely aligned with those estimates, suggesting that the ZIP code proxy is reasonably accurate in this application.²¹

Figure 2a plots a seven-day moving average of total daily card spending for households in the bottom vs. top quartile of ZIP codes based on median household income.²² The solid line shows data from January to August 2020, while the dashed line shows data for the same days in 2019 as a reference. Spending fell sharply on March 15, when the National Emergency was declared and the threat of COVID became widely discussed in the United States. Spending fell from \$7.9 billion per day in February to \$5.5 billion per day by the end of March (a 30% reduction) for high-income households; the corresponding change for low-income households was \$3.5 billion to \$2.7 billion (a 23% reduction).

Because high-income households both cut spending more in percentage terms and accounted

20. Of course, our national benchmarking exercise does not guarantee that our statistics capture economic activity in every subgroup accurately. We cannot benchmark most datasets at the local level: this is precisely the value of the private sector data that we introduce here. To assuage concerns about differential selection bias across regions, we show that each of our main results are obtained in multiple different data sources, likely because any biases due to non-representative sampling are small relative to the scale of changes induced by COVID-19.

21. Cox et al. (2020) report an eight percentage point (pp) larger decline in spending for the highest income quartile relative to the lowest income quartile in the second week of April. Our estimate of the gap is also eight pp at that point, although the levels of the declines in our data are slightly smaller in magnitude for both groups. The JPMorgan Chase data cannot themselves be used for the analysis that follows because there are no publicly available aggregated series based on those data at present.

22. We estimate total card spending by multiplying the raw totals in the Affinity Solutions data by the ratio of total spending on the categories shown in the last bar of Figure 1a in PCE to total spending in the Affinity data in January 2020.

for a larger share of aggregate spending to begin with, they account for a much larger share of the decline in total spending in the U.S. than low-income households. In Column 1 of Table 1a, we estimate that as of mid-April, top-quartile households accounted for 40% of the aggregate spending decline after the COVID shock, while bottom-quartile households accounted for only 13% of the decline. This gap grew even larger after stimulus payments began in mid-April. By the end of July, top-quartile households accounted for roughly half of the total spending decline in the U.S. (Table 1, Column 3) and were still spending 16% less than their January levels, whereas bottom-quartile households were spending almost the same amount they were in 2019. This heterogeneity in spending changes by income is much larger than that observed in previous recessions (Petev, Pistaferri, and Eksten 2011, Figure 6) and plays a central role in understanding the downstream impacts of COVID on businesses and the labor market, as we show below.

A potential concern with our card-based estimates of spending changes is bias from substitution out of cash purchases; for instance, if individuals sought to use more contactless methods to pay or began placing more orders online, trends in card spending might exhibit excess volatility relative to overall spending. To assess the importance of such substitution, we examine cash purchases using receipts data from CoinOut. Appendix Figure 2a plots aggregate cash purchases in the CoinOut data vs. aggregate card spending at grocery stores over time.²³ The time trends are very similar between the two series (with a signal correlation of 0.9 at the weekly level), showing a sharp spike in spending in late March (as households stocked up on groceries), followed by a more sustained increase in spending from the latter half of April. Although we cannot rule out small differential shifts into or out of cash, aggregate fluctuations in card spending do not appear to have been offset by opposite-signed changes in cash spending. Rather, households shifted spending similarly across both modes of payment.²⁴ We therefore proceed to focus on card spending as a proxy for total spending in the rest of our analysis.

Heterogeneity Across Sectors. Next, we disaggregate the change in total card spending across categories to understand why households cut spending so rapidly. In particular, we seek to distinguish two channels: reductions in spending due to loss of income vs. fears of contracting or spreading COVID.

The left bar in Figure 2b plots the share of the total decline in spending from the pre-COVID

23. We focus on grocery spending in the card data because cash spending in CoinOut is concentrated in certain sectors such as groceries; unfortunately, we are unable to disaggregate the CoinOut data by sector or align sectoral definitions more precisely across the datasets.

24. Appendix Figures 21b and 21c compare the patterns for spending in high- vs. low-rent areas; the patterns also appear quite similar between the CoinOut spending and grocery card spending across areas as well.

period to mid-April accounted for by various categories. Nearly three-fourths of the reduction in spending comes from reduced spending on goods or services that require in-person contact (and thereby carry a risk of COVID infection), such as hotels, transportation, and food services.²⁵ This is particularly striking given that these goods accounted for only one-third of total spending in January, as shown by the right bar in Figure 2b. Panel B of Table 1 shows that these gaps have only grown larger as the pandemic has progressed, as consumer spending increased above pre-pandemic levels for durable and non-durable goods by mid-July, but remained sharply depressed for in-person services.

Next, we zoom in to specific subcategories of spending that differ sharply in the degree to which they require physical interaction in Figure 2c. Spending on luxury goods such as installation of home pools and landscaping services – which do not require in-person contact – *increased* slightly after the COVID shock; by contrast, spending on restaurants, beauty shops, and airlines all plummeted sharply. Consistent with these substitution patterns, spending at online retailers increase sharply: online purchases comprised 11% of retail sales in 2019 vs. 22% in April and May of 2020 (Mastercard 2020).²⁶ A conventional reduction in income or wealth would typically reduce spending on all goods as predicted by their Engel curves (income elasticities); the fact that the spending reductions vary so sharply across goods that differ in terms of their health risks lends further support to the hypothesis that it is health concerns rather than a lack of purchasing power that drove spending reductions.

These patterns of spending reductions are particularly remarkable when contrasted with those observed in prior recessions. Figure 2d compares the change in spending across categories in national accounts data in the COVID recession and the Great Recession in 2009-10. In the Great Recession, nearly all of the reduction in consumer spending came from a reduction in spending on goods; spending on services was almost unchanged. In the COVID recession, 67% of the reduction in total spending came from a reduction in spending on services, as anticipated by Mathy (2020).

Heterogeneity by COVID Incidence. To further evaluate the role of health concerns, we next turn to directly examine the association between incidence of COVID across areas and changes in spending. Figure 3a presents a binned scatterplot of changes in spending from January to April vs. the rate of detected COVID cases by county. To construct this figure, we divide the x variable (COVID cases) into 20 bins, each of which contain 5% of the population, and plot the mean value of the x and y variables within each bin. Areas with higher rates of COVID infection experience

25. The relative shares of spending reductions across categories are similar for low- and high-income households (Appendix Figure 4); what differs is the level of spending reduction, as discussed above.

26. We are unable to distinguish online and in-store transactions in the Affinity Solutions data.

significantly larger declines in spending, a relationship that holds conditional on controls for median household income and state fixed effects (Appendix Figure 5).²⁷

To examine the mechanism driving these spending reductions more directly, in Figure 3b, we present a binned scatterplot of the amount of time spent outside home (using anonymized cell phone data from Google) vs. COVID case rates, separately for low- and high-income counties (median household income in the bottom vs. top income quartile). In both sets of areas, there is a strong negative relationship: people spend considerably less time outside home in areas with higher rates of COVID infection. The reduction in spending on services that require physical, in-person interaction (e.g., restaurants) is mechanically related to this simple but important change in behavior.

At all levels of COVID infection, higher-income households spend less time outside. Figure 3c establishes this point more directly by showing that time spent outside home falls monotonically with household income across the distribution. These results help explain why the rich reduce spending more, especially on goods that require in-person interaction: high-income people apparently self-isolate more, perhaps by working remotely or because they have larger living spaces.

In sum, disaggregated data on consumer spending reveals that spending in the initial stages of the pandemic fell primarily because of health concerns rather than a loss of current or expected income. Indeed, income losses were relatively modest because relatively few high-income individuals lost their jobs – as we show in Section III.C below – and lower-income households who experienced job loss had their incomes more than replaced by unemployment benefits (Ganong, Noel, and Vavra 2020). As a result, national accounts data actually show an *increase* in total income of 13% from March to April 2020. This result implies that the central channel emphasized in Keynesian models that have guided policy responses to prior recessions – a fall in aggregate demand due to a lack of purchasing power – was less important in the early stages of the pandemic, partly as a result of policies such as increases in unemployment benefits that offset lost earnings. Rather, the key driver of residual changes in aggregate spending is a contraction in firms’ ability to supply certain goods, namely services that carry no health risks.²⁸ We now show that this novel source of spending

27. Note that there is a substantial reduction in spending even in areas without high rates of realized COVID infection, which is consistent with widespread concern about the disease even in areas where outbreaks did not actually occur at high rates.

28. Of course, these results only apply to the period we study, contingent on the policies that were in place: the first three months after the COVID shock hit the U.S., when the federal government was injecting substantial income into the economy via increased UI benefits and stimulus checks to households. Recessions induced by pandemics could produce more traditional economic shocks with Keynesian spillovers across a wider set of sectors and areas as time passes or in the absence of income support for the unemployed, in which case tools such as stimulus and liquidity could become much more impactful (Guerrieri et al. 2020).

reductions leads to a distinct pattern of downstream impacts on businesses and the labor market, potentially calling for different policy responses than in prior recessions.

III.B Business Revenues

We now turn to examine how reductions in consumer spending affect business activity. Conceptually, we seek to understand how a change in revenue for a given firm affects its decisions: whether to remain open, how many employees to retain, what wage rates to pay them, how many new people to hire. Ideally, one would analyze these impacts at the firm level, examining how the customer base of a given firm affected its revenues and employment decisions. Lacking firm-level data, we use geographic variation as an instrument for the shocks to firm’s revenues. The motivation for this geographical approach is that spending fell primarily among high-income households in sectors that require in-person interaction, such as restaurants. Most of these goods are non-tradable products produced by small local businesses who serve customers in their local area.²⁹ We therefore use differences in average incomes and rents across ZIP codes as a source of variation in the magnitude of the spending shock that small businesses face.³⁰

Benchmarking. We measure small business revenues using data from Womply, which records revenues from credit card transactions for small businesses (as defined by the Small Business Administration). Business revenues in Womply closely track patterns in the Affinity total spending data, especially in sectors with a large share of small businesses, such as food and accommodation services (Appendix Figure 6).³¹

Heterogeneity Across Areas. We begin our analysis of the Womply data by examining how small business revenues changed in low- vs. high-income ZIP codes (formally, ZIP code Tabulation Areas - ZCTAs) from a baseline period prior the COVID shock (January 8 to March 8, 2020) to the weeks immediately after the COVID shock before the stimulus program began (March 9 to May 3, 2020). Figure 4 maps the change in small business revenue by ZIP code in three large metro areas: New York City, San Francisco, and Chicago (analogous ZIP-level maps for other cities are

29. 56% of workers in food and accommodation services and retail (two major non-tradeable sectors) work in establishments with fewer than 50 employees.

30. We focus on small businesses because their customers are typically located near the business itself; larger businesses’ customers (e.g., large retail chains) are more dispersed, making the geographic location of the business less relevant. One could also in principle use other groups (e.g., sectors) instead of geography as instruments. We focus primarily on geographic variation because the granularity of the data by ZIP code yields much sharper variation than what is available across sectors and arguably yields comparisons across more similar firms (e.g., restaurants in different neighborhoods rather than airlines vs. manufacturing).

31. In sectors that have a bigger share of large businesses – such as retail – the Womply small business series exhibits a larger decline during the COVID crisis than Affinity (or MRTS). This pattern is precisely as expected given other evidence that consumers shifted spending toward large online retailers such as Amazon (Alexander and Karger 2020).

available [here](#)). There is substantial heterogeneity in revenue declines across areas. For example, average revenue declines range from -67% (or below) in the hardest-hit (lowest decile) of ZIP codes to -15% (or above) in the top decile within New York City.³²

In all three cities, revenue losses are largest in the most affluent parts of the city. For example, small business lost 63% of their revenue in the Upper East Side in New York, compared with 39% in the East Bronx; 67% in Lincoln Park vs. 38% in Bronzeville on the South Side of Chicago; and 71% in Nob Hill vs. 33% in Bayview in San Francisco. Revenue losses are also large in the central business districts in each city (lower Manhattan, the Loop in Chicago, the Financial District in San Francisco), likely a direct consequence of the fact that many workers who used to work in these areas are now working remotely. But even within predominantly residential areas, businesses located in more affluent neighborhoods suffered much larger revenue losses, consistent with the heterogeneity in spending reductions observed in the Affinity data.³³ More broadly, cities that have experienced the largest declines in small business revenue on average tend to be affluent cities – such as New York, San Francisco, and Boston (Appendix Table 4, Appendix Figure 8).

Figure 5a generalizes these examples by presenting a binned scatter plot of percent changes in small business revenue vs. median household incomes, by ZIP code across the entire country. We observe much larger reductions in revenue at local small businesses in affluent ZIP codes. In the richest 5% of ZIP codes, small business revenues fell by 60%, as compared with 40% in the poorest 5% of ZIP codes.³⁴

As discussed above, spending fell most sharply not just in high-income areas, but particularly in high-income areas with a high rate of COVID infection. Data on COVID case rates are not available at the ZIP code level; however, one well established predictor of the rate of spread of COVID is population density: the infection spreads more rapidly in dense areas. Figure 5b shows that small business revenues fell more heavily in more densely populated ZIP codes.³⁵

32. Very little of this variation is due to sampling error: the reliability of these estimates across ZIP codes within counties exceeds 0.8, i.e., more than 80% of the variance within each of these maps is due to signal rather than noise.

33. We find a similar pattern when controlling for differences in industry mix across areas; for instance, the maps look very similar when we focus solely on small businesses in food and accommodation services (Appendix Figure 7).

34. Of course, households do not restrict their spending solely to businesses in their own ZIP code. An alternative way to establish this result at a broader geography is to relate small business revenue changes to the degree of income inequality across counties. Counties with higher Gini coefficients experienced large losses of small business revenue (Appendix Figure 9a). This is particularly the case among counties with a large top 1% income share (Appendix Figure 9b). Poverty rates are not strongly associated with revenue losses at the county level (Appendix Figure 9c), showing that it is the presence of the rich in particular (as opposed to the middle class) that is most predictive of economic impacts on local businesses.

35. Consistent with this pattern, total spending levels and time spent outside also fell much more in high population density areas.

Figure 5c combines the income and population density mechanisms by plotting revenue changes vs. median rents (for a two bedroom apartment) by ZIP code. Rents are a simple measure of the affluence of an area that combine income and population density: the highest rent ZIP codes tend to be high-income, dense areas such as Manhattan. Figure 5c shows a particularly steep gradient of revenue changes with respect to rents: revenues fell by less than 30% in the lowest-rent ZIP codes, compared with more than 60% in the highest-rent ZIP codes. This relationship is essentially unchanged when controlling for worker density in the ZIP code and county fixed effects (Table 2).

In Figure 5d, we examine heterogeneity in this relationship across sectors that require different levels of physical interaction: food and accommodation services and retail trade (which largely require in-person interaction) vs. finance and professional services (which largely can be conducted remotely). Revenues fall much more sharply for food and retail in higher-rent areas; in contrast, there is essentially no relationship between rents and revenue changes for finance and professional services. These findings show that businesses that cater *in person* to the rich are those that lost the most businesses. Naturally, many of those businesses are located in high-income areas given people’s preference for geographic proximity in consuming services.

As a result of this sharp loss in revenues, small businesses in high-rent areas are much more likely to close entirely. We measure closure in the Womply data as reporting zero credit card revenue for three days in a row. Appendix Figure 10 shows that 55% of small businesses in the highest-rent ZIP codes closed, compared with 40% in the lowest rent ZIP codes. The extensive margin of business closure accounts for most of the decline in total revenues.

Because businesses located in high-rent areas lose more revenue in percentage terms and tend to account for a greater share of total revenue to begin with, they account for a very large share of the total loss in small business revenue. More than half of the total loss in small business revenues comes from business located in the top-quartile of ZIP codes by rent; only 8% of the revenue loss comes from businesses located in the bottom quartile. We now examine how the incidence of this shock is passed on to their employees.

III.C Employment Rates

We study the impacts of COVID on employment rates using data from payroll companies. We begin by benchmarking these data sources to employment statistics from nationally representative surveys conducted by the Bureau of Labor Statistics and then disaggregate the data by wage level and geography to analyze how the shock in consumer spending and business revenue affected

employment rates.

Benchmarking. Figure 6a plots employment rates from the nationally representative Current Employment Statistics (a survey of businesses) and Current Population Survey (a survey of households) for all workers alongside our combined Paychex-Intuit-Earnin employment series, constructed as described in Section II.C. Our payroll-based series is broadly aligned with the survey-based measures, falling between estimates obtained from the two surveys.

Figure 6b examines how our series performs in matching national statistics on trends across sectors. For illustration, we focus on two sectors that experienced very different trajectories: food services, where employment fell heavily, and professional services, where it did not. In both cases, our Paychex-Intuit-Earnin series closely tracks data from the CES. Appendix Figure 12d shows more generally that changes in employment rates across sectors (two-digit NAICS) are very closely aligned in our series and the CES, with a correlation of 0.94

For comparison, we also examine trends in employment based on data from Homebase, a dataset that has been used to examine employment trends in the COVID recession in many studies. Homebase exhibits a much larger decline in employment than the other series (56.1% at the trough vs. 15.2% in the CES). This is primarily because 64% of individuals in the Homebase data work in the food services sector, which suffered particularly large employment losses as noted above; however, even within food services, Homebase exhibits a larger decline in employment at the trough (63.8%) relative to the CES (46.8%), as shown in Figure 6b. Because Homebase does not track overall national benchmarks on employment very closely, we do not use it for the analysis that follows, although we note that it exhibits qualitative patterns similar to the other series within food services.

In Appendix Figure 12, we compare trends by wage quartile in our data with estimates based on the Current Population Survey and estimates reported in Cajner et al. (2020), who report employment changes by wage quintile using data from ADP in the initial weeks after the COVID shock. We find broadly similar trends in all three datasets. We also examine employment changes by state and find that in almost all states (excluding Kansas, North Dakota, Hawaii, and Idaho), employment changes from January-April in our combined series align very closely with changes in the CES, with an overall correlation of 0.98.

Based on these benchmarking exercises, we conclude that our combined employment series provides a good representation of employment rates across sectors, wage groups, and geographic areas. Consistent with the results of Bartik et al. (2020) and Cajner et al. (2020), we find that wage rates remained relatively constant, at least in comparison to the sharp fluctuations in employment,

after the COVID shock for workers who retained their jobs (Appendix Figure 11). Additionally, changes in employment rates are virtually identical to changes in hours because the extensive margin accounts for the vast majority of hours reductions. As a result, the employment changes in Figure 6 are almost identical to observed changes in workers’ hours and earnings. We therefore focuses solely on employment changes in what follows to characterize the incidence of the COVID shock on workers.

Heterogeneity by Wage Rates. Figure 7a plots the combined employment series by wage quartile. To construct this figure, we first construct hourly wage quartiles based on fixed thresholds of the hourly wage distribution in 2019 ($< \$13.00$, $\$13.00$ - $\$18.18$, $\$18.18$ - $\$29.17$, $> \$29.17$). The solid lines plot total employment (based on repeated daily cross-sections) in each of these bins relative to the January baseline, based on the combined Paychex-Intuit-Earnin data. Consistent with the findings of Cajner et al. (2020) using ADP data, we find sharp heterogeneity in job losses by wage rate. Employment rates fell by 34% at the trough of the recession (April 15) for workers in the bottom wage quartile (i.e., the total number of jobs paying $< \$13$ /hour was 34% lower as of April 15 than in January). By contrast, employment rates fell by only 10% for those in the top wage quartile as of April 15.

High-wage workers not only were less likely to lose their jobs to begin with, but also experienced a much more rapid recovery. By late May, employment for high-wage workers had returned nearly to the pre-COVID baseline. But employment rates for low-wage workers remained 15% below baseline levels even as of late July. Using time sheet data from Kronos, and payroll data from firms with weekly paycycles in Paychex – both of which are available with a shorter lag than payroll-based employment data containing all paycycles – we construct a prediction of employment rates for low-wage workers up to September 25 as described in Section II.C (shown by the dashed lines in Figure 7a). These predictions suggest that the rate of recovery remained slow in September, and as a result employment levels for low-wage workers were likely to remain well below baseline levels even at the end of September.

In sum, COVID induced a short-term “V-shaped” recession for high-wage workers in terms of employment opportunities, but led to a much deeper and more prolonged recession for lower-wage workers. Why did employment trajectories for low-wage workers differ so sharply from those for high-wage workers? One potential explanation is that low-wage workers work in different sectors or areas that may have experienced larger reductions in consumer demand. We evaluate this hypothesis in Figure 7b by plotting employment for workers in the bottom wage quartile, reweighting

the series to match baseline employment shares by county and industry (2 digit NAICS) in the top wage quartile. This reweighting closes very little of the gap between the two series, showing that differences in industry and location do not explain the differences in employment trajectories.

Figure 7c provides a specific illustration of this result by showing trends in employment and spending in the retail trade sector. Total retail spending is nearly 10% higher as of August 15 relative to the pre-COVID baseline. Employment of high-wage workers is comparable to baseline levels, yet employment of low-wage workers is still down by slightly over 20%. More broadly, even though consumer spending was only 5.9% below baseline levels as of August 15, low-wage employment levels remained 19.2% lower.

One explanation for these patterns is that firms shifted their production processes to use more technology, reducing demand for routine occupations – a common phenomenon in earlier recessions, as documented by Jaimovich and Siu (2020). For example, retail spending may have shifted toward online retailers and larger firms that may use more capital (or imports) than low-wage labor in the United States to produce goods. Given prior evidence of significant hysteresis in labor markets (e.g., Yagan 2019), these results raise the possibility of another protracted “jobless recovery” absent efforts to help workers who have been displaced from their prior jobs.

Heterogeneity Across Areas. To shed further light on why employment rates for low-wage workers fell so much, we next turn to examine geographic heterogeneity in employment losses, in connection to the heterogeneity in spending changes and business revenue losses examined above. We begin by using the Earnin data – which is publicly available at the ZIP code level – to analyze heterogeneity across ZIP codes. Figure 8 maps changes in employment rates for low-wage workers at small- and mid-size businesses (fewer than 500 employees) by ZIP code in New York, San Francisco, and Chicago (analogous ZIP-level maps for other cities are available [here](#)).³⁶ The patterns closely mirror those observed for business revenues above. Employment rates for low-wage workers fell by more than 80% in the most affluent areas of these cities, as compared with 30% in the least affluent areas. We observe very similar spatial patterns when examining variation across commuting zones (aggregates of counties) at the national level using the combined Paychex-Intuit-Earnin data (Appendix Figure 13).

Figure 9a presents a binned scatter plot of changes in hours of work vs. median rents by employer ZIP code, by firm size. We see much larger reductions in hours of work for workers who

36. We focus on small and mid-size businesses here because larger firms exhibit significantly smaller declines in employment (Appendix Figure 14) and because, as noted above, their markets are likely to extend well beyond the ZIP code in which they are located.

work in high-rent areas than low-rent areas in all groups. Employment rates fell by more than 55% for workers in the smaller group of firms located in high-rent ZIP codes, as compared with 25% for workers in low-rent ZIP codes, supporting the view that the sharp reductions in business revenue in affluent areas induced firms to lay off low-wage workers.

Interestingly, we observe a similar gradient with respect to local rents for workers at very large firms: from 20% in the lowest-rent ZIPs to nearly 40% in the highest-rent ZIPs. This presumably reflects that fact that multi-establishment firms such as Starbucks face larger revenue losses at stores located in more affluent neighborhoods for the reasons documented above, which in turns induces them to reduce employment in those areas more heavily. While there is a similar gradient with respect to rent levels, the overall level of employment losses for workers at large firms is lower than at smaller firms. This may be because large firms lost less revenue as a result of the COVID shock given their line of business (e.g., fast food vs. sit-down restaurants) or have a greater ability to substitute to other modes of business (delivery, online retail).³⁷

Figure 9b replicates Figure 9a using our combined Paychex-Intuit-Earnin employment series – which is available only at the county level. We see a very similar pattern of larger losses in employment for low-wage (bottom quartile) workers in high-rent counties, although the magnitude of the gradient is attenuated as expected given the coarser geographic measure. Table 2b presents a set of regression specifications quantifying these impacts. Across a broad range of specifications, we see that low-wage workers consistently face larger employment losses in higher-rent areas.

Job Postings. Prior work suggests that the labor market impacts of the recession may depend as much upon job postings as they do on the rate of initial layoffs (e.g., Diamond and Blanchard 1989, Elsby, Michaels, and Ratner 2015). We therefore now turn to examine how the spending shocks and revenue losses have affected job postings. We measure job postings at the county level using data from Burning Glass, which prior work has shown is fairly well aligned with government statistics based on the Job Openings and Labor Turnover Survey (Carnevale, Jayasundera, and Repnikov 2014, Kahn, Lange, and Wiczer 2020).³⁸ We conduct this analysis at the county level, pooling firms of all sizes and sectors because workers can substitute across firms and areas when searching for a new job, making it less relevant which exact firm or ZIP code they work in.

Figure 9c presents a binned scatter plot of the change in job postings pre- vs. post-COVID vs.

37. We cannot measure changes in revenue by establishment for large firms because the Womply data on revenues only cover small businesses. Moreover, one would need data on revenues by *establishment* within large companies to conduct such an analysis.

38. Burning Glass measures the sum of job postings, whereas JOLTS measures job openings at a given point in time. Hence, jobs that are posted and quickly filled will be included in Burning Glass but not in JOLTS.

median rents by county for jobs that require minimal education. We find a pattern similar to what we find with current employment: job postings for lower-skilled workers in high-rent areas have fallen much more sharply (by approximately 30%) than for workers in lower-rent areas. Hence, low-wage workers in such areas are not only more likely to have lost their jobs to begin with, they also have poorer prospects of finding a new job. Figure 9d replicates Figure 9c for job postings that require higher levels of education. For this group, which is much more likely to be employed in tradable sectors that are less influenced by local conditions (e.g., finance or professional services), there is no relationship between local rents and the change in job postings, consistent with our findings above in Figure 5d.³⁹

Unemployment Rates. The low rates of job postings combined with high rates of job loss in affluent areas combined to create very tight labor markets that produce unemployment in such areas that are unprecedented in recent history. To illustrate this, we contrast rates of employment losses by county in the COVID recession (from Feb-April 2020) with the Great Recession (from 2007-2010) using statistics on employment from the Bureau of Labor Statistics.⁴⁰

Figure 10 shows that in the Great Recession, counties with lower median incomes tended to account for a greater share of job losses. In particular, the first set of bars in Figure 10 show that counties in the bottom quartile (25%) of household median income distribution comprised a disproportionate (30%) share of job losses. In contrast, in the recent recession they account for actually less than 25% of the job losses, consistent with the evidence above that employment losses from the COVID shock have been concentrated among low-income employees in affluent areas. In the final set of bars, we show that in the recent recession this has led to the surprising pattern that UI claims are almost equally likely to come from high versus low-income counties.⁴¹

Santa Clara county in California is the highest income county on the West Coast, yet 16% of its labor force claimed UI between March 15th to May 2nd. This claim rate is identical to the share of the labor force that claimed UI in Fresno CA, a low-income county in California’s Central Valley.

39. The magnitude of the reduction in job postings for highly educated workers is substantial, at approximately 27%. This contrasts with evidence that higher-skilled workers have experienced much lower rates of job loss to date, and suggests that unemployment rates could begin to rise even for higher-skilled workers going forward.

40. One notable feature of the COVID recession is that the increase in unemployment rates between February and April 2020 (11%) is only two-thirds as large as the decrease in employment (16%). The difference is due to a 5% decline in the labor force: many people lost their jobs but were not actively searching for a new job in the midst of the pandemic. In the three prior recessions, the labor force continued to grow by 0.3% to 0.8% annually. We therefore focus on the decline in employment rates to obtain comparable statistics on job loss across recessions.

41. Unlike our analyses of private data, the publicly released unemployment claims data do not allow us disaggregate changes in employment by individuals’ income or ZIP code. Given the evidence above that job losses are concentrated among low-wage workers in high-income areas, there is strong reason to believe that the unemployment claims in high-income counties are coming from lower-income individuals living in those counties.

Unemployment rates above 10% have happened regularly in Fresno during prior recessions, but are unprecedented in Santa Clara. In Montgomery County, MD, long one of the richest counties in the U.S., workers have historically been quite insulated from prior recessions. During the 1991 and 2001 recessions the unemployment rate in Montgomery remained 3%. In 2010 it only hit 6%, one of the lowest in the country. In May 2020 employment losses and unemployment claims in Montgomery exceeded 12% of the labor force, resembling many counties with much lower average incomes.

In the Great Recession, the areas of the country that experienced the largest increases in unemployment took many years to recover because workers did not move to find new jobs and job vacancies remained depressed in hard-hit areas well after the national recession ended (Yagan 2019). Appendix Figure 15 shows early signs of a similar pattern in this recession: job postings went up significantly starting in late May in the U.S., but remained significantly lower in high-rent counties than in low-rent counties (where postings recovered nearly to pre-COVID levels by the end of May). If this pattern persists going forward, the recovery for low-income workers may take the longest in the richest parts of the country.

III.D Spending by Low-Income Workers

We close our analysis by showing job loss induced by working for firms in affluent areas affected the consumption of low-income workers themselves. To do so, we return to the credit card spending data from Affinity Solutions and ask whether low-income individuals working in high-rent ZIP codes reduce spending more than those working in low-rent ZIP codes.

Because we cannot measure workplace location in the credit card data itself, we use data from the Census LEHD Origin-Destination Employment Statistics (LODES) database, which provides information on the matrix of residential ZIP by work ZIP for all workers in the U.S. in 2017. Using this matrix, we compute the average workplace median rent level for each residential ZIP. Figure 11a presents a binned scatter plot of changes in employment by *home* (residential) ZIP code and average workplace rent, restricting the sample to low-income (bottom income quartile) ZIP codes. This figure confirms that low-income individuals who work in high-rent areas are more likely to lose their jobs, verifying that the LODES data linked to residential ZIPs produce the same result as directly using workplace ZIP codes in the Earnin data.

Figure 11b replicates Figure 11a using spending changes from January 5-March 7 to April 8-28 on the y axis. Low-income individuals who work in high-rent ZIP codes cut spending by 35%

on average from the baseline period to mid-April 2020, compared with 15% for those working in low-rent ZIPs. In Appendix Table 5, we present a set of regression specifications showing that the relationship remains similar when we compare ZIP codes within the same county by including county fixed effects, control for rents in the home (residential) ZIP code, and include other controls. Intuitively, these results show that among two equally low-income ZIP codes in Queens, those who live in a ZIP code where many work in an affluent area (perhaps because of a proximate subway line into Manhattan) are more likely to lose their jobs and, as a result, cut their own spending more following the COVID shock.

These findings imply that low-income households were not fully insured against job loss, consistent with other data showing that food insecurity rose in the COVID pandemic (Bitler, Hoynes, and Schanzenbach 2020). This may be surprising given that unemployment benefits were increased substantially and most households also received stimulus payments in mid-April. One explanation for why those who lost their jobs reduced spending is that they were ineligible for or faced delays in signing up for unemployment insurance benefits (Farrell et al. 2020). Another possibility is that households recognized that increased government support would be temporary and began engaging in precautionary saving. In either case, it appears that the initial reduction in spending due to health concerns by high-income households ultimately led to less spending by low-income households who lost their jobs and had lower current or expected incomes.⁴²

IV Evaluation of Policy Responses to COVID-19

We have seen that a chain of events led to substantial employment losses following the COVID-19 shock: (1) reductions in spending by high-income individuals due to health concerns, (2) revenue losses for businesses catering to those customers, and (3) job losses for low-income workers working at those businesses. We now turn to study what type of policies can mitigate the economic impacts of the pandemic, focusing in particular on increasing employment among low-income workers. We study three sets of policies that target different points of the economic chain: (1) state-ordered business reopenings that remove barriers to economic activity; (2) stimulus payments to households, which aim to spur consumer spending and thereby increase employment; and (3) loans to small businesses, which provide liquidity to keep workers on payroll.

42. Of course, low-income households would have had to cut spending much more had the government not provided income support (Casado et al. 2020, Farrell et al. 2020).

IV.A State-Ordered Reopenings

One direct approach to changing consumer spending and employment is via executive orders. Many states enacted stay-at-home orders and shutdowns of businesses in an effort to limit the spread of COVID infection and later reopened their economies by removing these restrictions. We begin by examining how such executive orders affect economic activity, exploiting variation across states in the timing of shutdowns and reopenings. Throughout this section, we define the reopening date to be the day that a state *began* the reopening process (see Appendix F for details). In most states, reopening was a gradual process in which certain industries and types of businesses opened before others, but there was a lot of heterogeneity across states in the precise form that the reopening took. Our estimates should therefore be viewed as an assessment of the average impact of typical re-opening efforts on aggregate economic activity; we defer a more detailed analysis of how different types of re-openings affect different sectors (which can be undertaken with the data we have made publicly available) to future work.

We begin with a case study comparing Colorado and New Mexico that is representative of our broader findings. These two states both issued stay-at-home orders during the final week of March (New Mexico on March 24, Colorado on March 26). Colorado then partially reopened its economy, permitting a larger group of businesses to operate, on May 1, while New Mexico did not re-open until two weeks later, on May 16.⁴³

Figure 12a plots consumer spending (using the Affinity Solutions data) in Colorado and New Mexico. Spending evolved on a nearly identical path in these two states: in particular, there is no evidence that the earlier reopening in Colorado did anything to boost spending during the two intervening weeks before New Mexico reopened.

Figure 12b generalizes the case study in Figure 12a by studying partial reopenings in the 20 states that issued such orders on or before May 4. For each reopening date (of which there are five: April 20, 24th, and 27, as well as May 1 and 4), we compare the trajectory of spending in treated states to a group of control states selected from the group of 13 states that did not issue reopening orders until after May 18. We select the control states for each of the five reopening dates by choosing nearest-neighbor matches on pre-period levels of spending (relative to January) during the weeks ending March 31, April 7, and April 19. Appendix Table 6 lists the control states we use for each date. We then calculate unweighted means of the outcome variables in the control

43. Specifically, on May 1, Colorado allowed retail businesses to open to the public beyond curbside pick-up and delivery, and permitted personal services businesses to re-open.

and treatment states to construct the two series for each reopening date. Finally, we pool these five event studies together (redefining calendar time as time relative to the reopening date) to create Figure 12b.

Just as in the case study of Colorado vs. New Mexico, the trajectories of spending in the treated states almost exactly mirror that in the control states. We formalize the estimate from this design using a difference-in-differences design that compares the two weeks before the reopening in the treated states and two weeks after. The estimated effect is that reopenings led to a 1.04 p.p. increase in spending. This estimate also appears in Table 3, Column 1. Column 2 replicates that specification but focusing on the earlier reopenings, so that we can go out three weeks after the event; the estimate is even lower here, at just 0.18 p.p. The evidence does not suggest that reopenings increased spending, at least in the first several weeks.

Figure 12c shows a small positive effect of reopenings on employment (using data from Paychex and Earnin), with a 1.52 p.p. effect over two weeks that is statistically significant. Column 4 of Table 3 shows that this effect grows to 2.75 p.p. after three weeks. Columns 5 and 6 show that this effect is primarily concentrated among low-wage employees, although the estimates are somewhat imprecise. Figure 12d shows the likely driver of this small increase in employment: there was a 3.69 p.p. increase in merchants open after states allowed businesses to reopen (using data from Womply).

Despite these positive effects on employment, the reopenings account for a relatively small share of the overall variation in economic conditions across states. To demonstrate this, we first calculate the actual variance in spending levels across states. We then counterfactually add our estimate effect of reopening to all states that were not yet open as of May 18, and recalculate the variance. Figure 12e then plots the 1 minus the ratio of the counterfactual variance to the actual variance, which is a measure of the importance of early reopenings in explaining the variation in economic activity. These ratios are very low even for employment, showing that early reopenings did not play an important role in explaining why some states had stronger employment trajectories than others.⁴⁴ These results are consistent with the findings of Goolsbee and Syverson (2020) and Lin and Meissner (2020), who use a state-border discontinuity design and find no impact of stay-at-home orders on job losses.

44. We emphasize that these results apply to *average* employment rates for *low-income* workers and are thus not inconsistent with evidence of modest impacts in specific subsectors, particularly at higher wage levels, as identified e.g., by Cajner et al. (2020).

Why did these reopenings have so little immediate impact on economic activity?⁴⁵ The evidence in Section 3 suggests that health concerns among consumers were the primary driver of the sharp decline in economic activity in March and April. Consistent with that evidence, spending fell sharply in most states *before* formal state closures (Appendix Figure 16). If health concerns are the core driver of reductions in spending rather than government-imposed restrictions, governments may have limited capacity to restore economic activity through reopenings, especially if those reopenings are not interpreted by consumers as a clear signal of reduced health risks.

IV.B Stimulus Payments to Households

The Coronavirus Aid, Relief, and Economic Security (CARES) Act made direct payments to nearly 160 million people, totaling \$267 billion as of May 31, 2020. Individuals earning less than \$75,000 received a stimulus payment of \$1,200; married couples earning less than \$150,000 received a payment of \$2,400; and households received an additional \$500 for each dependent they claimed. These payments were reduced at higher levels of income and phased out entirely for households with incomes above \$99,000 (for single filers without children) or \$198,000 (for married couples without children). The vast majority of these stimulus payments were deposited on exactly April 15, 2020, while some households received payments on April 14 (Appendix Figure 17).

The goal of these stimulus payments was to increase consumer spending and restore employment.⁴⁶ Was the stimulus effective in achieving these goals? In this section, we analyze this question using high-frequency event studies examining spending and employment changes in the days surrounding April 15, comparing outcomes for lower-income and higher-income households.

Impacts on Consumer Spending. We begin in Figure 13a by plotting a weekly moving average of spending changes relative to mean levels in January for low-income (bottom income quartile) vs. high-income (top income quartile ZIP codes) households. As noted above, high-income households decreased spending by more than low-income households in the immediate aftermath of the COVID shock; in the week ending April 13th, spending in top-income-quartile households was down by 37% relative to pre-COVID levels, as compared with 28% for bottom-income-quartile households. Starting on April 15, spending rose very sharply for those in the bottom income quartile, increasing by over 15 percentage points within a week. Spending among top-income-quartile households increased as well, but by only about 7 percentage points. This simple analysis suggests that the

45. Reopenings could have a lagged effect on spending, particularly if they serve as a signal of changes in health risks; going forward, the real-time data in the tracker can be used to assess such lagged impacts.

46. The Congressional Budget Office (2020) estimates that these payments will cost \$293 billion, a considerably larger sum than similar direct stimulus in 2001 and 2008.

stimulus payments had a large positive effect on spending, especially for low-income families.⁴⁷

To estimate the causal effect of the stimulus payments more precisely, we use a regression discontinuity estimator with the daily spending data.⁴⁸ Figures 13b and 13c plot daily spending levels relative to baseline for low- and high-income households, respectively, for the month of April. Spending levels jumped sharply from April 13th to 15th. Fitting a linear approximation to the points on either side of the stimulus, we estimate that spending levels rose discontinuously on April 15 by 26pp in low-income households and 9pp in high-income households.⁴⁹ Both effects are statistically significantly different from 0, as well as from each other. Panel A of Table 4 shows these regression discontinuity estimates under a variety of bandwidths. These findings are consistent with Baker et al. (2020) and Karger and Rajan (2020), who use individual transaction data on incomes and spending patterns of approximately 15,000 primarily low-income individuals to estimate a large and immediate effect of receiving the stimulus check on spending, especially among the very poorest households.

In Figures 13d and 13e, we investigate the composition of goods on which households spent their stimulus checks. We pool all households in these figures to maximize precision. Figure 13d shows that spending on durable goods rose by 21 pp following the arrival of the stimulus payments and further increased thereafter, rising well above pre-crisis levels. But Figure 13e shows that spending on in-person services rose by only 7 pp, remaining more than 50% below pre-crisis levels. Durable goods accounted for 44% of the recovery in spending levels from the beginning to the end of April, despite accounting for just 23% of pre-crisis spending. In-person services accounted for just 18% of the recovery, despite making up 32% of pre-crisis spending (Appendix Figure 18).⁵⁰ These results show that the stimulus increased the overall level of spending, but did not increase spending in the sectors where spending fell most following the COVID shock (Figure 2b). As a result, the stimulus did not channel money back to the businesses that lost the most revenue as a result of the COVID shock.

Impacts on Business Revenue Across Areas. Next, we investigate how the stimulus program

47. We expect the stimulus program to have a smaller impact on high-income households for three reasons. First, lower-income households simply received more money than high-income households. Second, low-income households spend half as much as high-income households prior to the COVID shock (Figure 2a), and hence one would expect a larger impact on their spending levels as a percentage of baseline spending. Finally, many studies have found higher marginal propensities to consume (MPCs) among lower-income households, who are often more liquidity constrained.

48. We use the raw daily data, not the 7-day moving average.

49. We omit the partially treated date of April 14 (denoted by a hollow dot) since a small fraction of stimulus payments arrived on that day when estimating this RD specification.

50. The other major spending categories (non-durable goods and remote services) each accounted for 19% of the recovery and 23% and 21% of pre-crisis spending, respectively.

affected business revenues across areas. In particular, did the businesses that lost the most revenue – those in high-rent areas – gain business as a result of the stimulus? Figures 14a and 14b replicate the analysis above using Womply data on small business revenues as the outcome, separately for lowest-rent-quartile and highest-rent-quartile ZIP codes. We see a sharp increase of 21 pp in revenues in small businesses in low-rent neighborhoods exactly at the time when households received stimulus payments. In contrast, Panel B shows a small, statistically insignificant increase in revenues of 4 pp for small businesses in high-rent areas.

This geographic heterogeneity illustrates another important dimension in which the stimulus did not channel money back to the business that lost the most revenue from the COVID shock. In fact, the stimulus actually *amplified* the difference in small business revenue losses rather than narrowing it across areas. Those in low-rent areas have nearly returned to pre-crisis levels following the stimulus payments, while those in high-rent areas remained nearly 40% down relative to January levels in the second half of April (Figure 14c, solid lines). Panel B of Table 4 shows these regression discontinuity estimates under a variety of bandwidths.

Impacts on Low-Income Employment. Finally, we investigate whether the increase in spending induced by the stimulus increased employment rates, as one would expect in a traditional Keynesian stimulus. Here, we do not use the RD design as we do not expect employment to respond immediately to increased spending. Instead, we analyze the evolution of employment of low-income workers in the Earnin data in low vs. high-rent ZIP codes over time in Figure 14c (dashed lines). In high-rent areas, low-wage employment remains 45% below pre-COVID levels – perhaps not surprisingly, since revenues have not recovered significantly there. But even in low rent areas, employment has recovered only partially, despite the fact that small business revenues have reverted to pre-COVID baseline levels. It is unclear why revenues and employment both *fell* in tandem at very similar rates when the COVID shock hit, but revenues recovered much more quickly than employment in low-rent areas. One possibility is that businesses have reopened temporarily with a minimal staff (Lazear, Shaw, and Stanton 2016) and are planning to recall or hire new workers going forward. A more worrisome possibility is a “jobless” recovery, in which economic activity shifts away from in-person labor intensive production, reducing employment opportunities in the longer term (Berger 2012).

In summary, our analysis suggests that stimulus substantially increased total consumer spending but did not directly undo the initial spending reductions by returning money back to the businesses that lost the most revenue. This impact contrasts with the canonical theoretical motivation for

stimulus spending in response to shocks. In particular, Farhi and Werning (2016) show that optimal macroprudential policy involves a stimulus that increases spending in sectors and areas whose demand is depressed. In a frictionless model where businesses and workers could costlessly reallocate their capital and labor to other sectors, the reallocation of spending might have no consequence for employment levels. But if workers’ ability to switch jobs is constrained – e.g., because of job-specific skills that limit switching across industries or costs that limit moving across geographic areas, as suggested by Yagan (2019) – the ability of the stimulus to foster a uniform recovery in employment to pre-COVID levels is likely to be hampered, perhaps explaining why employment levels remained well below baseline even as total spending recovered after April 15 (Figure 7).

IV.C Loans to Small Businesses

We now turn to evaluating the Paycheck Protection Program (PPP), a policy that sought to reduce employment losses by providing financial support to small businesses. Congress appropriated nearly \$350 billion for loans to small businesses in an initial tranche that was paid beginning on April 3, followed by another \$175 billion in a second round beginning on April 27. The program offered loan forgiveness for businesses that maintained sufficiently high employment (relative to pre-crisis levels). The stated primary purpose of the PPP was to encourage businesses to maintain employment even as they lost revenue (House Committee on Small Business 2020). For example, the Small Business Administration (Small Business Administration 2020) emphasized the employment impacts of the PPP as a key measure of the program’s success, noting that the PPP “ensure[d] that over approximately 50 million hardworking Americans stay[ed] connected to their jobs” based on self-reports of the number of jobs retained by firms that received the PPP.

Here, we study the impacts of the PPP on employment directly using payroll data, exploiting the fact that eligibility for the PPP depended on business size. Firms with fewer than 500 employees before the COVID crisis qualified for PPP loans, while those with more than 500 employees generally did not.⁵¹ One important exception to this rule is the food service industry, which was treated differently because of the prevalence of franchises. We therefore omit the food services sector from the analysis that follows.

51. The eligibility rules vary across industries, with some exceptions that allow larger firms to obtain loans. Appendix Figure 19 plots a histogram of the exact size cutoffs weighting by employees in the national sample in ReferenceUSA data, restricting to workers in companies with 300-700 employees. More than 90% of employees work at firms that face the 500 employee threshold. In addition to employment thresholds, firms may also qualify based on revenue thresholds set by the Small Business Administration; however, using the distribution of firm size and revenue from Reference USA, we estimate that in practice the size threshold is the binding constraint for the vast majority of firms. Given these results, we use a pre-COVID employee size cutoff of 500 to define treatment and control groups.

We estimate the causal effect of the PPP on employment rates at small businesses using a difference-in-differences research design, comparing trends in employment for firms below the 500 employee cutoff (the treated group) vs. those above the 500 employee cutoff (the control group) before vs. after April 3, when the PPP program began. Figure 15a plots the average change in employment rates (inferred from payroll deposits) relative to January for firms in the Paychex-Earnin data employing 100-500 employees, which were eligible for PPP loans, vs. firms employing 501-800 employees, which were generally ineligible for PPP loans.⁵² To adjust for the fact that industry composition varies across firms of different sizes, we reweight firms within each decile to match the average (2 digit NAICS) industry composition in January in the sample as a whole when computing mean employment rates by firm size. We also residualize employment rates by county x wage quartile x week fixed effects, to account for the differential time patterns of employment rates by county and wage quartile shown in Section III.C.

Before April 3, trends in employment are similar among eligible vs. ineligible firms, showing that larger businesses provide a good counterfactual for employment trends one would have observed in smaller firms absent the PPP program. After April 3, employment continues to follow almost an identical trajectory in the treated group (<500 employees) and the control group, implying that the PPP program had little impact on employment at small businesses under the identification assumption that employment trends in the two groups would have remained similar absent the PPP. Notably, even after PPP implementation, employment levels for low-wage workers remain well below baseline – showing that the program did not close a significant share of the loss in jobs resulting from the COVID shock.

Figure 15b plots the change in employment from January 8-January 31 to June 1-June 23 by firm size bin. The decline in employment is quite similar across firm sizes, and in particular is not markedly smaller for firms below the 500 employee eligibility threshold.⁵³ Appendix Figure 20 shows that we obtain very similar results in the Kronos data: the employment trajectory of ineligible firms closely matches the employment trajectory of eligible firms, both prior to and after the beginning of the PPP.

In Table 5, we quantify the impacts of the PPP using a set of regression models of the form:

$$\text{Emp}_{scqt} = \alpha_{cqt} + \delta \text{Eligible}_s + \gamma \text{Post-PPP}_t + \beta_{DD} \text{Eligible}_s \cdot \text{Post-PPP}_t + \varepsilon_{scqt}, \quad (1)$$

52. Since Intuit consists primarily of firms with fewer than 20 employees, we omit it from this analysis.

53. Because of differences in the measurement of firm sizes in our data and the SBA data used to determine PPP eligibility (see below), there is no sharp discontinuity in eligibility at the 500 cutoff. Hence, we do not interpret this plot using an RD design, but rather view it as showing that our estimates are insensitive to the bandwidth used to define the treatment and control groups in the DD analysis.

where Emp_{scqit} is the change in employment within each firm size category $s \times$ county $c \times$ wage quartile $q \times$ 2-digit NAICS industry i on week t cell, relative to January 11-January 31, Eligible_s is an indicator variable for whether firm size is 500 or fewer employees in the pre-COVID period, Post-PPP_t is an indicator variable for the date being on or after 3 April 2020, and α_{cqt} represents a county-week-wage quartile fixed effect. We estimate this regression on the sample of firms with 100-800 employees using data from March 11 to August 15. We reweight by two-digit NAICS code so that the distribution of industries in the below-500 and above-500 employee groups match the overall distribution of industries in January 2020. We cluster standard errors at the county-by-industry-by-firm-size level to permit correlation in errors across firms and over time within counties and estimate the regression using OLS, weighting by the total number of employees in the cell from January 11-31, 2020.

Column 1 presents the baseline estimate obtained by estimating (1) of $\beta_{DD} = 1.05$ ($s.e. = 1.62$), an estimate that matches the figure plotted in Figure 15a. The mean decline in employment among firms in the control group to August 15 was 18.7%, implying that the PPP saved 5.6% of the jobs that would otherwise have been lost. In Column 2, we reduce the bandwidth to focus more narrowly around the 500-employee size threshold; the estimates remains statistically indistinguishable from that in Column 1. Columns 3 and 4 replicate the specification in Column 1, using data only from Earnin and Kronos, respectively.⁵⁴ Since these data sources only cover workers in the lowest wage quartile, we include fixed effects only at the county-by-week level. The estimates from these data sources are also similar to our baseline estimate.

Note that our difference-in-differences research design identifies the causal effect of the PPP on eligible firms under the assumption that the PPP did not have a causal effect on employment at PPP-ineligible firms. However, it is possible that the PPP reduced employment at ineligible firms (relative to the no-PPP counterfactual) through an employment substitution channel: ineligible firms might have hired workers laid off from eligible firms in the absence of the PPP. In the presence of such substitution, our DD estimate would overstate the causal effect of the PPP on employment at small businesses, providing an upper bound for the partial equilibrium impact of the PPP (ignoring general equilibrium effects that may have influenced consumer demand and employment at all firms).

Measurement Error in Firm Sizes. Our measures of firm size – which are based on employment levels in 2018 from the ReferenceUSA database for the Earnin sample and 2019 employment for the

54. Our data-use agreements do not permit us to report results based solely on Paychex data.

Paychex and Kronos data – do not correspond precisely to the measures used by the Small Business Administration to determine PPP eligibility. Such measurement error in firm size attenuates the estimates of β_{DD} obtained from (1) relative to the true causal effect of PPP eligibility because some of the firms classified as having more than 500 employees may have actually received PPP (and vice versa). We estimate the degree of this attenuation bias by matching our data on firm sizes to data publicly released by the Small Business Administration (SBA) on a selected set of PPP recipients and assessing the extent to which firms are mis-classified around the threshold. We estimate that our reduced-form estimates are attenuated by factor of 1.9 based on this matched data (see Appendix D for details). Under standard assumptions required to obtain a local average treatment effect in the presence of non-compliance – no direct effect of being classified as having more than 500 workers independent of the PPP and a monotonic treatment effect – we can estimate the LATE of the PPP on employment rates by dividing the raw estimates reported in Table 1 by $0.977 - 0.454 = 52.2\%$ (Angrist, Imbens, and Rubin 1996). This gives us a final preferred point estimate for the effect of PPP eligibility on employment of 2.02 percentage points.

Costs Per Job Saved. Using SUSB data, we calculate that approximately 53.6 million workers work at firms eligible for PPP assistance, excluding firms in NAICS 72.⁵⁵ Under the assumption that the PPP’s effects on firms with between 100 and 500 employees are the same in percentage terms as the PPP’s effects on all eligible firms, our baseline estimates in the combined Paychex-Earnin data (Column 1 of Table 5), adjusted for attenuation bias, imply that the PPP saved $0.02 \times 53.6M = 1.08$ million jobs from April through August 15. Given a total expenditure on the PPP program of \$486 billion through August 8 (excluding firms in food services), this translates to an average cost per job saved by the PPP of \$451,000.⁵⁶ Even at the upper bound of the 95% confidence interval for employment impact, we estimate a cost per job saved of \$114,000. For comparison, mean annual earnings for workers at PPP-eligible firms are only \$45,000.

Why did the PPP have relatively small effects on employment rates despite having a very high takeup rate among small businesses? One potential explanation is that the loans were taken by firms that did not intend to layoff many employees to begin with, i.e. firms that were inframarginal

55. To compute this estimate, we use 2017 SUSB data on firm sizes. We first exclude NAICS 72 firms. We then calculate total employment in firms with fewer than 500 employees as 52.0 million workers. Following Autor et al. (2020), we adjust for 3% growth in private sector payrolls since 2017 to arrive at a final estimate of 53.6 million workers in eligible firms outside NAICS 72.

56. To compute total PPP expenditure on non-NAICS 72 firms, we first use publicly released data on loan recipients to calculate that 92.1% of total PPP expenditure was received by non-NAICS 72 firms. We then multiply this share by total PPP expenditure as of 8 August to reach an estimate of \$486 billion in non-NAICS 72 firms. If the treatment effect of the PPP program on food services was the same in percentage terms as in other sectors, we estimate a total of 1.26 million jobs saved by the PPP.

recipients of loans. Consistent with this hypothesis, Granja et al. (2020) show that states and congressional districts that experienced more job losses prior to April 3 actually received *fewer* PPP loans. Moreover, PPP loans also were not distributed to the industries most likely to experience job losses from the COVID crisis. For example, firms in the professional, scientific, and technical services industry received a greater share of the PPP loans than accommodation and food services (SBA 2020). Yet accommodation and food services accounted for half of the total decline in employment between February and March (prior to PPP enactment) in BLS statistics, while employment in professional, scientific and technical services accounted for less than 5% of the decline.

Although the PPP had modest impacts on employment – consistent with the estimates from studies using other data sources – the program may have had other important benefits, such as reducing the rate of small business closures. As emphasized by Hubbard and Strain (2020), if the disruption costs of closing and restarting businesses are sufficiently large, the PPP may have still have significant benefits over time – an important question for future research.

V Conclusion

Transactional data held by private companies have great potential for measuring economic activity, but to date have been accessible only internally within companies or through contracts to work with confidential microdata. In this paper, we have constructed a public platform to measure economic activity at a high-frequency, granular level using data from private companies. By carefully aggregating and masking the underlying micro data, we construct series that can be released publicly without disclosing sensitive information, yet are well suited to answer a variety of research questions.

We apply these new data to analyze the economic impacts of COVID-19 in the first three months after the pandemic began to spread in the U.S. (March 15 - June 15, 2020). We find that during this period, COVID-19 induced high-income households to self-isolate and sharply reduce spending in sectors that require physical interaction. This spending shock in turn led to losses in business revenue and layoffs of low-income workers at firms that cater to high-income consumers, ultimately reducing their own consumption levels. Because the root cause of the shock is self-isolation driven by health concerns, there is limited capacity to restore economic activity without addressing the virus itself. In particular, we find that state-ordered reopenings of economies have only modest impacts on economic activity; stimulus checks increase spending particularly among low-income households, but very little of the additional spending flows to the businesses most affected by the

COVID shock; and loans to small businesses have little impact on employment rates. Our analysis therefore suggests that the most effective approach to mitigating economic hardship in the midst of a pandemic may be to provide benefits to those who have lost their incomes to mitigate consumption losses while investing in public health measures restore consumer confidence and ultimately increase spending.

We focused in this paper on the short-run economic consequences of the COVID-19 pandemic. However, such shocks can also have long-lasting scarring effects that warrant attention. Private sector data can be useful in measuring these impacts in real time as well. As an illustration, Figure 16 plots weekly student progress (lessons completed) on Zearn, an online math platform used by many elementary school students as part of their regular school curriculum. Children in high-income areas experienced a temporary reduction in learning on this platform when the COVID crisis hit and schools shifted to remote instruction, but soon recovered to baseline levels. By contrast, children in lower-income areas remained 50% below baseline levels persistently. Although this platform captures only one aspect of education, these findings raise the concern that pandemics may reduce social mobility and ultimately further amplify inequality by having particularly negative effects on human capital development for lower-income children.

More broadly, beyond its implications for the economics of pandemics, our analysis demonstrates two broad ways in which the public platform constructed here provides a new tool for empirical macroeconomics. First, the data can be used to learn rapidly from sub-national heterogeneity, as different places, sectors, and subgroups are often hit by different shocks and pursue different local policy responses. This approach can permit rapid diagnosis of the root factors underlying an economic crisis. Second, the data permit rapid policy evaluation – often within two or three weeks of implementation – opening a path to fine-tuning policy responses in an evidence-based manner.

The advantage of constructing a public platform to support such analyses rather than working directly with the underlying confidential data held by private sector firms is that it permits a much broader range of downstream work along these lines. For example, the data on the platform are now being used by local policymakers to inform local policy responses and forecast tax revenue impacts (e.g., [Maine](#), [Missouri](#), [Kansas](#), and [Texas](#)). They are also being used by Congressional staff to design federal policies, e.g. predicting the impacts and costs of policies targeted based on business revenue losses or other outcomes (RESTART Act 2020). And they are being used by other researchers to analyze a broad range of issues: constructing more accurate price indices that account for changes in consumption bundles (Cavallo 2020), analyzing the effects of political

views on economic outcomes (Makridis and Hartley 2020), estimating the effects of the Paycheck Protection Program on small firms’ employment decisions (Granja et al. 2020), and estimating the impacts of changes in unemployment benefits on aggregate spending (Casado et al. 2020).

The platform built here can be viewed as a prototype for a system of “real time national accounts” using administrative data from the private sector, much as the Bureau of Economic Analysis, building on a prototype developed by Kuznets (1941), instituted a set of systematic, recurring surveys of businesses and households that are the basis for the National Income and Product accounts. Our analysis demonstrates that even this first prototype yields timely insights that are not apparent in existing data, suggesting that a broader, more refined platform that aggregates data from additional private companies has great potential for improving our understanding of economic activity and policymaking going forward.

References

- Abraham, Katharine G, Ron S Jarmin, Brian Moyer, and Matthew D Shapiro (ed.) 2019. *Big Data for 21st Century Economic Statistics*. NBER Book Series Studies in Income / Wealth.
- Administration, Small Business. 2020. *Paycheck Protection Program (PPP) Report*. Technical report. May. <https://www.sba.gov/document/report--paycheck-protection-program-ppp-report>.
- Aladangady, Aditya, Shifrah Aron-Dine, Wendy Dunn, Laura Feiveson, Paul Lengermann, and Claudia Sahm. 2019. "From Transactions Data to Economic Statistics: Constructing Real-time, High-frequency, Geographic Measures of Consumer Spending." *NBER Working Paper No. 26253* (September). doi:[10.3386/w26253](https://doi.org/10.3386/w26253). <http://www.nber.org/papers/w26253>.
- Alexander, Diane, and Ezra Karger. 2020. "Do stay-at-home orders cause people to stay at home? Effects of stay-at-home orders on consumer behavior." *Federal Reserve Bank of Chicago Working Paper No. 2020-12* (April). doi:[10.21033/wp-2020-12](https://doi.org/10.21033/wp-2020-12). <https://www.chicagofed.org/publications/working-papers/2020/2020-12>.
- Allcott, Hunt, Levi Boxell, Jacob Conway, Billy Ferguson, Matthew Gentzkow, and Benjamin Goldman. 2020. "Economic and Health Impacts of Social Distancing Policies during the Coronavirus Pandemic." *Available at SSRN* 53, no. 3 (May): 571–630. <https://ssrn.com/abstract=3610422>.
- Allen, Danielle, Sharon Block, Joshua Cohen, Peter Eckersley, and Meredith Rosenthal. 2020. "Roadmap to Pandemic Resilience: Massive Scale Testing, Tracing, and Supported Isolation (TTSI) as the Path to Pandemic Resilience for a Free Society." *Edmond J. Safra Center For Ethics At Harvard University*.
- Altonji, Joseph, Zara Contractor, Lucas Finamor, Ryan Haygood, Ilse Lindenlaub, Costas Meghir, Cormac O'Dea, Dana Scott, Liana Wang, and Ebonya Washington. 2020. "Employment Effects of Unemployment Insurance Generosity During the Pandemic." *Yale University Manuscript*.
- Angrist, Joshua D, Guido W Imbens, and Donald B Rubin. 1996. "Identification of causal effects using instrumental variables." *Journal of the American statistical Association* 91 (434): 444–455.
- Austin, Benjamin A, Edward L Glaeser, and Lawrence H Summers. 2018. "Jobs for the Heartland: Place-based policies in 21st century America." *NBER Working Paper No. 24548* (April). <https://www.nber.org/papers/w24548>.
- Autor, David, David Cho, Leland D Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz. 2020. *An Evaluation of the Paycheck Protection Program Using Administrative Payroll Microdata*. Technical report.

- Baker, Scott R, R. A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine” Yannelis. 2020. “Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments.” *NBER Working Paper No. 27097*.
- Bartik, Alexander W., Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matt Unrath. 2020. “Measuring the labor market at the onset of the COVID-19 crisis.” *Brookings Papers on Economic Activity* (June). <https://www.brookings.edu/wp-content/uploads/2020/06/Bartik-et-al-conference-draft.pdf>.
- Bartlett, Robert P, and Adair Morse. 2020. “Small Business Survival Capabilities and Policy Effectiveness: Evidence from Oakland.” *NBER Working Paper No. 27629*.
- Bennet, M. 2020. “S 3814-RESTART Act.” *Senate - Finance Committee*.
- Berger, David. 2012. “Countercyclical restructuring and jobless recoveries.”
- Bitler, Marianne P, Hilary W Hoynes, and Diane Whitmore Schanzenbach. 2020. “The social safety net in the wake of COVID-19.” *Brookings Papers on Economic Activity* 13.
- Blanchard, Olivier, and Lawrence Katz. 1992. “Regional Evolutions.” *Brookings Papers on Economic Activity* 1992 (1): 1–61.
- Board of Governors of the Federal Reserve System. 2019. *The 2019 Federal Reserve Payments Study*.
- Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. “The U.S. Labor Market during the Beginning of the Pandemic Recession.” *Working Paper* (May).
- Cajner, Tomaz, Leland D Crane, Ryan A Decker, Adrian Hamins-Puertolas, and Christopher Kurz. 2019. *Improving the Accuracy of Economic Measurement with Multiple Data Sources: The Case of Payroll Employment Data*. Technical report. National Bureau of Economic Research.
- Carnevale, Anthony P, Tamara Jayasundera, and Dmitri Repnikov. 2014. “Understanding online job ads data.” *Georgetown University, Center on Education and the Workforce, Technical Report* (April).
- Casado, Miguel Garza, Britta Glennon, Julia Lane, David McQuown, Daniel Rich, and Bruce A Weinberg. 2020. “The Effect of Fiscal Stimulus: Evidence from COVID-19.”
- Cavallo, Alberto. 2020. “Inflation with Covid Consumption Baskets.”
- Chen, Haiqiang, Wenlan Qian, and Qiang Wen. 2020. “The Impact of the COVID-19 Pandemic on Consumption: Learning from High Frequency Transaction Data.” *Working Paper* (April). doi:<http://dx.doi.org/10.2139/ssrn.3568574>. <https://ssrn.com/abstract=3568574>.

- Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *The Quarterly Journal of Economics*.
- Chiou, Lesley, and Catherine Tucker. 2020. "Social Distancing, Internet Access and Inequality." *NBER Working Paper No. 26982* (April). doi:[10.3386/w26982](https://doi.org/10.3386/w26982). <http://www.nber.org/papers/w26982>.
- Congressional Budget Office. 2020. *Preliminary Estimate of the Effects of H.R. 748, the CARES Act*.
- Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, and Fiona Greig. 2020. "Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data." *Brookings Papers on Economic Activity*.
- Deming, David, and Lisa B. Kahn. 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals." *Journal of Labor Economics* 36 (S1): S337–S369. doi:[10.1086/694106](https://doi.org/10.1086/694106). <https://doi.org/10.1086/694106>.
- Diamond, Peter, and OJ Blanchard. 1989. "The beveridge curve." *Brookings Papers on Economic Activity* 1:1–76.
- Dunn, Abe, Kyle Hood, and Alexander Driessen. 2020. "Measuring the Effects of the COVID-19 Pandemic on Consumer Spending Using Card Transaction Data." *National Bureau of Economic Research*.
- Ehrlich, Gabriel, John Haltiwanger, Ron Jarmin, David Johnson, and Matthew D Shapiro. 2019. "Re-engineering Key National Economic Indicators." In *Big Data for 21st Century Economic Statistics*. University of Chicago Press.
- Elsby, Michael WL, Ryan Michaels, and David Ratner. 2015. "The Beveridge curve: A survey." *Journal of Economic Literature* 53 (3): 571–630.
- Farhi, Emmanuel, and Iván Werning. 2016. "A theory of macroprudential policies in the presence of nominal rigidities." *Econometrica* 84 (5): 1645–1704.
- Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, and Joseph Vavra. 2020. *Consumption Effects of Unemployment Insurance during the COVID-19 Pandemic*. Technical report. JPMorgan Chase Institute.
- Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer. 2015. "The next generation of the Penn World Table." *American Economic Review* 105 (10): 3150–82.
- Ganong, Peter, Pascal J Noel, and Joseph S Vavra. 2020. *US Unemployment Insurance Replacement Rates During the Pandemic*. Technical report. National Bureau of Economic Research.

- Gindelsky, Marina, Jeremy Moulton, and Scott A Wentland. 2019. “Valuing housing services in the era of big data: A user cost approach leveraging Zillow microdata.” In *Big Data for 21st Century Economic Statistics*. University of Chicago Press.
- Goldfarb, Avi, and Catherine Tucker. 2020. “Which Retail Outlets Generate the Most Physical Interactions?” *NBER Working Paper No. 27042* (April). doi:[10.3386/w27042](https://doi.org/10.3386/w27042). <http://www.nber.org/papers/w27042>.
- Goolsbee, Austan, and Chad Syverson. 2020. *Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020*. Working Paper, Working Paper Series 27432. National Bureau of Economic Research, June. <http://www.nber.org/papers/w27339>.
- Granja, João, Christos Makridis, Constantine Yannelis, and Eric Zwick. 2020. “Did the Paycheck Protection Program Hit the Target?” *NBER Working Paper No. 27095* (May). doi:[10.3386/w27095](https://doi.org/10.3386/w27095). <http://www.nber.org/papers/w27095>.
- Greene, Claire, and Joanna Stavins. 2020. *2019 Diary of Consumer Payment Choice*. Technical report. Federal Reserve Bank of Atlanta.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. *Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?* Working Paper, Working Paper Series 26918. National Bureau of Economic Research, April. doi:[10.3386/w26918](https://doi.org/10.3386/w26918). <http://www.nber.org/papers/w26918>.
- Hershbein, Brad, and Lisa B. Kahn. 2018. “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings.” *American Economic Review* 108, no. 7 (July): 1737–72. doi:[10.1257/aer.20161570](https://doi.org/10.1257/aer.20161570). <https://www.aeaweb.org/articles?id=10.1257/aer.20161570>.
- House Committee on Small Business. 2020. *Oversight of the Small Business Administration and Department of Treasury Pandemic Programs: Hearing Before The House Committee on Small Business, 116th Cong. (Testimony of Steven Mnuchin)*.
- Hubbard, Glenn, and Michael R Strain. 2020. “Has the Paycheck Protection Program Succeeded?” *Brookings Institution*.
- Jaimovich, Nir, and Henry E Siu. 2020. “Job polarization and jobless recoveries.” *Review of Economics and Statistics* 102 (1): 129–147.
- Kahn, Lisa B, Fabian Lange, and David G Wiczer. 2020. “Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims.” *NBER Working Paper No. 27061* (April). doi:[10.3386/w27061](https://doi.org/10.3386/w27061). <http://www.nber.org/papers/w27061>.
- Kaplan, Greg, and Giovanni L Violante. 2014. “A model of the consumption response to fiscal stimulus payments.” *Econometrica* 82 (4): 1199–1239.

- Karger, Ezra, and Aastha Rajan. 2020. "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments." *FRB of Chicago Working Paper*.
- Kurmann, André, Etienne L   , and Lien Ta. 2020. "The Impact of COVID-19 on U.S. Employment and Hours: Real-Time Estimates with Homebase Data" (May). http://www.andrekurmann.com/hb_covid.
- Kuznets, Simon. 1941. *National Income and Its Composition, 1919-1938*. New York: National Bureau of Economic Research.
- Lazear, Edward P., Kathryn L. Shaw, and Christopher Stanton. 2016. "Making Do with Less: Working Harder during Recessions." *Journal of Labor Economics* 34 (S1): S333–S360. doi:10.1086/682406. eprint: <https://doi.org/10.1086/682406>. <https://doi.org/10.1086/682406>.
- Lin, Zhixian, and Christopher M Meissner. 2020. "Health vs. Wealth? Public Health Policies and the Economy During Covid-19." *NBER Working Paper No. 27099* (May). doi:10.3386/w27099. <http://www.nber.org/papers/w27099>.
- Makridis, Christos, and Jonathan Hartley. 2020. "The Cost of Covid-19: A Rough Estimate of the 2020 US GDP Impact."
- Mastercard. 2020. *Mastercard Recovery Insights: The Shift to Digital*.
- Mathy, Gabriel. 2020. *The COVID-19 Epidemic will be the First Services Recession and it Could be a Bad One*.
- Mian, Atif, and Amir Sufi. 2009. "The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis." *The Quarterly Journal of Economics* 124 (4): 1449–1496.
- Mongey, Simon, Laura Pilossoph, and Alex Weinberg. 2020. "Which Workers Bear the Burden of Social Distancing Policies?" *NBER Working Paper No. 27085* (May). doi:10.3386/w27085. <http://www.nber.org/papers/w27085>.
- Paychex. 2020. *Small Business Employment Watch*. <https://www.paychex.com/employment-watch/#/>, April.
- Petev, Ivaylo, Luigi Pistaferri, and Itay Saporta Eksten. 2011. *Consumption and the Great Recession: An analysis of trends, perceptions, and distributional effects*.
- Romer, Paul. 2020. "Roadmap to Responsibly Reopen America." *roadmap.paulromer.net*. <http://roadmap.paulromer.net>.
- Small Business Administration. 2020. *Joint Statement by SBA Administrator Jovita Carranza and U.S. Treasury Secretary Steven T. Mnuchin Regarding Enactment of the Paycheck Protection Program Flexibility Act*.

- Summers, Robert, and Alan Heston. 1991. “The Penn World Table (Mark 5): an expanded set of international comparisons, 1950–1988.” *The Quarterly Journal of Economics* 106 (2): 327–368.
- U.S. Bureau of Economic Analysis. 2020. *National Income and Product Accounts*. Data retrieved from U.S. Bureau of Economic Analysis, National Income and Product Accounts.
- Villas-Boas, Sofia B, James Sears, Miguel Villas-Boas, and Vasco Villas-Boas. 2020. “Are We #StayingHome to Flatten the Curve?” *UC Berkeley: Department of Agricultural and Resource Economics CUDARE Working Papers* (April). <https://escholarship.org/uc/item/5h97n884>.
- Yagan, Danny. 2019. “Employment Hysteresis from the Great Recession.” *Journal of Political Economy* 127 (5): 2505–2558.

Online Appendix

A Automated Data Processing Pipeline

This appendix provides an overview of the automated pipeline we built to ingest raw data, process it to construct aggregate statistics, and then release those statistics publicly. This automated pipeline allows us to typically post updated statistics within one business day of receiving the raw data. By automating the data processing to the extent possible, we aim to post data as close to real-time as possible, while maintaining the quality of the data and minimizing the manual upkeep required. The primary source of lags in the posted data is therefore driven by lags in the underlying data generating processes: for example, card transactions can take up to a week to settle and employment income is typically paid in biweekly or monthly payrolls. We summarize our data engineering methods here for those who may be interested in setting up similar infrastructure in other contexts.

Step 1: Data Ingestion. In order to flexibly accommodate diverse data sources, with varying secure file transfer methods and update frequencies, we operate a server in the cloud that pulls updated data from each source on a regular interval. We receive data updates from private companies on a daily, weekly or monthly cadence. Many companies have unique policies and requirements for securing data transfers, so we write scripts to intake this data using a variety of secure file transfer services (e.g. Amazon S3 buckets and SFTP servers). We also download or scrape a variety of publicly available statistics from the web, such as unemployment insurance claims and COVID-19 case counts.

Three main challenges arise when handling this large volume of frequently updated data: storing, syncing, and version controlling the data we receive. We store all the raw data we receive as flat files in a data lake (an Amazon S3 bucket). We use object storage rather than a database or a more customized storage service (such as Git LFS) to minimize storage costs while maximizing our flexibility to ingest incoming data which arrives in numerous formats that may change over time. We version control each snapshot of the data we download within the same Git repository that stores our code using a tool called [DVC](#) (“Data Version Control”). DVC creates a pointer to a hash of the raw data for each data file or folder (in other words, a shortcut to the files in the data lake), which we version control in Git and update every time new data is downloaded. This associates each snapshot of data with the code that existed at the time it was processed, and allows us to easily roll back our code and data simultaneously to any prior state. DVC also facilitates syncing

the raw data from the data lake by efficiently downloading the data that is associated with each pointer in the Git repository.

Step 2: Data Processing. For each dataset, we have an automated pipeline of programs that process and transform the raw data into the public datasets that we post online. We use an automated build tool to organize and execute this collection of programs. We mostly process the data using Stata and execute our automated builds within Stata using the `-project-` command developed by Robert Picard.

This data processing step generates two outputs: (1) a set of CSV files that contain all the data to be posted publicly and (2) a quality control report. The quality control report is a document that allows analysts to quickly assess any notable deviations in the data and determine whether the updated data requires further review before being publicly released. Each report flags three types of changes that would require manual review: revisions made to previously posted data, large deviations in newly reported data, or newly missing data. The report also contains a series of tables and figures that preview the data and highlight any changes in the newly processed data.

Each time new data is ingested, the data processing step is run automatically. If it runs to completion, a Git pull request is generated with DVC pointers to the newly updated raw data alongside a link to the quality control report. If the data processing fails (for example, because the structure of the raw data has changed), an error report is generated. At this point, we pause and perform a manual review before posting the new data online. If the data processing failed or if any changes were detected in the quality control report that require further review, we manually investigate and write new code as needed, then re-process the data and inspect the updated quality control report before proceeding.

After reviewing and approving the quality control report, we merge the Git pull request containing the new data, which automatically triggers the final Data Release step. This manual review and approval is therefore the only manual step in the data processing pipeline.

Step 3: Data Release. Once the processed data is ready for release, our scripts automatically post the updated data to two public destinations. First, we sync the updated data into the database powering our online data visualization website built by DarkHorse Analytics (www.tracktherecovery.org). While doing so, we also update the “last updated” and “next expected update” dates on the website. Second, we upload the CSV files containing all the updated data to our “data downloads” page. The updated visualizations and data downloads are then both immediately available for public use.

B Consumer Spending Series Construction

This appendix provides greater detail on the construction of the consumer spending series in the Affinity Solutions data.

We receive data from Affinity Solutions in cells corresponding to the intersection of (i) county by (ii) income quartile by (iii) industry by (iv) week, where cells where fewer than five unique cards transacted are masked. Income quartile is assigned based on ZCTA of residence using 2014-2018 ACS estimates of median household income along with population size. We use population weights when defining quartile thresholds so that each income quartile has the same number of individuals. ZCTA income quartile and county are both determined by the cardholder’s residence.

We adjust the raw data we receive from Affinity Solutions to address two challenges: (1) changes in the customer base over time and (2) a data quality issue which creates spurious increase in consumer spending.

Changing Customer Base. One of the main practical challenges with this private-sector data source is that their customer base changes over time. These changes manifest themselves in the form of large, sudden changes in both the amount of spending per day and the number of unique cards transacting per day. We identify these sudden changes systematically by regressing the number of transacting cards on our date variable separately for each year-by-county, and then implementing a Supremum Wald test for a structural break at an unknown break point.

950 counties have a structural break where the p-value of the test is less than 5×10^{-8} . For counties with only one break below this threshold, we correct our estimate using chain weighting. For this procedure, we first compute the state-level percent change from week-to-week excluding all counties with a structural break (using the national series for DC and states for which all counties have a structural break). If we identify a structural break in week t , we impute spending levels in weeks $t-1$, t , and $t+1$ levels because we don’t know the precise date when the structural break occurred (e.g. it might have occurred on the 2nd day of $t-1$ or the 6th day of t). To keep our estimates conservative, when there is a change in coverage we adjust the series to be in line with the lower level of coverage. For example, say a county has n active cards, until in week t the county has $3n$ cards. In week $t-2$, the county would have a level of n cards, its reported value. In week $t-1$, if counties in the rest of the state had a 5% increase in the number of cards, we would impute the county with a break to have a level of $1.05n$ cards in week $t-1$. In week t , if counties in the rest of the state had a 10% increase in the number of cards, we would impute t to have a level of

$(1.10) * (1.05n) = 1.155n$. Likewise, if counties in the rest of the state had an 8% decrease in the number of cards in week $t+1$, we would impute $t+1$ to have a level of $(0.92) * (1.155n) = 1.0626n$. At this point, state-level fluctuations no longer impact the series, and we use the reported percent change each week to adjust this number for card coverage. For counties with multiple breaks, we omit the county from our series; this happens in 98 cases. Furthermore, we do not remove any counties where the structural break occurred between March 10th and March 31st of 2020 because the consumer spending response to the coronavirus epidemic was so strong that in many places it could be mistaken as a structural break. We also impute the number of cards by averaging the preceding and succeeding week for the week of Christmas, since holiday spending spikes are also sometimes mistaken for a structural break.

We implement a structural-break correction for three counties: Philadelphia County, Pennsylvania (county fips of 42101); Washington, District of Columbia (11001); Jefferson County, Kentucky (21111). For Philadelphia and Washington, we implement a correction by estimating a regression discontinuity at the date of the break, and then adding the RD estimate to the series prior to the structural break. The structural break in Jefferson county occurs on January 7th of 2020, and so there are not enough days on the left-hand side to implement the RD correction. Consequently, we assign the January 7th value to each day between January 1st and January 6th.

Spurious Increases in Consumer Spending. There is an unreasonably large spike in consumer spending between January 15th, 2019 and January 17th, 2019 that is not found in other data series, and so we believe it is not representative of true economic activity. This spike in national consumer spending is not driven by specific regions nor sectors. We deal with this data quality issue by replacing each impacted day with the average spending on $t - 7$, $t + 7$, and $t + 14$, where t is the impacted day. A similar problem arises in the “Accommodations and Food Services” sector in Richmond City County, Virginia where spending increases by over 80 times on May 23rd, 2019 relative to nearby days. We implement a similar procedure replacing the impacted day with the average spending on $t - 14$, $t - 7$, $t + 7$, and $t + 14$, where t is the impacted day.

C Small Business Revenue and Small Businesses Open Series Construction

This appendix provides further details on our construction of small business revenue and small businesses open series in Womply data.

Construction of Small Business Revenue and Small Businesses Open Series. We receive Wom-

ply data on total revenue and number of open businesses at the date x ZIP code x firm category level. We crosswalk from ZIP codes to counties using the geographic definitions described in Appendix F to aggregate the series to the county, state and national level. We then collapse the Womply data to aggregate spending and total small businesses open within each date x NAICS x geography x ZIP income quartile, creating ZIP income quartiles as described in Appendix B. We take a seven-day lookback moving average of each series, and norm its series relative to its level over the period January 4-31.

Masking. To preserve the privacy of firms in the data and to avoid displaying noisy estimates for small cells, we mask Womply series that report less than \$250,000 in total revenue during the base period of January 4-31.

Anomalous Data. Our quality-control process checks for anomalous variations in the Womply raw data. We detect two cases of large, sudden spikes in revenue, without an accompanying increase in small businesses open, driven by a single category in a single ZIP code. In Detroit, MI, “Healthcare and Medical Centers” in ZIP code 48150 increases by 34,000% on June 21 relative to the level one week prior on June 14. This leads to an anomalous revenue spike for the entire city of Detroit. Since this firm category makes up a small percentage of revenue (about 1% of Michigan’s revenue during the base period), we drop the revenue category in this specific ZIP code from our series. In Ohio, there is an increase in revenue driven heavily by “Retail and Clothing Revenue” in ZIP code 43125 and “Professional Services Revenue” in ZIP code 43201. Collectively, these two ZIP codes experience a revenue increase of more than 600% in late May that drives spikes in total revenue in both Columbus and Ohio (since these two ZIP codes represent more than 15% of the total Ohio revenue during the base period). We drop these two ZIP codes from our series and exclude Franklin County (i.e. Columbus) from our county and city level aggregation to prevent displaying a selected distribution of firm types in the small business revenue series data. However, we include the remaining ZIP codes in our state and national level aggregations of the Ohio revenue.

There are also several cases of single-day, downward spikes of negative revenue within a given firm category x ZIP code. We treat these cases as outliers, and replace these instances of negative revenue with zero revenue.⁵⁷

Delayed Processing of Payments. Due to differences in the speed at which data providers share

57. More generally, negative revenue may appear in the Womply data due to returns and refunds. There are a number of cases of observed negative revenue, especially during March 2019, due to consumers seeking returns or refunds on certain products. We include these cases in the Womply series, but exclude large single-day occurrences of negative revenue.

their data with Womply, the most recent date as of a given data refresh is typically incomplete. If left unaddressed, there would appear to be a decline in small business revenue and small businesses open in the most recent data. The full Womply data has generally been received by two days after it is first provided. As a conservative approach, we delay publishing Womply data for four days after it is first provided.

D Employment Series Construction

This appendix provides further details on how we construct various employment series analyzed in the paper.

Construction of Paychex Employment Series. We receive Paychex data at the county x industry x 2019 hourly wage quartile x 2019 firm size x pay frequency x week of payroll processing level. We first create a daily series of paychecks processed on each date by linearly interpolating daily values between each week in each county x 2-digit NAICS code x 2019 hourly wage quartile x 2019 firm size bin x pay frequency cell. In order to construct a series of employment as of each date, rather than paychecks being processed as of each date, we take two steps.

First, we construct a series of pay periods ending as of each date. We take a separate approach for paychecks following regular weekly cycles (i.e. weekly and bi-weekly paychecks) and for paychecks following a cycle based on fixed calendar dates (i.e. semi-monthly and monthly paychecks). For weekly and bi-weekly payfrequencies, we use data provided by Paychex on the distribution of (date at which payroll is processed – last date in pay period), for weekly and bi-weekly payfrequencies to distribute paychecks to the last date of the corresponding pay period, treating the distribution of (date at which payroll is processed – last date in pay period) as constant across geographies and NAICS codes. For monthly and semi-monthly payfrequencies, where cycles regularly occur on fixed calendar dates (e.g. the 15th and 30th of each month for semi-monthly paycycles), we assume that the last date within each pay period is the closest preceding calendar date that is the 15th or the 30th day of the month (semi-monthly paycycles) or the 30th day of the month (monthly paycycles). In each case, we interpolate values around public holidays

Second, to construct a series of employment as of each date, we record a worker as being employed for the full duration of the paycycle up until the last date in their pay period, under the assumption that workers are employed for each day during their pay period. We then collapse the data to the level of county x industry x 2019 hourly wage quartile x 2019 firm size x pay x date.

Finally, we take steps to prevent the introduction of new Paychex clients from artificially creating breaks in the employment series at smaller levels of geography. We begin by calculating the share of county employment in January 2020 accounted for by each industry x firm size bin within each county. Next, we calculate the change in employment relative to January for each county x industry x firm size bin, and multiply this change by the share of total employment in the respective counties, creating an employee-weighted employment series for each county x industry x firm size bin cell. We denote a county x industry x firm size bin cell as an “influential cell” if the county contains 40 or fewer than unique county x industry x firm size bin cells, and the cell accounts for over 10% of employment in the county at any date in 2020, or if the county contains greater than 40 unique county x industry x firm size bin cells, and the cell accounts for over 5% of employment in the county at any date in 2020. We drop influential cells that record a change in employment relative to January 2020 of at least +50% on any date, on the basis that such a trend likely arises due to changes in Paychex’s client base rather than true employment changes. Of the 34,018 unique county x industry x firm size bin cells in the Paychex data, fewer than 500 cells are affected by this procedure.

Construction of Earnin Employment Series. We obtain anonymized microdata at the worker level from Earnin. We construct our analysis sample by restricting the sample to workers who are paid on a weekly or bi-weekly paycycle; these categories account for 92% of paychecks. We also restrict the sample to workers who are active Earnin users, with non-missing earnings and hours worked over the last 28 days. Next, we exclude workers whose reported income over the prior 28 days is greater than \$50,000/13 (corresponding to an income of greater than \$50,000 annually).

We then restrict the sample to workers who are in paid employment. Users may continue to use Earnin after they have been laid off; we exclude payments which Earnin classifies as unemployment payments, either based on the user’s registration with Earnin as being unemployed, or based on the string description of the transaction. Where a user has previously been unemployed, but stops receiving unemployment checks after a certain date, we treat the user as having been re-employed if they receive a payment amount of \$200 within the two weeks following their last unemployment check. Using this approach, we find that 90% of Earnin users are re-employed within fourteen days of receiving their last unemployment check.

We use external data sources to gather further information on firm size and industry. To obtain information on industry, we use a custom-built crosswalk created by Digital Divide Data which contains NAICS codes for each employer in the Earnin data with more than ten Earnin

users. To obtain information on firm size, we crosswalk Earnin employers to ReferenceUSA data at the firm location level by spatially matching Earnin employers to ReferenceUSA firms. We begin by geocoding Earnin addresses to obtain latitudes and longitudes for each Earnin employer. We then remove common prefixes and suffixes of firm names, such as “inc” and “associated”. Next, we compute the trigram similarities between firm names for all Earnin and ReferenceUSA firms within twenty-five miles of another. We then select one “match” for each Earnin firm within the ReferenceUSA data, among the subset of firms within one mile. We first match Earnin employers to ReferenceUSA firms if the firms are within one mile of one another, and share the same firm name. Second, where no such match is available, we choose the geographically closest firm (up to a distance of one mile) among all firms with string similarities of over 0.6. Third, where no such match is available, we match an Earnin employer to the ReferenceUSA employer within twenty-five miles with the highest trigram string similarity, provided that the employer has a trigram string similarity of 0.9. We then compute the modal parent-firm match in the ReferenceUSA data for each parent-firm grouping in Earnin. Where at least 80% of locations within a parent-firm grouping in Earnin are matched to a single parent-firm grouping in the ReferenceUSA data, we impute that parent-firm to every Earnin location. In total, we match around 70% of Earnin employers to ReferenceUSA firms.

Earnin data are observed at the ZIP code level. We crosswalk from ZIP Codes to counties using the geographic definitions described in Appendix F to aggregate the series to the county, state and national level.

We construct an employment series in the Earnin data from our analysis sample as follows. In the paycheck-level data, we observe the worker’s paycycle frequency. We use paycycle frequency to construct an employment series by assuming that workers are employed throughout the full duration of their paycycle. That is, we assume that a worker paid every two weeks has been fully employed for the two weeks prior to receiving their paycheck. To account for the delay in receipt of paychecks, we shift the Earnin series back by one week. We then take the count of employed individuals across the Earnin sample as our measure of employment. We take a 7-day moving average to form our Earnin employment series, and express the series as a change relative to January 4-31.

We also receive transaction-level data on all payments received, which we use to measure the receipt of stimulus checks. We classify a transaction as a stimulus check if the transaction (1) has a string description containing words indicating that it is a stimulus transaction, such as “IRS” or “Economic Impact Payment”, (2) is of an amount that could be received as stimulus under the

CARES Act, and (3) was received after 10 April 2020.

Comparison of Construction of Earnin and Paychex Employment Series to Cajner et al. (2020).

In both the Earnin and Paychex datasets, we construct daily employment series using data on paychecks. Our treatment of paycheck data is similar to the treatment of paycheck data in Cajner et al. (2020), who estimate employment based on paycheck deposits using firm-level data from ADP. Cajner et al. (2020) define employment within a week as the count of paychecks that are processed during that week. For businesses which do not process payroll every week (e.g. businesses whose workers are paid every two weeks), Cajner et al. (2020) impute the count of paychecks in the “missing” week using the number of paychecks in the next period in which the businesses processes payroll.

Because the Earnin data are available at the worker level, we do not observe whether a business as a whole does not process payroll every week. However, under the assumption that all workers within a business are paid on the same paycycle, our worker-level approach of distributing paychecks uniformly over the paycycle matches the approach in Cajner et al. (2020) of imputing employment based on the next week in which paychecks are observed. The two primary differences between our treatment of paycycles and the treatment in Cajner et al. (2020) are that we use a 7-day moving average, whereas Cajner et al. (2020) use a 14-day moving average, and that we treat that the last date of the employment period as seven days prior to the receipt of the paycheck, whereas Cajner et al. (2020) observe the pay period directly. The seven-day lag accounts for delays between the end of a worker’s pay period, which is the event observed in Cajner et al. (2020), and the date on which paychecks are received by workers, which is the event observed in the Earnin data.

Because the Paychex data are not available at the firm level, we are not able to directly implement the approach in Cajner et al. (2020) of imputing employment using the count of paychecks in the “missing” week for firms that do not process payroll on a weekly basis. Instead, we make the conceptually similar assumption that workers are employed throughout the full duration of their paycycle, such that we can infer the full set of dates on which an individual worked by observing the last date of each of their pay periods and their pay frequency. Under the assumptions that all workers within a given firm are paid according to the same paycycle, our approach of inferring employment based on last date of pay period matches the approach in Cajner et al. (2020) of imputing employment based on the next week in which paychecks are observed. A further difference is that pay period is observed in Cajner et al. (2020); by contrast, in the Paychex data, pay periods are imputed using payroll processing date and the distribution of (payroll processing date – last date

in pay period). Finally, Cajner et al. (2020) use a 14-day moving average, whereas we use a 7-day moving average.

Construction of ZIP Code-Level Low-Income Employment Series. As ZIP code is not observed in Paychex and Intuit, we separately construct ZIP code-level employment using the Earnin data only. We construct our analysis sample as above. To account for the noisier data at the ZIP code-level, we norm the ZIP code-level changes relative to a pre-period of January 5 - March 7. We suppress estimates for ZIP codes with fewer than 100 worker-days observed over this period.

Construction of Paycheck Protection Program Analysis Series. To construct the series that we use to analyze the effects of the Paycheck Protection Program, we first collapse the Earnin and Paychex datasets down to the County x 2-digit NAICS x Income Quartile x Eligible (i.e., baseline employment less than 500) x Week level. For the combined Earnin-Paychex series, since Earnin is only representative of Q1 workers, we downweight Q1 employment in Paychex and Earnin such that the share of Q1 employment when summing Paychex and Earnin is equal to the share of Q1 employment in Paychex. Within each dataset, we then reweight on NAICS codes such that the worker-weighted distribution of NAICS codes within eligible vs. ineligible firms in January 2020 matches the worker-weighted pooled distribution of NAICS codes in 2020 in that dataset. For the combined dataset, we reweight each employment variable (i.e. Earnin and Paychex) on NAICS codes, then reweight on datasets within eligible vs. ineligible, such that the share of Earnin vs. Paychex within eligible vs. ineligible firms in January 2020 matches the worker-weighted pooled distribution of NAICS codes in 2020 in that dataset.

Assessing Mismeasurement of Firm Sizes using SBA data. We assess the degree of misclassification of PPP eligibility in our sample by merging publicly available data on PPP recipients from the SBA to data on firm sizes from ReferenceUSA, which form our measure of firm size in the Earnin data. To construct SBA data on PPP recipients, we restrict our attention to firms receiving loans of at least \$150,000; the names and addresses of these firms are publicly available from the SBA. We first geocode addresses recorded in SBA and ReferenceUSA data to obtain a latitude and longitude for each firm. We then compute the trigram similarities between firm names for all SBA and ReferenceUSA firms within twenty-five miles of another. We then select one “match” for each PPP recipient within the ReferenceUSA data, among the subset of firms within twenty-five miles, following the procedure described above in our merge of Earnin data to ReferenceUSA data. For firms with loans of above \$150,000, exact loan size is not observed; we impute loan size as the midpoint of loan range.

We use the SBA data on PPP recipients to estimate the first-stage of our difference-in-differences design, i.e. how much more PPP assistance firms classified as having 100-500 employees in our sample received relative to those classified as having more than 500-800 employees.⁵⁸ The SBA released firm names and ZIP codes of PPP recipients receiving over \$150,000 in loans, which represent 72.8% of total PPP expenditure. Of the roughly 660,000 PPP recipients of these loans, we merge around 69% of firms and 74% of total expenditure to firm size data from ReferenceUSA.⁵⁹ In this matched subset, we find that mean PPP expenditure per worker is \$2,150 for firms we classify as having 100-500 employees and \$1,000 per worker for firms with 500-800 employees (excluding firms in the food services industry). Given that we match only 74% of the publicly available PPP expenditure to our data and the publicly available data covers only 73% of total PPP expenditure, this implies that firms measured as having 100-500 employees in our sample received $\frac{\$2,150}{0.74 \times 0.73} = \$3,990$ of PPP assistance per worker, while firms with 500-800 employees received $\frac{\$1,000}{0.74 \times 0.73} = \$1,860$ in PPP assistance per worker.⁶⁰ Using data released by the SBA on loans given to all firms (including firms receiving less than \$150,000), we calculate that PPP assistance to eligible firms with between 100 and 500 employees (excluding NAICS 72) is \$4,080 per worker on average.⁶¹ Hence, firms with 501-800 workers in the ReferenceUSA data (the control group) were effectively treated at an intensity of $\frac{\$1,860}{\$4,080} = 45.4\%$, whereas firms with 100-500 workers in the ReferenceUSA data (the treatment group) were treated at an intensity of $\frac{\$3,990}{\$4,080} = 97.7\%$. Thus, inflating our baseline reduced-form estimates by $\frac{1}{0.977 - 0.454} = 1.92$ yields estimates of the treatment effect of PPP eligibility adjusted for attenuation bias due to mismeasurement of firm size.

58. We assess our first-stage using data on firm size from ReferenceUSA, which is the source for firm sizes in the Earnin data. We assume that the degree of misclassification in the merged ReferenceUSA-Earnin data matches the degree of misclassification in Paychex data, where firm sizes are computed internally using employment in 2019.

59. Precise loan amounts are released for loans of under \$150,000, whereas loan ranges are released for loans over \$150,000. For loans where only a loan range is available, we impute loan amount using the midpoint of minimum and maximum of loan range.

60. This calculation assumes that the degree of misclassification of eligibility among identifiable PPP recipients matches the degree of misclassification of eligibility in the broader ReferenceUSA sample.

61. To compute this statistic, we first calculate the share of total loan amounts received by non-NAICS 72 firms in the publicly released SBA data. We begin by imputing precise loan amount as the midpoint of minimum and maximum of loan range, where precise loan amount is not released. We then calculate the share of loans in firms with firm size between 100 and 500, in NAICS codes other than NAICS 72, under the assumption that our merge rate is constant by firm size. Using this approach, we calculate that 13.1% of PPP loan spending was allocated to non-NAICS 72 firms with between 100 and 500 employees. We then rescale the total PPP expenditure to the end of June, \$521 billion, by 0.131 to arrive at an estimate of \$69.22 billion in PPP loan spending to non-NAICS 72 firms with 100-500 employees. Finally, we divide \$69.22 billion by the number of workers at non-NAICS 72 firms with 100-500 employees to arrive at an estimate of loan spending per worker.

E Public Data Sources

This appendix provides further details on our use of public data sources on unemployment benefits, COVID-19 incidence, and mobility measures.

Unemployment Benefit Claims. We collect county-level data by week on unemployment insurance claims starting in January 2020 from state government agencies since no weekly, county-level national data exist. Location is defined as the county where the filer resides. We use the initial claims reported by states, which sometimes vary in their exact definitions (e.g., including or excluding certain federal programs). In some cases, states only publish monthly data. For these cases, we impute the weekly values from the monthly values using the distribution of the weekly state claims data from the Department of Labor (described below). We construct an unemployment claims rate by dividing the total number of claims filed by the 2019 Bureau of Labor Statistics [labor force](#) estimates. Note that county-level data are available for 22 states, including the District of Columbia.

We also report weekly unemployment insurance claims at the state level from the Office of Unemployment Insurance at the Department of Labor. Here, location is defined as the state liable for the benefits payment, regardless of the filer’s residence. We report both new unemployment claims and total employment claims. Total claims are the count of new claims plus the count of people receiving unemployment insurance benefits in the same period of eligibility as when they last received the benefits.

COVID-19 Data. We report the number of new COVID-19 cases and deaths each day using publicly available data from the New York Times available at the county, state and national level.⁶² We also report daily state-level data on the number of tests performed per day per 100,000 people from the [COVID Tracking Project](#).⁶³ For each measure - cases, deaths, and tests – we report two daily series per 100,000 people: a seven-day moving average of new daily totals and a cumulative total through the given date.

Google Mobility Reports. We use [data](#) from Google’s COVID-19 Community Mobility Reports to construct measures of daily time spent at parks, retail and recreation, grocery, transit locations, and workplaces.⁶⁴ We report these values as changes relative to the median value for the corresponding

62. See the New York Times data [description](#) for a complete discussion of methodology and definitions. Because the New York Times groups all New York City counties as one entity, we instead use case and death data from New York City Department of [Health data](#) for counties in New York City.

63. We use the Census Bureau’s 2019 population estimates to define population when normalizing by 100,000 people. We suppress data where new counts are negative due to adjustments in official statistics.

64. Google Mobility trends may not precisely reflect time spent at locations, but rather “show how visits and length

day of the week during the five-week period from January 3rd - February 6, 2020. Details on place types and additional information about data collection is available from [Google](#). We use these raw series to form a measure of time spent outside home as follows. We first use the American Time Use survey to measure the mean time spent inside home (excluding time asleep) and outside home in January 2018 for each day of the week. We then multiply time spent inside home in January with Google’s percent change in time spent at residential locations to get an estimate of time spent inside the home for each date. The remainder of waking hours in the day provides an estimate for time spent outside the home, which we report as changes relative to the mean values for the corresponding day of the week in January 2018.

F Key Dates and Geographic Definitions

In this appendix, we provide additional details about how we define key dates and geographic units used in our analysis.

Key Dates for COVID-19 Crisis. The Economic Tracker includes information about key dates relevant for understanding the impacts of the COVID-19 crisis. At the national level, we focus on three key dates:

- First U.S. COVID-19 Case: 1/20/2020
- National Emergency Declared: 3/13/2020
- CARES Act Signed in to Law: 3/27/2020

At the state level we collect information on the following events:

- Schools closed statewide: Sourced from COVID-19 Impact: School Status Updates by MCH Strategic Data, available [here](#). Compiled from public federal, state and local school information and media updates.
- Nonessential businesses closed: Sourced from the Institute for Health Metrics and Evaluation state-level data (available [here](#)), who define a non-essential business closure order as: "Only locally defined 'essential services' are in operation. Typically, this results in closure of public spaces such as stadiums, cinemas, shopping malls, museums, and playgrounds. It also includes restrictions on bars and restaurants (they may provide take-away and delivery

of stay at different places change compared to a baseline.” We call this “time spent at a location” for brevity.

services only), closure of general retail stores, and services (like nail salons, hair salons, and barber shops) where appropriate social distancing measures are not practical. There is an enforceable consequence for non-compliance such as fines or prosecution."

- Stay-at-home order goes into effect: Sourced and verified from the New York Times reopening data, available [here](#), and hand-collection from local news and government sources where needed.
- Stay-at-home order ends: Sourced and verified from the New York Times reopening data, available [here](#), and hand-collection from local news and government sources where needed. Defined as the date at which the state government lifted or eased executive action or other policies instructing residents to stay home. We code "regional" and "statewide" expiry of stay-at-home orders separately. A "regional" expiration of a stay-at-home orders occurs when a stay-at-home order expires in one region within a state, but not everywhere within the state. A "statewide" expiration of a stay-at-home order occurs when a stay-at-home order first expired throughout a whole state, either due to a statewide change in policy, or due to the stay-at-home order in each county having expired.
- Partial business reopening: Sourced and verified from the New York Times reopening data, available [here](#), and hand-collection from local news and government sources where needed. Defined as the date at which the state government allowed the first set of major industries to reopen (non-essential retail or manufacturing in nearly every case). Deviations from the New York Times reopening data are deliberate and usually involve our regional classification or our inclusion of manufacturing. A "regional" reopening occurs when businesses are allowed to reopen in one region within a state, but not everywhere within the state. A "statewide" reopening occurs when businesses are allowed to reopen throughout a whole state, either due to a statewide change in policy, or due to restrictions being eased in each individual county.

Geographic Definitions. For many of the series we convert from counties to metros and ZIP codes to counties. We use the HUD-USPS ZIP code Crosswalk Files to convert from ZIP code to county. When a ZIP code corresponds to multiple counties, we assign the entity to the county with the highest business ratio, as defined by HUD-USPS ZIP Crosswalk. We generate metro values for a selection of large cities using a custom metro-county crosswalk, available in Appendix Table 7. We assigned metros to counties and ensured that a significant portion of the county population was

in the metro of interest. Some large metros share a county, in this case the smaller metro was subsumed into the larger metro. We use the Uniform Data Systems (UDS) Mapper to crosswalk from ZIP codes to ZCTAs.

Table 1
Changes in Consumer Spending by Sector and Income Quartile

| Dep. Var.: | Change in Mean Consumer Spending Per Day (\$ Billions) Relative to 2019 Level | | | Level of Mean Consumer Spending Per Day (\$ Billions) |
|---|---|------------------------|------------------------|---|
| | Change as of April 8-14 | Change as of June 8-14 | Change as of July 8-14 | Level as of January 4-31 2020 |
| | (1) | (2) | (3) | (4) |
| Panel A: Consumer Spending by Income Quartile | | | | |
| Pooled, All Income Quartiles | -7.84 | -3.09 | -2.86 | 22.05 |
| Low-Income | -1.00 (12.77%) | -0.21 (6.92%) | -0.24 (8.48%) | 3.35 (15.18%) |
| Q2 | -1.60 (20.34%) | -0.53 (17.18%) | -0.56 (19.57%) | 4.94 (22.41%) |
| Q3 | -2.10 (26.75%) | -0.83 (26.81%) | -0.76 (26.58%) | 5.97 (27.07%) |
| High-Income | -3.15 (40.14%) | -1.52 (49.09%) | -1.30 (45.38%) | 7.79 (35.33%) |
| Panel B: Consumer Spending by Sector | | | | |
| | Overall Sector Decomposition | | | |
| Durable Goods | -0.83 (10.62%) | 0.96 (-42.46%) | 0.89 (-45.64%) | 4.94 (22.65%) |
| Non-Durable Goods | -0.63 (8.02%) | 0.07 (-2.96%) | 0.19 (-9.86%) | 4.86 (22.28%) |
| Remote Services | -1.18 (15.00%) | -0.16 (6.90%) | -0.22 (11.48%) | 4.45 (20.41%) |
| In-Person Services | -5.11 (65.12%) | -2.88 (127.34%) | -3.03 (156.11%) | 6.94 (31.79%) |
| | In-Person Services Sub-Sector Decomposition | | | |
| Hotels & Food | -1.96 (38.34%) | -1.14 (39.56%) | -1.17 (38.64%) | 2.61 (37.59%) |
| Transportation | -1.50 (29.42%) | -1.07 (37.33%) | -1.10 (36.17%) | 1.74 (25.02%) |
| Health Care | -0.51 (9.96%) | -0.11 (3.89%) | -0.15 (4.84%) | 0.85 (12.30%) |
| Recreation | -0.43 (8.35%) | -0.32 (11.26%) | -0.33 (10.94%) | 0.51 (7.28%) |
| Other In-Person Services | -0.71 (13.92%) | -0.23 (7.97%) | -0.29 (9.41%) | 1.23 (17.80%) |

Notes: This table shows change in and levels of national mean daily spending between 2019 and 2020 for selected dates, ZCTA income quartiles, and categories of goods. Panel A shows changes in and levels of consumer spending by income quartile, in the Affinity Solutions data. Column (1) shows the change in consumer spending as of the week April 8-14. The first row shows the change in national spending on all goods. To compute this change, we begin by calculating total daily spending in the Affinity Solutions data for each day in 2019 and 2020. We then scale the 2020 (2019) values of daily total spending by multiplying by the ratio of January 2020 total spending for components of PCE that are likely captured in credit/debit card spending (shown in the last bar of Figure 1a) to the January 2020 (2019) total spending in the Affinity data. We then calculate the change in total spending between 2019 and 2020 for the period April 8-14 as ((Spending in April 8 through April 14 2020) - (Spending in April 8 through April 14 2019)) - ((Spending in January 4 through January 31 2020) - (Spending in January 4 - January 31 2019)). The second, third, fourth and fifth rows of Panel A replicate the first row, restricting to ZCTAs in a given income quartile. The decline in spending in each income quartile is expressed in percentage terms as a share of the national decline in brackets underneath each row. Columns (2) and (3) replicate column (1), calculating the change as of June 8-14 (column (2)) and July 8-14 (column (3)) respectively. Column (4) shows mean daily national spending over the period January 4-31 2020 for each income quartile. Panel B replicates Panel A in the Affinity data for various categories of goods instead of income quartiles. The first four rows of Panel B show the change in consumer spending across four broad categories of goods: durable goods (row 1), non-durable goods (row 2), remote services (row 3) and in-person services (row 4); for details of the definitions of these categories, see notes to Figure 2. The change in consumer spending summed across these four categories may not add to 100% because of spending on uncategorized goods. Rows 5-9 show the change in consumer spending within five components of in-person services: hotels and food (row 5), transportation (row 6), health care (row 7), recreation (row 8), and other in-person services (row 9).

Table 2
Changes in Business Revenue and Employment in Small Businesses

| <i>Panel A: Changes in Business Revenue</i> | | | | | |
|--|---|-----------------|------------------|------------------|-----------------|
| Dep. Var.: | % Change in Changes in Small Business Revenue | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Median 2BR Rent (per thousand dollars) | -13.00 (0.38) | -7.69 (0.51) | -14.07 (0.70) | -10.99 (0.73) | -0.06 (0.01) |
| Log of Density of High Wage Workers | | -1.16 (0.08) | | -2.11 (0.13) | -1.30 (0.10) |
| County FEs | | | X | X | |
| State FEs | | | | | X |
| Level of Observation | ZIP code | ZIP code | ZIP code | ZIP code | County |
| Observations | 16396 | 15907 | 16396 | 15907 | 2759 |
| <i>Panel B: Changes in Low-Wage Employment</i> | | | | | |
| Dep. Var.: | % Change in Changes in Low-Wage Employment | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Median 2BR Rent (per thousand dollars) | -12.73 (0.53) | -9.56 (0.71) | -11.58 (0.97) | -10.00 (1.02) | -9.97 (1.56) |
| Log of Density of High Wage Workers | | -0.88 (0.11) | | -1.25 (0.19) | -1.07 (0.17) |
| County FEs | | | X | X | |
| State FEs | | | | | X |
| Level of Observation | ZIP code | ZIP code | ZIP code | ZIP code | County |
| Observations | 13440 | 12885 | 13440 | 12885 | 2108 |

Notes: This table shows OLS regressions of average percentage changes in small business revenue by ZIP code (using Womply data) and small business low-wage employment (using Earnin data) on average ZIP code median two-bedroom rent. Standard errors are reported in parentheses. The dependent variable is scaled from 0 to 100, such that, for example, the coefficient of -13.00 in Column (1), Panel A implies that a \$1000 increase in monthly two-bedroom rent is associated with a 13% larger drop in total revenue. Panel A estimated the changes in small business revenue while Panel B does the same for changes in small businesses low-wage employment. In both cases Columns (1)-(4) are estimated at the ZIP code level, Column (1) shows the baseline regression without any controls while Columns (2)-(4) add county fixed effects and the log of the density of high wage workers. Column (5) is estimated at the county level, with State fixed effects and controlling for the log of the density of high wage workers.

Table 3
Causal Effects of Re-Openings on Economic Activity: Event Studies

| Outcome: | Spending (%) | | Employment (%) | | Low-Wage Employment (%) | High-Wage Employment (%) | Merchants Open (%) | | Time Away From Home (%) | |
|---|----------------|----------------|----------------|----------------|-------------------------|--------------------------|--------------------|----------------|-------------------------|----------------|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| Dif-in-Dif Effect of Reopening: | 1.04 (0.93) | 0.18 (1.34) | 1.44 (0.68) | 2.66 (1.05) | 0.59 (0.38) | 1.74 (1.4) | 3.67 (1.35) | 5.01 (1.77) | 0.76 (0.56) | 0.92 (0.98) |
| State-Week | 208 | 318 | 200 | 264 | 248 | 248 | 248 | 324 | 112 | 138 |
| Weeks on either side of reopening included: | 2 | 3 | 2 | 3 | 2 | 2 | 2 | 3 | 2 | 3 |
| Decline (Peak to Trough): | 33.6 | | 19.4 | | 31.0 | 9.7 | 36.3 | | 23.5 | |

Notes: This table displays changes to outcomes as a result of reopenings using an event study design looking at states that reopened non-essential businesses between April 20 and May 4. Each state that reopens is matched to multiple controls states that did not reopen but had similar trends of the outcome variable during the weeks preceding the reopening. Standard errors are reported in parentheses. Columns (1) - (2) look at changes in consumer spending using seasonally-adjusted data from Affinity Solutions. Consumer spending is normalized by its level over the period January 4-31, and seasonally adjusted using 2019 data, as described in Section II.A. Columns (3) - (4) look at employment data from Womply and Paychex (see Appendix D for details). Columns (5) and (6) restrict this to workers in the bottom and top quartile of earnings respectively. Columns (7) and (8) look at merchants that reopen, using data from Womply. Columns (9) and (10) look at time spend away from home, using data from Google. Columns (2), (4), (8) and (10) estimate the coefficient on the reopening using the three weeks preceding and following each reopening; all other columns use two weeks instead. Columns (1), (3), and (7) correspond to the specifications displayed in Figures 12B, 12C, and 12D respectively.

Table 4
Regression Discontinuity Estimates of Stimulus Payments on Spending

| | (1) Q1 ZIP codes | (2) Q1 ZIP codes | (3) Q4 ZIP codes | (4) Q4 ZIP codes |
|--|---------------------|---------------------|---------------------|---------------------|
| <i>Panel A: Impact of Stimulus Payments on Consumer Spending</i> | | | | |
| Dep. Var.: _____ | Spending | | | |
| RD Effect of Stimulus: | 0.26 (0.07) | 0.38 (0.10) | 0.09 (0.04) | 0.16 (0.05) |
| Window: | April 1 - April 30 | April 7 - April 21 | April 1 - April 30 | April 7 - April 21 |
| <i>Panel B: Impact of Stimulus Payments on Revenue</i> | | | | |
| Dep. Var.: _____ | Revenue | | | |
| RD Effect of Stimulus: | 0.18 (0.10) | 0.21 (0.17) | 0.01 (0.06) | -0.08 (0.10) |
| Window: | April 1 - April 30 | April 7 - April 21 | April 1 - April 30 | April 7 - April 21 |

Notes: This table shows regressions of changes to outcomes as a result of stimulus payments using a regression discontinuity design around the April 14 payment date. Standard errors are reported in parentheses. Panel A looks at changes in consumer spending using seasonally-adjusted data from Affinity Solutions (where the seasonal adjustment consists of dividing spending on each day by the average level of spending in January and then residualised by day of week and first of the month fixed effects). We estimate the fixed effects using data from January 1, 2019, to May 10, 2019. Columns (1)-(2) looks at changes to spending for cardholders living in ZIP codes in the bottom quartile of the distribution of ZIP code median household income (based on data from the 2014-2018 ACS), while Columns (3)-(4) look cardholders living in the quartile of ZIP codes with the highest median incomes. This panel corresponds to the specifications displayed in Figures 12B and 12C. Panel B looks at changes in consumer spending using seasonally-adjusted data from Womply (cleaned the same way as the Affinity solutions data). Columns (1)-(2) looks at changes to business revenue for ZIP codes in the bottom quartile of the distribution of ZIP code median rent for a two bedroom apartment (based on data from the 2014-2018 ACS), while Columns (3)-(4) look at businesses in the quartile of ZIP codes with the highest median two bedroom rents. This panel corresponds to the specifications displayed in Figures 13A and 13B. In both panels, Columns (1) and (3) include all of April 2020 in the regression specification, while Columns (2) and (4) restrict to only considering dates within one week of the stimulus payment date.

Table 5
Estimated Effects of the Paychex Protection Program on Employment

| | Outcome Variable: Employment | | | |
|-------------|---|---|---|---|
| | Combined Paychex and Earnin Data | | Earnin Data | Kronos Data |
| | Baseline Estimate (100-800 Employees) (1) | Small Bandwidth (300-700 Employees) (2) | Baseline Estimate (100-800 Employees) (3) | Baseline Estimate (100-800 Employees) (4) |
| DD Estimate | 1.05 (1.60) | -0.78 (2.03) | 1.08 (0.99) | -0.03 (2.12) |

Notes: This table shows difference-in-difference estimates of the effect of PPP eligibility on employment. The outcome variable is employment at the county x 2-digit NAICS x income quartile x eligibility x week level, expressed as a percentage change relative to a pre-period of January 11-31 2020. Columns (1)-(2) present regressions in combined Paychex-Earnin data. In the baseline estimate in column (1), we begin by restricting to firms with between 100 and 800 employees. We then reweight firms on 2-digit NAICS codes such that the (worker-weighted) distribution of 2-digit NAICS codes within eligible Paychex (Earnin) firms matches the national distribution of 2-digit NAICS codes among Paychex (Earnin) firms in the period January 11-31 2020. Next, to combine the datasets, we reweight such that the (worker-weighted) share of each dataset is constant in eligible vs. ineligible firms. We then sum employment across datasets at the county x 2-digit NAICS x income quartile x eligibility x week level. We then regress change in employment on PPP eligibility, county x worker income quartile x week fixed effects, and an interaction term for PPP eligibility and the date being after April 3, clustering on county x industry. The DD estimate presents the coefficient and standard error on the interaction term for PPP eligibility and the date being after April 3. Column (2) replicates Column (1), restricting to firms with between 300 and 700 employees. Column (3) replicates Column (1) in Earnin data. As we treat all Earnin workers as belonging to the first quartile, we use county x week FEs, rather than county x worker income quartile x week FEs. Column (4) replicates Column (3) in Kronos data.

Appendix Table 1
Hourly Wage Rates By Industry

| NAICS Code | NAICS Description | 2019 BLS Wages | | | Median in Private Datasets | | | |
|------------|--|-------------------------------------|-------------------------------------|----------------------------|-----------------------------|------------------------------|-----------------------------|----------------------------|
| | | 10th Percentile (Pre Tax) (1) | 25th Percentile (Pre Tax) (2) | Median (Pre Tax) (3) | Earnin (Post Tax) (4) | Homebase (Pre Tax) (5) | Paychex (Pre Tax) (6) | Intuit (Pre Tax) (7) |
| 22 | Utilities | 18.56 | 26.82 | 38.06 | 15.00 | | 32.15 | |
| 55 | Management of Companies and Enterprises | 16.09 | 22.42 | 34.74 | 12.34 | | 21.34 | |
| 54 | Professional, Scientific, and Technical Services | 14.85 | 21.62 | 34.00 | 12.63 | 13.00 | 35.36 | |
| 51 | Information | 12.90 | 19.56 | 32.13 | 12.49 | | 33.48 | |
| 52 | Finance and Insurance | 14.25 | 18.40 | 27.42 | 12.77 | | 33.22 | |
| 21 | Mining, Quarrying, and Oil and Gas Extraction | 15.36 | 19.11 | 25.82 | 15.69 | | 32.76 | |
| 61 | Educational Services | 11.54 | 16.18 | 24.47 | 13.25 | 11.50 | 25.02 | |
| 23 | Construction | 13.78 | 17.51 | 23.92 | 13.94 | | 28.30 | |
| 42 | Wholesale Trade | 12.30 | 15.73 | 22.05 | 11.79 | | 28.37 | |
| 48-49 | Transportation and Warehousing | 12.07 | 15.49 | 20.89 | 13.20 | 15.00 | 24.83 | |
| 31-33 | Manufacturing | 12.36 | 15.35 | 20.77 | 12.66 | | 26.20 | |
| 53 | Real Estate and Rental and Leasing | 11.31 | 14.14 | 19.31 | 12.64 | | 26.13 | |
| 62 | Health Care and Social Assistance | 11.18 | 13.59 | 19.27 | 11.68 | 14.00 | 26.42 | |
| 81 | Other Services (except Public Administration) | 9.73 | 12.02 | 16.57 | 10.97 | 14.00 | 21.78 | |
| 56 | Administrative Support | 10.33 | 12.26 | 15.71 | 11.82 | | 24.06 | |
| 71 | Arts, Entertainment, and Recreation | 9.21 | 11.17 | 14.09 | 10.38 | 12.00 | 21.15 | |
| 11 | Agriculture, Forestry, Fishing and Hunting | 11.28 | 11.89 | 13.38 | 11.56 | | 20.37 | |
| 44-45 | Retail Trade | 9.49 | 11.18 | 13.36 | 9.76 | 12.00 | 21.15 | |
| 72 | Accommodation and Food Services | 8.68 | 9.61 | 11.81 | 9.26 | 11.00 | 16.32 | |
| | All | | | 18.25 | 12.34 | 12.00 | 25.69 | 25.68 |
| | Industry-Weighted Average of BLS Median Wages | | | | 19.22 | 14.23 | 20.46 | 22.17 |

Notes: This table reports wages at various percentiles for two-digit NAICS sectors. 2019 BLS Wages (1-3) come from the May 2019 Occupational Employment Statistics and are inflated to 2020 dollars using the Consumer Price Index. Columns (4-7) report median wages in four private employment datasets, Earnin, Homebase, Paychex, and Intuit. In Earnin and Homebase, the median wage is the 50th percentile of hourly wages for workers of the given industry during the pre-COVID period (January 8th - March 10th). In Paychex, the median wage is the 50th percentile of county-level average hourly wages for the given industry during the pre-COVID period (January 8th - March 10th). In Intuit, the median wage is the 50th percentile of county-level average hourly wages during the pre-COVID period (January - February). In Earnin (4), wages are calculated by dividing the payment deposited in the individual's bank account by hours worked and are thus post-tax. Homebase, Paychex, and Intuit wages are pre-tax. Industries missing from the Homebase data are left blank. The industry-weighted average of BLS median wages are calculated using the industry shares for the relevant dataset.

Appendix Table 2
Demographic Characteristics of Zearn Users

| | Zearn Users (1) | US Population (2) |
|-------------------------------------|--------------------|----------------------|
| <i>Panel A: Income</i> | | |
| ZIP Median Household Income | | |
| 25th Percentile | 43,766 | 45,655 |
| Median | 54,516 | 57,869 |
| 75th Percentile | 70,198 | 77,014 |
| Number of ZIP codes | 5,148 | 33,253 |
| Number of People | 803,794 | 322,586,624 |
| <i>Panel B: School Demographics</i> | | |
| Share of Black Students | | |
| 25th Percentile | 1.4% | 1.5% |
| Median | 5.6% | 5.8% |
| 75th Percentile | 21.3% | 19.1% |
| Share of Hispanic Students | | |
| 25th Percentile | 4.3% | 5.6% |
| Median | 10.9% | 15.0% |
| 75th Percentile | 35.7% | 40.6% |
| Share of Students Receiving FRPL | | |
| 25th Percentile | 33.8% | 28.2% |
| Median | 55.5% | 50.1% |
| 75th Percentile | 78.5% | 74.8% |
| Number of Schools | 8,801 | 88,459 |
| Number of Students | 767,310 | 49,038,524 |

Notes: This table reports demographic characteristics for US schools. Household income percentiles are calculated using the 2017 median household income in each school's ZIP code. The share of students who are Black, Hispanic, or receive Free or Reduced Price Lunch (FRPL) in a given school are calculated using school demographic data from the Common Core data set from MDR Education, a private education data firm. Percentile distributions for each demographic variable are calculated separately and weighted by the number of students in each school. Column (1) reports school characteristics for students using Zearn, while Column (2) reports income data for the entire US population and shares of students who are Black, Hispanic, or receive FRPL for all US elementary school students.

Appendix Table 3
Cities with Largest Small Business Revenue Losses Following COVID Shock

| City | State | % Change in Small Bus. Revenue (Womply) | % Change in Low-Wage Worker Hours, Small Restaurants/Retail (HomeBase) | % Change in Low- Wage Worker Hours (Earnin) |
|---------------|----------------------|---|--|---|
| (1) | (2) | (3) | (4) | (5) |
| New Orleans | Louisiana | -80.8% | -76.6% | -60.9% |
| Washington | District of Columbia | -72.9% | -73.2% | -60.2% |
| Honolulu | Hawaii | -62.7% | -75.8% | -25.3% |
| Miami | Florida | -62.2% | -68.7% | -51.1% |
| Boston | Massachusetts | -60.6% | -79.5% | -60.9% |
| Philadelphia | Pennsylvania | -58.7% | -66.6% | -51.8% |
| Fresno | California | -58.7% | -60.7% | -36.6% |
| San Jose | California | -58.6% | -61.5% | -51.9% |
| New York City | New York | -57.0% | -78.7% | -63.4% |
| Las Vegas | Nevada | -56.1% | -66.4% | -53.0% |

Notes: This table shows the ten cities with the largest small business revenue declines as measured in the Womply data (among the fifty largest cities in the U.S.). The decline is defined as net revenue normalized by revenue in 2019 from March 25th 2020 to April 14th 2020 over the normalized net revenue from Jan 8th to March 10th 2020. The changes in low-wage worker hours (both for small restaurants/retail - HomeBase and in general - Earnin) are defined as the change in hours from March 25th 2020 to April 14th 2020 relative to total hours from Jan 8th to March 10th 2020.

Appendix Table 4

Association Between Changes in Consumer Spending Home Area and Workplace Area Rents

| Dep. Var.: % Change in Total Credit Card Spending | (1) | (2) | (3) |
|---|---------------------|---------------------|---------------------|
| Median Workplace 2BR Rent | -0.0129 (0.0006) | -0.0089 (0.0012) | -0.0121 (0.0039) |
| Median Home 2BR Rent | | -0.0065 (0.0017) | |
| <i>Controls:</i> | | | |
| County Fixed Effects | | | X |
| Observations | 8,934 | 6,682 | 8,934 |

Notes: This table shows OLS regressions of average percentage changes in consumer spending by ZIP code code (using data from Affinity Solutions) on average workplace ZIP code code median two-bedroom rent. Standard errors are reported in parentheses. Workplace ZIP code code rent is computed by using data from the Census LEHD Origin-Destination Employment Statistics (LODES) database as described in the text. The dependent variable is scaled from 0 to 100 such that, for example, the coefficient of -0.0129 in Column (1) implies that a \$100 increase in monthly workplace rent is associated with a 1.2% larger drop in total spending. Column (1) shows the baseline regression without any controls, Column (2) adds median home two bedroom rent and Column (3) adds county level fixed effects.

Appendix Table 5
List of Partial Re-Openings and Control States for Event Study

| Date | States that Re-Opened | Affinity Controls | Earnin Controls | Womply Controls | Google Controls |
|------------------|------------------------|--|---|--|---|
| April 20th, 2020 | South Carolina | California, Connecticut, Delaware, District Of Columbia, Florida, Hawaii, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Missouri, Nebraska, New Jersey, New Mexico, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | California, Connecticut, Delaware, District Of Columbia, Florida, Hawaii, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | Indiana, Missouri, Nebraska, New Mexico, Oregon, South Dakota |
| April 24th, 2020 | Alaska, Georgia | California, Connecticut, Delaware, Florida, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Missouri, Nebraska, New Mexico, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | Delaware, Indiana, Louisiana, Missouri, Nebraska, New Mexico, South Dakota, Virginia, Wisconsin |
| April 24th, 2021 | Minnesota, Mississippi | Illinois, Nebraska, Pennsylvania, South Dakota, Virginia, Wisconsin | Delaware, District Of Columbia, Illinois, Maryland, Nebraska, New Mexico, South Dakota, Virginia, Wisconsin | California, Connecticut, Delaware, District Of Columbia, Illinois, Maryland, Nebraska, New Jersey, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin | Delaware, Illinois, Nebraska, New Mexico, Pennsylvania, South Dakota, Virginia, Wisconsin |

Notes: This table lists the treatment and control states for each opening date in Figures 12B-12D

Appendix Table 6
City to County Crosswalk

| City Name | State Name | County | County Fips Code |
|------------------|----------------------|----------------------|------------------|
| Los Angeles | California | Los Angeles | 6037 |
| New York City | New York | Richmond | 36085 |
| New York City | New York | Kings | 36047 |
| New York City | New York | Queens | 36081 |
| New York City | New York | New York | 36061 |
| New York City | New York | Bronx | 36005 |
| Chicago | Illinois | Cook | 17031 |
| Houston | Texas | Harris | 48201 |
| Phoenix | Arizona | Maricopa | 4013 |
| San Diego | California | San Diego | 6073 |
| Dallas | Texas | Dallas | 48113 |
| Las Vegas | Nevada | Clark | 32003 |
| Seattle | Washington | King | 53033 |
| Fort Worth | Texas | Tarrant | 48439 |
| San Antonio | Texas | Bexar | 48029 |
| San Jose | California | Santa Clara | 6085 |
| Detroit | Michigan | Wayne | 26163 |
| Philadelphia | Pennsylvania | Philadelphia | 42101 |
| Columbus | Ohio | Franklin | 39049 |
| Austin | Texas | Travis | 48453 |
| Charlotte | North Carolina | Mecklenburg | 37119 |
| Indianapolis | Indiana | Marion | 18097 |
| Jacksonville | Florida | Duval | 12031 |
| Memphis | Tennessee | Shelby | 47157 |
| San Francisco | California | San Francisco | 6075 |
| El Paso | Texas | El Paso | 48141 |
| Baltimore | Maryland | Baltimore | 24005 |
| Portland | Oregon | Multnomah | 41051 |
| Boston | Massachusetts | Suffolk | 25025 |
| Oklahoma City | Oklahoma | Oklahoma | 40109 |
| Louisville | Kentucky | Jefferson | 21111 |
| Denver | Colorado | Denver | 8031 |
| Washington | District of Columbia | District Of Columbia | 11001 |
| Nashville | Tennessee | Davidson | 47037 |
| Milwaukee | Wisconsin | Milwaukee | 55079 |
| Albuquerque | New Mexico | Bernalillo | 35001 |
| Tucson | Arizona | Pima | 4019 |
| Fresno | California | Fresno | 6019 |
| Sacramento | California | Sacramento | 6067 |
| Atlanta | Georgia | Fulton | 13121 |
| Kansas City | Missouri | Jackson | 29095 |
| Miami | Florida | Dade | 12086 |
| Raleigh | North Carolina | Wake | 37183 |
| Omaha | Nebraska | Douglas | 31055 |
| Oakland | California | Alameda | 6001 |
| Minneapolis | Minnesota | Hennepin | 27053 |
| Tampa | Florida | Hillsborough | 12057 |
| New Orleans | Louisiana | Orleans | 22071 |
| Wichita | Kansas | Sedgwick | 20173 |
| Cleveland | Ohio | Cuyahoga | 39035 |
| Bakersfield | California | Kern | 6029 |
| Honolulu | Hawaii | Honolulu | 15003 |
| Boise | Idaho | Ada | 16001 |
| Salt Lake City | Utah | Salt Lake | 49035 |
| Virginia Beach | Virginia | Virginia Beach City | 51810 |
| Colorado Springs | Colorado | El Paso | 8041 |
| Tulsa | Oklahoma | Tulsa | 40143 |

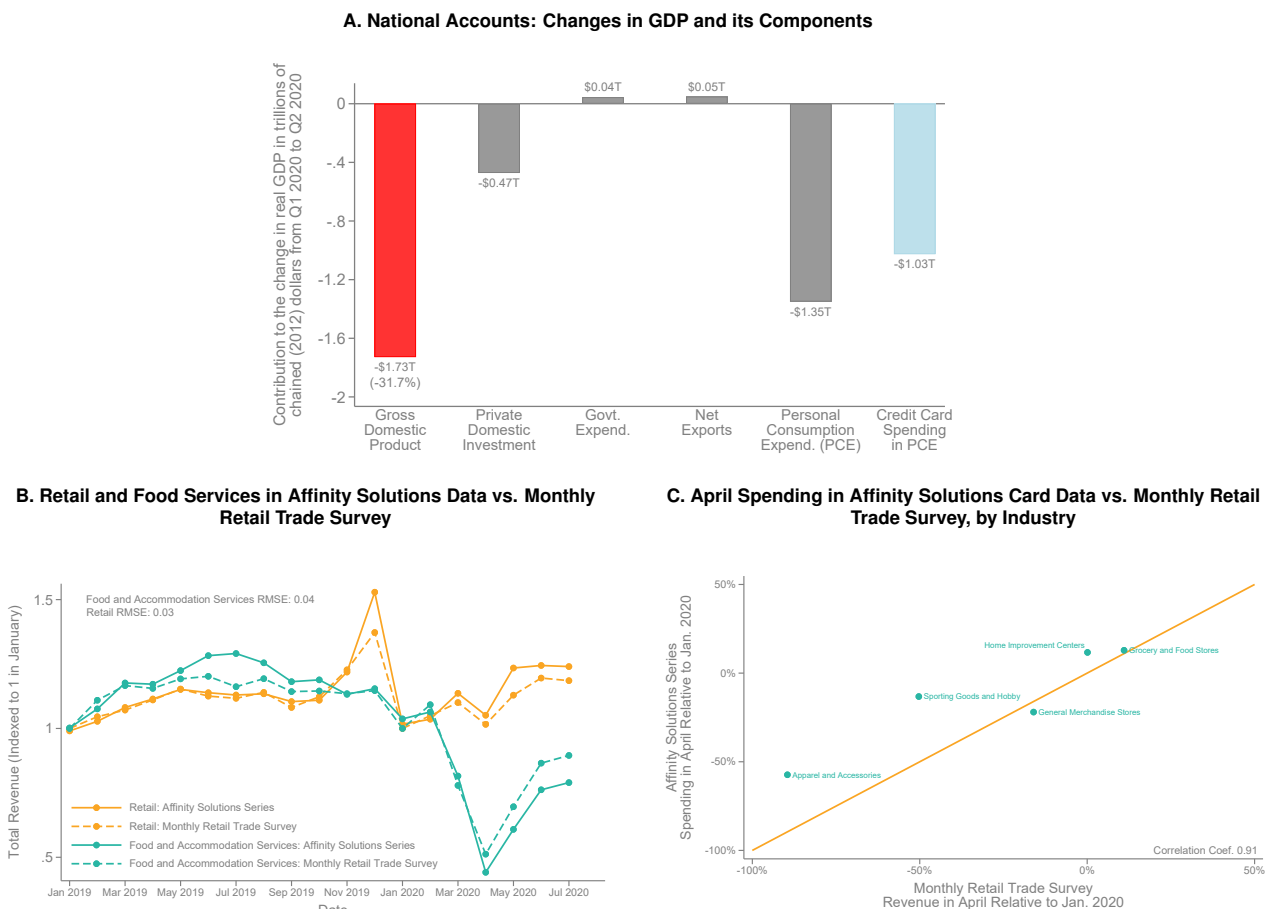
Notes: This table shows our metro area (city) to county crosswalk. We assigned metros to counties and ensured that a significant portion of the county population was in the metro of interest. Some large metros share a county, in this case the smaller metro was subsumed into the larger metro.

Appendix Table 7
Industry Employment Shares Across Data Sets

| NAICS Code | NAICS Description | QCEW Industry Shares | | Industry Shares in Private Datasets | | | |
|------------|--|-------------------------|---------------------------|-------------------------------------|----------------|--------|--------|
| | | QCEW All Establishments | QCEW Small Establishments | Homebase | Paychex-Earnin | Kronos | Intuit |
| | | (1) | (2) | (3) | (4) | (5) | (6) |
| 11 | Agriculture, Forestry, Fishing and Hunting | 0.84 | 1.04 | | 0.61 | | 0.83 |
| 21 | Mining, Quarrying, and Oil and Gas Extraction | 0.55 | 0.43 | | 0.21 | | 0.08 |
| 22 | Utilities | 0.44 | 0.29 | | 0.17 | 0.79 | 0.17 |
| 23 | Construction | 5.72 | 7.62 | | 6.35 | 1.13 | 7.32 |
| 31-33 | Manufacturing | 10.27 | 5.16 | | 8.48 | 22.14 | 2.24 |
| 42 | Wholesale Trade | 4.72 | 5.99 | | 5.78 | | 1.66 |
| 44-45 | Retail Trade | 12.48 | 14.06 | 11.28 | 8.32 | 3.72 | 4.65 |
| 48-49 | Transportation and Warehousing | 4.30 | 2.82 | 0.87 | 2.26 | 10.39 | 1.58 |
| 51 | Information | 2.29 | 1.64 | | 1.63 | | 0.94 |
| 52 | Finance and Insurance | 4.83 | 4.60 | | 3.57 | 5.88 | 1.72 |
| 53 | Real Estate and Rental and Leasing | 1.71 | 2.90 | | 3.08 | | 1.85 |
| 54 | Professional, Scientific, and Technical Services | 7.63 | 8.97 | 2.78 | 12.12 | | 11.84 |
| 55 | Management of Companies and Enterprises | 1.93 | 0.79 | | 0.45 | | 0.15 |
| 56 | Administrative Support | 7.25 | 5.30 | | 6.61 | | 5.00 |
| 61 | Educational Services | 2.39 | 1.53 | 3.62 | 2.43 | 1.11 | 1.18 |
| 62 | Health Care and Social Assistance | 16.16 | 13.16 | 5.34 | 15.21 | 22.71 | 5.71 |
| 71 | Arts, Entertainment, and Recreation | 1.78 | 1.64 | 2.07 | 2.17 | 1.77 | 1.30 |
| 72 | Accommodation and Food Services | 11.04 | 15.60 | 49.17 | 11.09 | 10.20 | 2.61 |
| 81 | Other Services (except Public Administration) | 3.57 | 6.21 | 1.83 | 8.74 | | 5.96 |
| 99 | Unclassified | 0.11 | 0.24 | 23.04 | 0.72 | 20.15 | 43.2 |

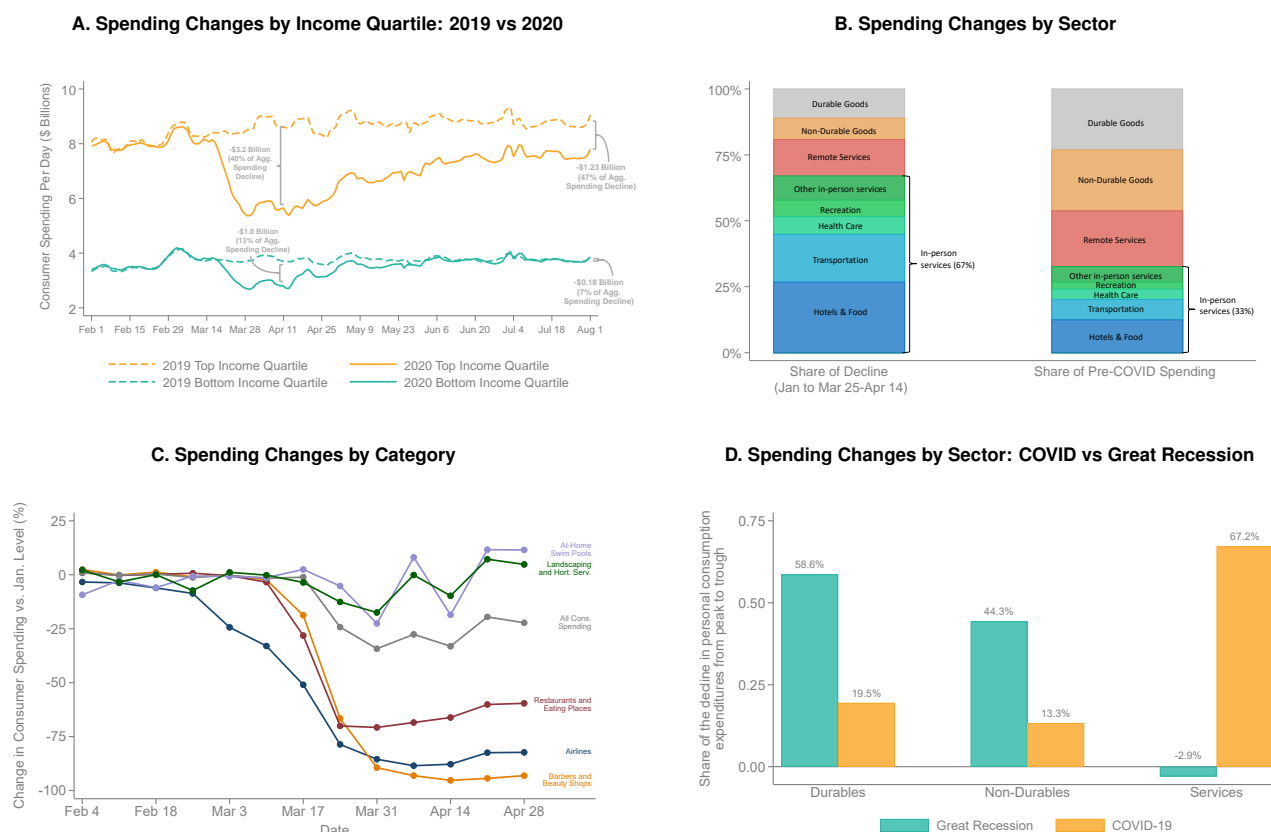
Notes: This table reports the NAICS two-digit industry mix for four private employment-based datasets compared with the Quarterly Census of Employment and Wages (QCEW), an administrative dataset covering the near-universe of firms in the United States. Columns (1) - (6) indicate the share of employees in the given dataset who work in the specified sector. In columns (1) and (2), we construct data for all establishments and small establishments using employment data from the Q1 2019 QCEW. Small establishments are defined as having fewer than 50 employees. In columns (3) - (6), we construct employment shares for the private datasets. For Homebase, Paychex-Earnin, Kronos, and Intuit we use January 2020 employment to do so. We define employment in Homebase as the number of unique individuals working a positive number of hours in the month. We define employment in Paychex-Earnin as a weighted sum of employment in Paychex (the total number of employees recorded by Paychex clients in the month) and employment in Earnin (the total number of worker-days in the month). We define employment in Kronos as the total number of employee punches in the month. We define employment in Intuit as the total number of employees recorded by Intuit clients in the month. Industries missing from any of the private data sources are left blank.

FIGURE 1: Consumer Spending in National Accounts vs. Credit and Debit Card Data



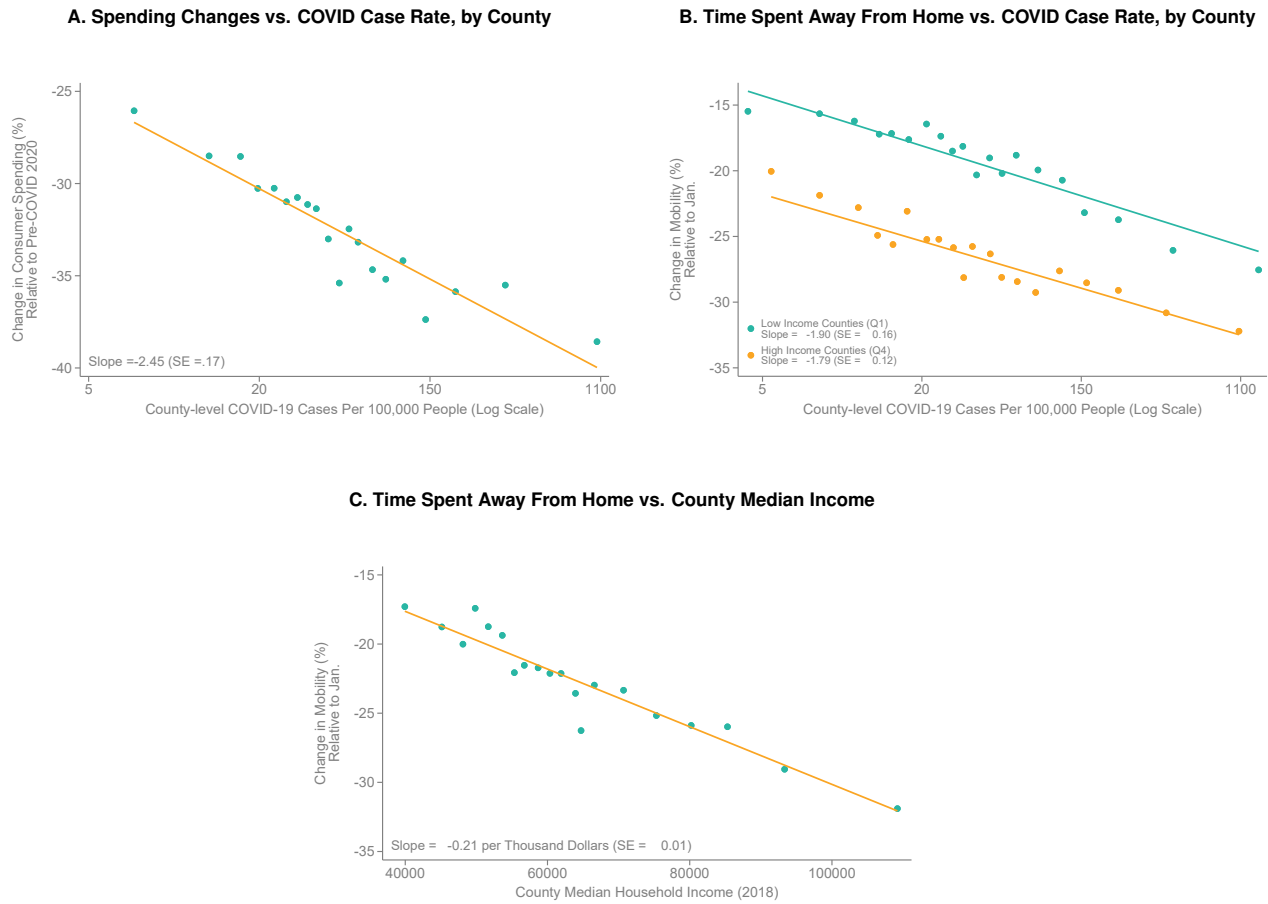
Notes: This figure compares changes in consumer spending in national income and product accounts (NIPA) data to measures of consumer spending recorded on debit and credit cards. Panel A summarizes NIPA data (Tables 1.1.2, 1.1.6 and 2.3.2), comparing Q1-2020 to Q2-2020. The first bar shows the seasonally adjusted change from Q1-2020 to Q2-2020 in real GDP in chained (2012) dollars (-\$1.73T). In parentheses under the first bar is the compound annual growth rate corresponding to this one-quarter change in real GDP (-31.7%). Bars two through five show the contribution to the change in real GDP of its components. These contributions are estimated by multiplying the change in real GDP (-\$1.73T) by the contributions to the percent change in real GDP given in Table NIPA 1.1.2. The final bar shows the contribution of components of Personal Consumption Expenditures (PCE) that are likely to be captured in credit card spending (-\$1.03T). This includes all components of PCE except for motor vehicles and parts, housing and utilities, health care, and the final consumption expenditures of nonprofit institutions serving households. This bar is computed by multiplying the change in PCE (-\$1.35T) by the contributions to the percent change in PCE given in NIPA Table 2.3.2 (excluding the aforementioned subcategories). Panels B and C report average daily spending for each month in the Affinity Solutions credit and debit card data and the Monthly Retail Trade Survey (MRTS), a government survey providing estimates of sales at retail and food services stores across the United States. The retail series for panel B restricts to retail trade sectors (NAICS code 44-45) excluding motor vehicles (NAICS code 441) and gas (NAICS code 447). The food and accommodation services series of panel B includes food and accommodation services in the Affinity Solutions data (NAICS codes 721 and 722), but restricts to food services in the MRTS (NAICS code 722). All four series are normalized relative to January of the corresponding year. The top left corner displays the root mean squared error (RMSE) of a regression of indexed MRTS monthly spending on indexed Affinity Solutions monthly spending for both retail and food and accommodation services. Panel C restricts to industries where the industry definitions in the Affinity Solutions series align closely with a three digit NAICS code surveyed in the MRTS. Data source: Affinity Solutions

FIGURE 2: Changes in Consumer Spending During COVID Crisis



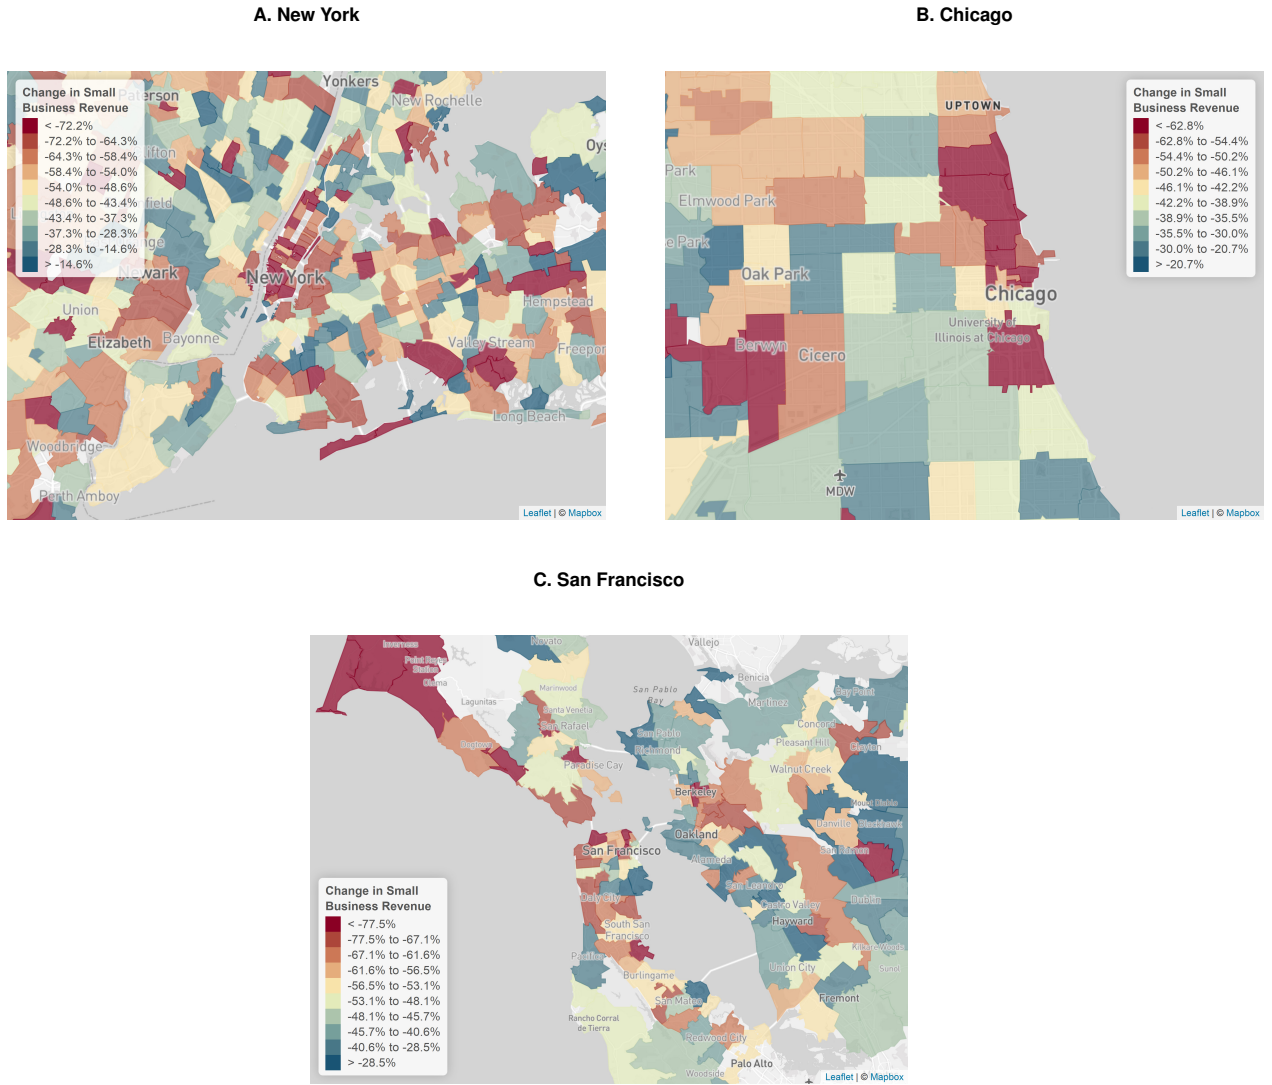
Notes: This figure disaggregates spending changes by income and sector in the COVID crisis using debit and credit card data from Affinity Solutions. Panel A plots a weekly series of consumer spending for cardholders residing in ZIP codes in the top and bottom quartiles of the distribution of ZIP code median household income (defined using population-weighted data from the 2014-2018 ACS). We scale the 2020 (2019) series by multiplying by the ratio of January 2020 (2019) total spending for components of PCE that are likely captured in credit/debit card spending (shown in the last bar of Figure 1a) to the January 2020 (2019) total spending in the Affinity Solutions data. We impute the value plotted for February 29, 2019 with the average of February 22, 2019 and March 7, 2019. Panels B and C show industry level spending changes using the Affinity Solutions data. Panel B disaggregates spending changes (left bar) and pre-COVID levels (right bar) by sector. The left bar in Panel B shows the share of the total decline in spending which can be attributed to the different sectors. The total decline is defined as ((Spending in March 25 through April 14 2020) - (Spending in March 26 through April 15 2019)) - ((Spending in January 8 through January 28 2020) - (Spending in January 8 - January 28 2019)). The second bar shows the share of spending in January 8-28 of 2020 for each sector. We define durable goods as the following MCC groups: motor vehicles, sporting goods and hobby, home improvement centers, consumer electronics, and telecommunications equipment. Non-durable goods include wholesale trade, agriculture, forestry and hunting, general merchandise, apparel and accessories, health and personal care stores, and grocery stores. Remote services include utilities, professional/scientific services, public administration, administration and waste services, information, construction, education, and finance and insurance. In-person services include real estate and leasing, recreation, health care services, transportation and warehousing services, and accommodation and food, as well as barber shops, spas, and assorted other services. Non-durables consist of 5.2% of the decline as show in the left-hand side bar and 23.0% of January spending. Excluding grocery stores from non-durable spending, non-durables constitute 11.6% of the decline and 10.5% of January spending. Panel C compares trends in consumer spending for six specific categories of goods. Consumer spending for each good is normalized by its level over the period January 4-31, and seasonally adjusted using 2019 data, as described in Section II.A. Panel D decomposes the change in personal consumption expenditures (PCE) in the COVID-19 Recession and the Great Recession using NIPA data (Table 2.3.6U). PCE is defined here as the sum of services, durables and non-durables in seasonally adjusted, chained (2012) dollars. For COVID-19 (Great Recession), the peak is defined as January 2020 (December 2007) and the trough is April 2020 (June 2009). Data source: Affinity Solutions

FIGURE 3: Association Between COVID-19 Incidence, Spending, and Time Outside Home



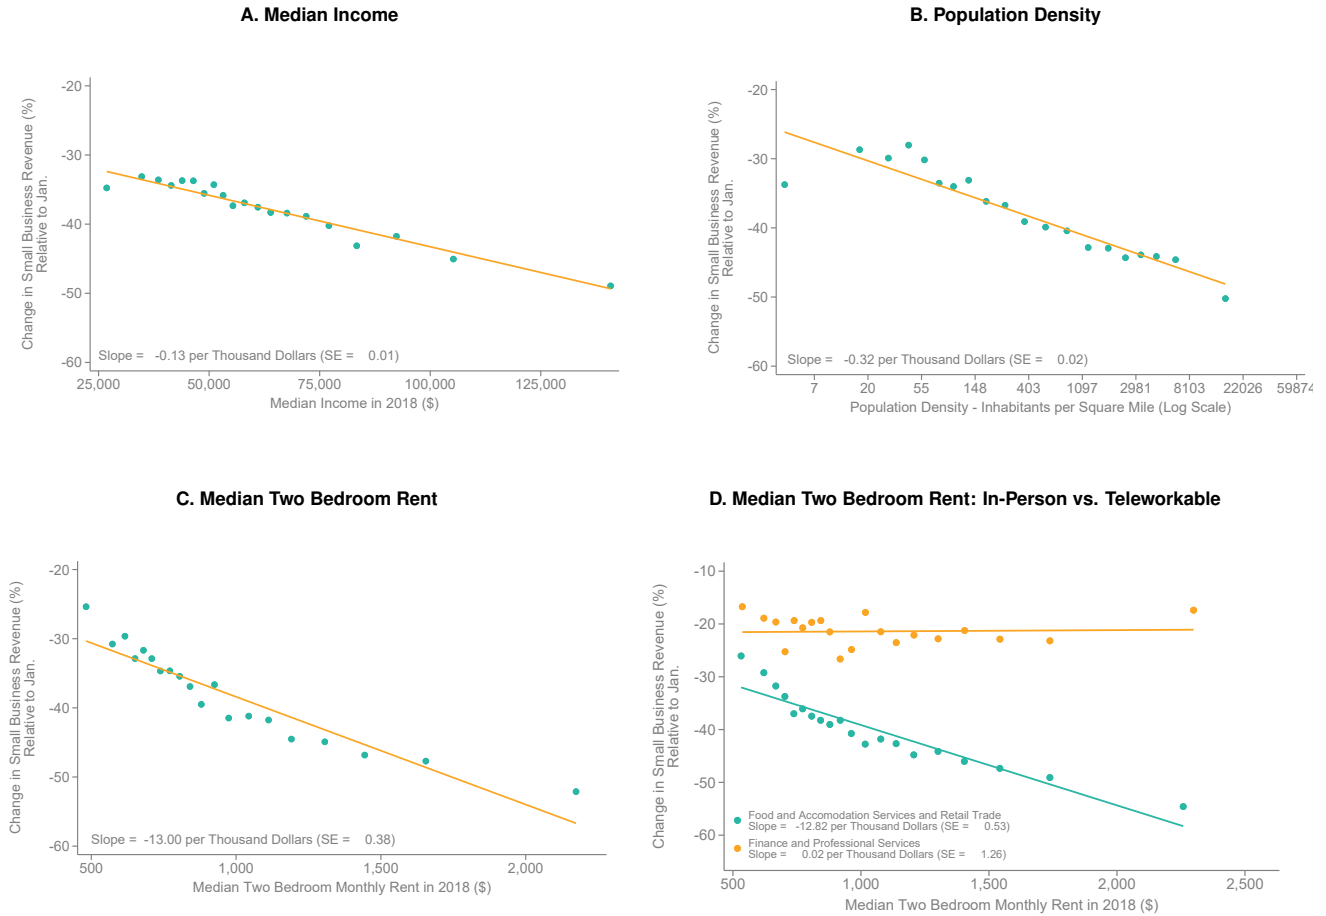
Notes: This figure presents three county-level binned scatter plots. To construct each binned scatter plot, we divide the data into twenty equal-sized bins, ranking by the x-axis variable and weighting by the county's population, and plot the (population-weighted) means of the y-axis and x-axis variables within each bin. Panel A presents a binned scatter plot of the change in average weekly consumer spending from the base period (January 8 - January 28) to the two-week period from April 1 - April 14 (using data from Affinity Solutions) vs. the county's COVID per-capita case rate over the two week period from April 1 - April 14. Panel B presents a binned scatter plot of the change in time spent outside the home in each county between January and the three-week period from March 25 - April 14 (using Google Mobility data) vs. the county's COVID case rate, separately for low and high-income counties. We define low-income and high-income counties as those with median household income in the bottom 25% and top 25% of all counties respectively, as measured in the 2014-2018 ACS (weighted by county population). Panel C presents a binned scatter plot of the change in time spent outside home in each county between January and the three-week period from March 25 - April 14 vs. the county's median household income. Data sources: Affinity Solutions, Google Mobility

FIGURE 4: Changes in Small Business Revenues by ZIP code




Notes: This figure plots percentage small business revenue declines during the COVID crisis by ZIP code from March 25 through April 14 in the MSAs corresponding to New York-Newark-Jersey City, NY-NJ-PA MSA (Panel A), Chicago-Naperville-Elgin, IL-IN-WI MSA (Panel B), and San Francisco-Oakland-Hayward, CA MSA (Panel C). The change in revenue is determined in each ZIP code using data from Womply. We normalize weekly revenue by the average weekly revenue from January 1 to 28 for both 2019 and 2020, and then divide the 2020 proportion by the 2019 proportion. We calculate the signal-to-noise ratio by regressing normalized weekly revenue on an indicator variable denoting whether the week is after March 9, 2020, within each ZIP code, denoting the coefficient and standard error on this indicator variable in each ZIP code as β_z and SE_z , respectively. We then calculate the signal-to-noise ratio as $1 - \sum SE_z^2 / \sum \beta_z^2$. The signal variance to total variance ratios for the panels are 0.83 (New York), 0.89 (Chicago), and 0.73 (San Francisco). These maps must be printed in color to be interpretable; dark red colors represent areas with larger revenue declines, while dark blue colors represent areas with smaller declines. Data source: Womply

FIGURE 5: Changes in Small Business Revenues vs. ZIP code Characteristics



Notes: Panels A-C present binned scatter plots showing the relationship between changes in small business revenue using data from Womply vs. various ZIP Code-level characteristics. Binned scatter plots are constructed as described in Figure 3. We calculate changes in business revenue in each ZIP code using average normalized weekly revenue, as defined in Figure 4, from March 25 to April 14. We exclude data from ZIP codes in which changes are larger than 200% or where variance of normalized revenue exceeded 900%. To preserve the privacy of firms included in the data, ZIP Code-by-industry cells with average weekly revenue of less than \$4,250 during the base period were also excluded. The x variable in Panel A is median household income at the ZIP code level, from the 2014-2018 ACS. The x variable in Panel B is the log number of inhabitants per square mile in the ZIP code the 2014-18 ACS. The x variable in Panel C is the median rent for a two-bedroom apartment in the ZIP code in the 2014-2018 ACS. Panel D replicates Panel C for two sectors: in person services, defined as Food and Accommodation (NAICS 72) and Retail Trade (NAICS 44 and 45), vs. sectors in which workers are more likely to be able to telework, defined as Finance and Professional Services (NAICS 52 and NAICS 54). Data source: Womply


FIGURE 6: Changes in Employment Rates Over Time




Notes: This figure compares employment changes relative to January 2020 within various datasets. In Panel A, we combine Paychex, Earnin, and Intuit data to construct a daily private employment series for all industries. We then construct an employment index by averaging employment over the prior seven days and then norming to the average value of employment over the period January 4 - January 31, 2020. The Current Employment Statistics (CES) data, as well as the Current Population Survey (CPS) are available monthly, so we plot changes in each month relative to January 2020. The CES is a monthly survey of firms at the parent level. The CPS is a monthly survey of households, which we then adjust to match a payroll definition of employment. The CES reports employment for the pay period including the 12th of each month, and the CPS is fielded during the week of the 19th of each month, so we plot these monthly series on the 15th of the month. Panel B replicates the combined private employment series and the CES series from figure A, but instead restricts to employment in the two-digit NAICS sector 72, Accommodations and Food Services and NAICS supersector 60, Professional and Business Services. In addition, we plot a series for NAICS 72 firms in the Homebase data. Data sources: Paychex, Earnin, Intuit, Homebase

FIGURE 7: Changes in Employment by Income Quartile and Consumer Spending


A. Changes in Employment by Income Quartile



B. Changes in Employment by Income Quartile and Consumer Spending, All Industries

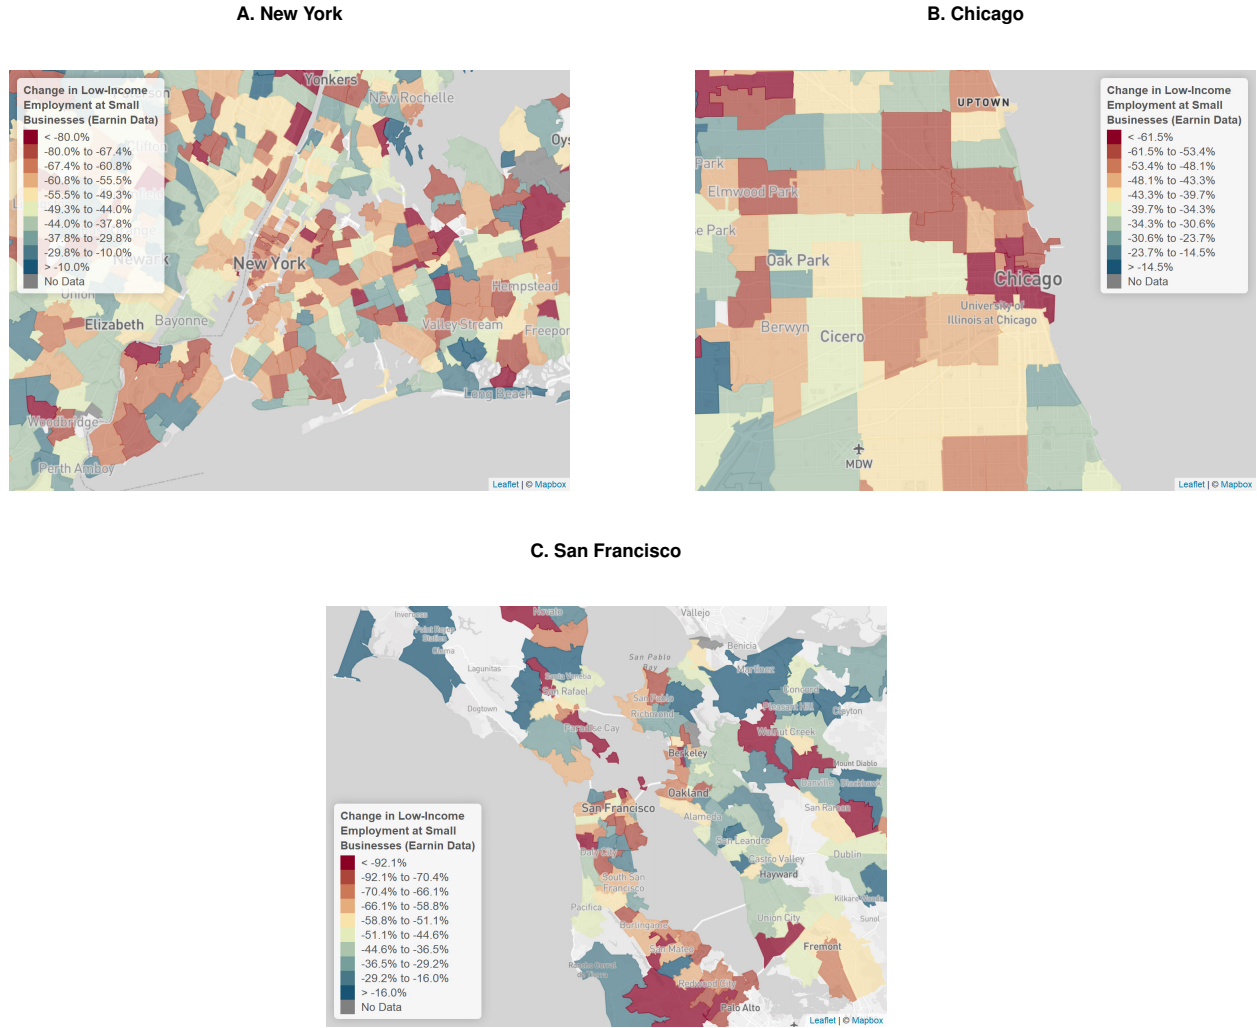


C. Changes in Employment by Income Quartile and Consumer Spending, Retail Trade



Notes: This figure shows changes in employment by income quartile and consumer spending relative to January 2020. In each panel, we present a daily employment series constructed by combining Paychex, Earnin and Intuit data; for details, see notes to Figure 6. Panel A shows changes in employment at the national level since January 2020, split by income quartile. The solid series shows changes in combined Paychex-Earnin-Intuit data, and the dotted series from August 22 to September 25 is forecasted using Kronos data and Paychex data from firms with weekly paycycles.. To construct this forecast, we regress the combined Paychex-Earnin-Intuit series on the de-seasonalized Kronos series for the same date (t), the Paychex weekly series for the same date (t), and the Paychex weekly series for three weeks prior ($t - 7$), ($t - 14$), ($t - 21$). We then use the resulting coefficients to predict values of combined Paychex-Earnin-Intuit employment. Panels B and C compare changes in employment by income quartile to changes in consumer spending. In both panels, the solid spending series is constructed using Affinity Solutions data. Panel B shows changes in spending and employment across all industries. To construct the employment series in Panel B, we begin with combined Paychex-Earnin-Intuit employment at the county x 2-digit NAICS code x income quartile x date level. We restrict the sample to county x 2-digit NAICS code cells which have positive first-quartile and fourth-quartile employment in the period January 4-31 2020; this sample restriction excludes 0.6% of worker-days from the sample. We then calculate the change in employment since January 4-31 2020 in each county x 2-digit NAICS code x income quartile cell as of each date, winsorizing at the 99th percentile (weighted by total employment in the period January 4-31 2020). To construct first-quartile (fourth-quartile) employment, we take the mean change in first-quartile employment as of each date, weighting by first-quartile (fourth-quartile) employment in January 4-31 in each cell. To construct first-quartile employment reweighted to match fourth-quartile employment on NAICS code x county, we take the mean change in first-quartile employment as of each date, weighting by fourth-quartile employment in January 4-31 in each cell. Panel C shows changes in spending and employment in NAICS 44-45 (Retail Trade) since January 4-31 2020. Data sources: Paychex, Intuit, Earnin, Kronos, Affinity Solutions


FIGURE 8: Changes in Employment Rates by ZIP code




Notes: This figure replicates Figure 4 using the changes in employment from Earnin. We focus on small and medium-sized businesses, defined as firms with at most 500 employees (as measured in the ReferenceUSA Data). For users whose employer cannot be matched to ReferenceUSA data on firm sizes, we restrict to users whose employer is in the fourth decile or below of firms in the Earnin data, in terms of number of Earnin users working for the firm; the median firm size for the fourth decile of Earnin employers is roughly 300 employees. We calculate normalized weekly employment as total employment in each week (measured by positive payroll deposits for an individual) divided by total average weekly employment in the period January 4-31, 2020. We calculate the signal-to-noise ratios in the same way as in Figure 4; these ratios are 0.79 in New York, 0.59 in Chicago, and 0.67 in San Francisco. These maps must be printed in color to be interpretable; dark red colors represent areas with larger employment declines, while dark blue colors represent areas with smaller declines. Data source: Earnin

FIGURE 9: Changes in Employment Rates and Job Postings vs. Rent


A. Employment in Earnin Data vs. Median Rent, by ZIP



B. Employment in Combined Paychex/Intuit/Earnin Data vs. Median Rent, by County



C. Job Postings for Low-Education Workers vs. Median Rent, by County




D. Job Postings for High-Education Workers vs. Median Rent, by County




Notes: This figure shows binned scatterplots of the relationship between changes in employment rates and median rents (Panels A-B) and job postings and median rent (Panels C-D), by ZIP code. Binned scatter plots are constructed as indicated in Figure 3 by binning ZIP codes (Panel A) and counties (Panels B-D) based on their median rent into 20 equally sized bins and computing the mean change in the outcome variable within each bin. Panel A shows the relationship between average employment rates in the Earnin data and median rents, separately for small firms (less than or equal to 500 employees), medium-sized firms (between 500 and 10,000 employees), and large firms (more than 10,000 employees). Panel B replicates Panel A using as the outcome variable combined employment in the Paychex/Earnin/Intuit combined data; for further details on the construction of this combined series, see Appendix D. Both panels measure the percentage change in hours/employment from January 31, 2020 to April 8-28th, 2020. Panel C shows the relationship between the percentage change in job postings for workers with minimal or some education and median 2 bedroom rent, by county. Panel D replicates Panel C, changing the y variable to the change in job postings for workers with moderate, considerable, or extensive education. Job postings data comes from Burning Glass. Solid lines are best-fit lines estimates using OLS regression, except in Panel C, where we use a lowess fit. Data sources: Earnin, Paychex, Intuit, Burning Glass Technologies

FIGURE 10: Geography of Unemployment in the Great Recession vs. COVID Recession



Notes: This figure displays the share of job losses occurring in low vs high income counties in the Great Recession and the COVID recession. To construct the first set of four bars, we first calculate national employment loss between 2007 and 2010 using data from the BLS. We then group counties into (population-weighted) quartiles by median income, and compute the share of employment loss that occurred in counties in each quartile of the distribution of county median income. The second set of bars replicates the first set of bars using total job losses that occurred between January 2020 and April 2020. The third set of bars reports the share of total initial UI claims between March 15 and April 12, 2020 across counties in different income quartiles. In the first set of bars, county median income is calculated using the 2006 ACS; in the second and third sets of bars, county median income is calculated using the 2014-2018 ACS. In the third bar, we only include counties that are in states which report county-level UI claims data, which comprise 53% of the U.S. population.


FIGURE 11: Changes in Employment and Consumer Spending vs. Workplace Rent for Low-Income Households




Notes: This figure plots changes in employment (Panel A) and consumer spending (Panel B) by ZIP code vs. the average median 2 bedroom rent in the *workplace* ZIP codes of individuals who live in a given ZIP code. The sample is restricted to individuals who live in ZIP codes in the bottom quartile of the household income distribution. We construct the average median 2 bedroom rent variable by combining data on the matrix of home residence by workplace ZIP codes taken from Census' LEHD Origin-Destination Employment Statistics (LODES) with data on median rents from the 2014-2018 ACS. In particular, we assign median rents from the ACS to each ZIP code of workplace in the LODES data and then collapse workplace rents to each home ZIP code, weighting by the number of jobs in each workplace ZIP code. In Panel A, the change in employment variable is based on payroll data from Earnin mapped to workers' home ZIP codes. The change is computed as the percentage change in total employment from Jan 5-Mar 7 to April 8-28, 2020. In Panel B, the spending change variable is based on data from Affinity Solutions on total card spending, mapped to the cardholder's residential ZIP. The change is computed from the period of Jan 5-Mar 7 to Mar 22-April 14, 2020. Data sources: Earnin, Affinity Solutions

FIGURE 12: Causal Effects of Re-Openings on Economic Activity: Event Studies


A. Case Study on Business Re-Openings: Colorado vs New Mexico




B. Re-Opened States vs. Control States: Consumer Spending




C. Re-Opened States vs. Control States: Employment



D. Re-Opened States vs. Control States: Merchants Open

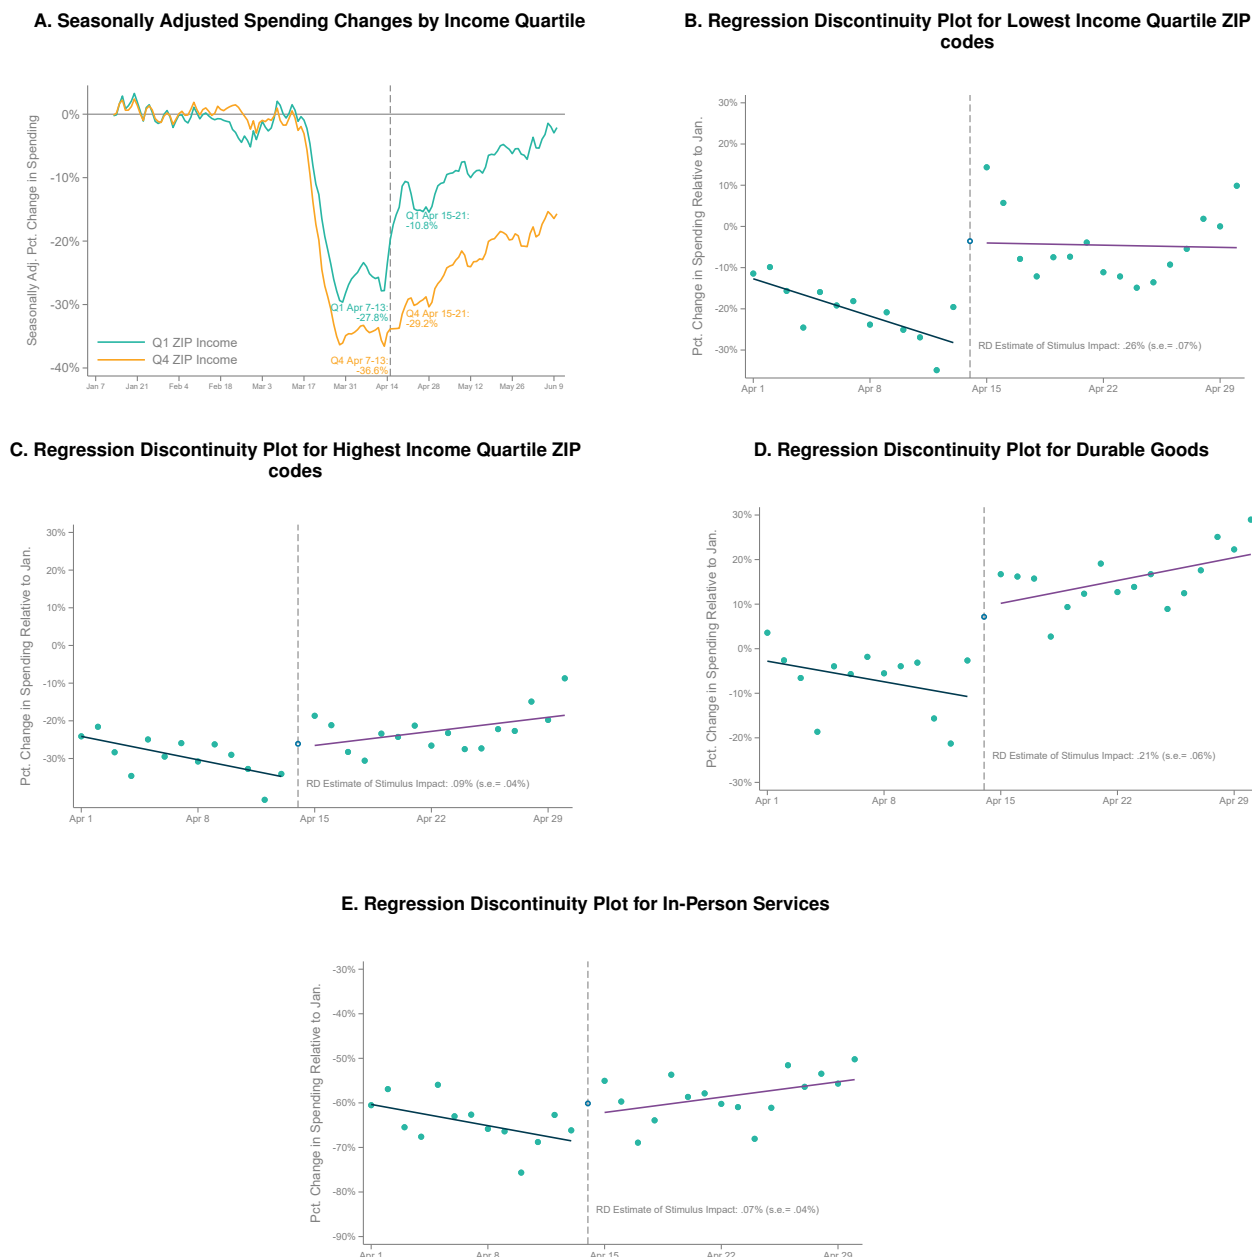


E. Variance Explained by Re-Openings



Notes: Panels A and B show percent change in consumer spending in the Affinity Solutions data. Consumer spending is normalized by its level over the period January 4-31, and seasonally adjusted using 2019 data, as described in Section II.A. Panel A shows the series for both New Mexico and Colorado; Colorado partially reopened non-essential businesses on May 1, while New Mexico did not do so until May 16. Panel B presents an event study of states that partially reopened non-essential businesses between April 20th and May 4th, compared to a matched control group. We construct the control group separately for states on each opening day and then stack the resulting event studies to align the events. Panel C replicates Panel B but instead plots the percent change in employment of workers using combined Paychex-Earnin-Intuit data. Panel D replicates Panel B but instead plots the percent change in open merchants using Womply data. In Panels B-D, we report the coefficient from a difference-in-differences regression comparing treated vs. untreated states in the two weeks following vs. the two weeks prior to the partial re-opening. Panel E reports the share of variance in outcomes explained by reopenings as of May 18. To estimate this, we first calculate the variance of outcome levels across states on May 18, 2020. Then, we add the estimated effect of reopening for a given outcome to all states not open on May 18 (adding only half of the effect is the state opened between May 11 and May 18). This effect is the difference-in-difference estimate reported in Panels B-D. We then recalculate the variance in this counterfactual in which all states had reopened. The share of variance explained by reopenings for each outcome is defined as $1 - (\text{counterfactual variance} / \text{actual variance})$. Data sources: Affinity Solutions, Paychex, Earnin, Intuit, Womply


FIGURE 13: Impact of Stimulus Payments on Consumer Spending




Notes: This figure studies the effect of the stimulus payments made on April 15, 2020 on credit and debit card spending. Panel A plots the percent change in seasonally-adjusted consumer spending in the Affinity Solutions data for cardholders living in ZIP codes in the bottom and top quartiles of the distribution of ZIP code median household income (based on data from the 2014-2018 ACS). Estimates are seasonalized relative to 2019 values as described in notes to Figure 12. In Panels B-E, each point is the national level of spending on that day divided by the average level of spending in January 2020 in the Affinity Solutions data. The points are residualized by day of week and first of the month fixed effects, which we estimate using data from January 1, 2019, to May 10, 2019. In each panel, we also report regression discontinuity estimates of the jump in spending on April 15, using a linear control function before and after April 15 (shown by the solid best fit lines), excluding the partially treated day of April 14, shown by the hollow-point and demarcated by the dashed vertical line. Panel B focuses on cardholders living in ZIP codes in the lowest median household income quartile. Panel C replicates B for cardholders in highest income quartile ZIP codes. Panel D pools all cardholders and examines spending on durable goods, defined in the notes for Figure 2. Panel E considers spending on in-person services, also defined in the notes for Figure 2. Data source: Affinity Solutions

FIGURE 14: Impact of Stimulus Payments on Business Revenue and Employment


A. Regression Discontinuity Plot for Lowest Rent Quartile ZIP codes



B. Regression Discontinuity Plot for Highest Rent Quartile ZIP Codes




C. Revenue and Worker Earnings Changes Among Small Businesses, by ZIP code Rent Quartile




Notes: Panels A and B of this figure study the effect of the stimulus payments made on April 15, 2020 on small business revenues using data from Womply. These panels are constructed in exactly the same way as Panels B and C of Figure 13 except that (1) we use revenue instead of spending as the outcome and (2) we split ZIP codes into quartiles based on median rent for a two-bedroom apartment instead of median household income. Panel C plots the percent change in seven-day moving averages of small-business revenue using Womply data and changes in employment rates using Earnin data, by ZIP code rent-quartile. The employment series restricts to smaller businesses in the Earnin sample, defined by parent employer size being at most 500 employees. The revenue series is seasonally-adjusted by dividing the percentage change from January to each calendar date in 2020 by the corresponding change in 2019; for details, see Appendix C. and The employment series is plotted as a raw change relative to the mean value in January 2020; for details, see Appendix D. Data sources: Womply, Earnin

FIGURE 15: Impact of Paycheck Protection Program on Employment

A. Change in Employment by PPP Eligibility, All Industries Excl. NAICS 72 (Combined Paychex and Earnin Data)



B. Change in Employment by Firm Size, All Industries Excl. NAICS 72 (Combined Paychex and Earnin Data)




Notes: This figure shows the effects of the Paycheck Protection Program on employment. Each panel excludes workers in the Accommodation and Food Services sector (NAICS 72). In both panels, we construct employment series using combined Paychex and Earnin data. Panel A compares employment trends among firms with 100-500 employees, which were eligible for PPP loans, vs. firms with 501-800 employees, which were ineligible for PPP loans. To construct employment trends, we begin by collapsing each dataset to the county x 2-digit NAICS code x income quartile x eligible x date level, where firms are classified as eligible for the PPP if they employ 500 or fewer workers. We then reweight on 2-digit NAICS code and data source such that the industry and data source composition in each size category matches the pooled distribution of industry composition and data source between January 11 and January 31 2020. We then express employment in each county x 2-digit NAICS code x income quartile x eligible cell at each date as a change since January 11-31 2020, winsorizing at the 99th percentile. To construct the “control” series for ineligible firms with between 501 and 800 employees, we take the mean value of employment in each week, expressed as a change relative to January 11-31 2020, and weighting by employment over the period January 11-31 in each county x 2-digit NAICS code x income quartile x eligible cell. To construct the “treatment” series, for each week, we regress change in employment at the county x 2-digit NAICS code x income quartile x eligible level on county x income quartile x week fixed effects, a dummy for the firm being eligible for the PPP. We then add the point estimate for the coefficient on eligibility to the control series in order to form the treatment series. Finally, we recenter each series so that mean change in employment since January 11-31 is 0 percentage points over the period 12 February to 18 March. The grey dashed line corresponds to April 3, 2020, the first day for enrollment in the Paycheck Protection Program (PPP). The annotation on the right of the panel describes the estimated effect on the interaction between eligibility and the date being after April 3 2020, in a regression of change in employment since January on eligibility for the PPP, the date being on or after 3 April 2020, the interaction between eligibility and the date being after April 3 2020, and quartile x county x week fixed effects. Panel B presents a binned scatterplot of changes in employment between the period January 11-31 and the period June 1-23 against firm size. To construct changes in employment by firm size, we first classify firms in bins of size 50 according to their parent employer size. Next, we calculate the mean change in employment among firms in each bin. Finally, we plot this mean change in employment against mean firm size in each bin. Data sources: Paychex, Earnin

FIGURE 16: Effects of COVID on Educational Progress by Income Group




Notes: This figure plots educational progress on the Zearn online math platform for schools located in ZIP codes in the bottom, middle (quartiles 2 and 3), and top quartile of the distribution of median household income (including all ZIP Codes, whether containing Zearn users or not, and weighting by the population in each ZIP Code). Student progress is defined as the number of accomplishment badges earned in Zearn in each week, and the figure plots changes in this metric relative to the mean value between Jan 6-Feb 7, 2020. The sample includes all classes with more than 10 students using Zearn during the base period and at least five users in every week during the reference period. We aggregate data to the ZIP code level weighting by the average number of students using the platform at each school during the reference period. Data source: Zearn

APPENDIX FIGURE 1: Industry Shares of Consumer Spending and Business Revenues Across Datasets




Notes: Panel A shows the NAICS two-digit industry mix for two private business credit card transaction datasets compared with the Quarterly Services Survey (QSS), a survey dataset providing timely estimates of revenue and expenses for selected service industries. Subsetting to the industries in the QSS, each bar represents the share of revenue in the specified sector during Q1 2020. We construct spending and revenue shares for the private datasets, Affinity and Womply, by aggregating firm revenue (from card transactions) in January through March of 2020. Panel B shows the NAICS three-digit industry mix for the same two private datasets compared with the Monthly Retail Trade Survey (MRTS), another survey dataset which provides current estimates of sales at retail and food services stores across the United States. Subsetting to the industries in the MRTS, each bar represents the share of revenue in the specified sector during January 2020. We construct revenue shares for the private datasets, Affinity and Womply, by aggregating firm revenue (from card transactions) in January 2020. Data sources: Affinity Solutions, Womply

APPENDIX FIGURE 2: Cash Spending in CoinOut Transactions Data vs. Affinity Grocery Spending



Notes: This figure displays a comparison between the CoinOut series and Affinity Grocery Expenditure Series. Panel A plots 7-day moving averages of these two series overlaid at the national level. Panels B and C plot total expenditure in CoinOut and total grocery expenditure in Affinity by county rent quartile, compared to January averages. Panel B displays total transaction amounts in CoinOut, while Panel C displays grocery expenditure in Affinity. The signal correlation between the two datasets at the national level is 0.90 at the weekly level between January 1st and June 9th. This correlation is computed by collapsing both datasets to the national weekly level, where values in each week are expressed as a percentage change from the January average. The correlation between both datasets is taken at this level. To adjust to get the signal correlation, correlations are also taken within each dataset, correlating the values in each week with the values in the following week. The correlation between datasets is then divided by the square root of the product of the within-dataset correlations to get the signal correlation. Data sources: CoinOut, Affinity Solutions

APPENDIX FIGURE 3: Industry Shares of Job Postings in Burning Glass and Job Openings in JOLTS




Notes: This Figure displays the NAICS two-digit industry mix of job postings in Burning Glass and job openings in JOLTS, the Job Openings and Labor Turnover Survey data provided by the U.S. Bureau of Labor Statistics, in January 2020. Data source: Burning Glass

APPENDIX FIGURE 4: Spending Changes by Sector and Income Quartile




Notes: This figure displays the change in spending by sector for the four quartiles of ZIP code median household income (constructed using 2014-2018 ACS population and income estimates). These sectors were constructed by grouping together similar merchant category codes, not all merchant category codes were used in this plot. The change in spending displayed is $(\text{the log difference-in-difference of spending} - 1) \times 100$, where the pre-period used is January 8th-28th and the post-period is March 25th-April 14th. Data source: Affinity Solutions

APPENDIX FIGURE 5: Spending Changes vs COVID Cases, by County



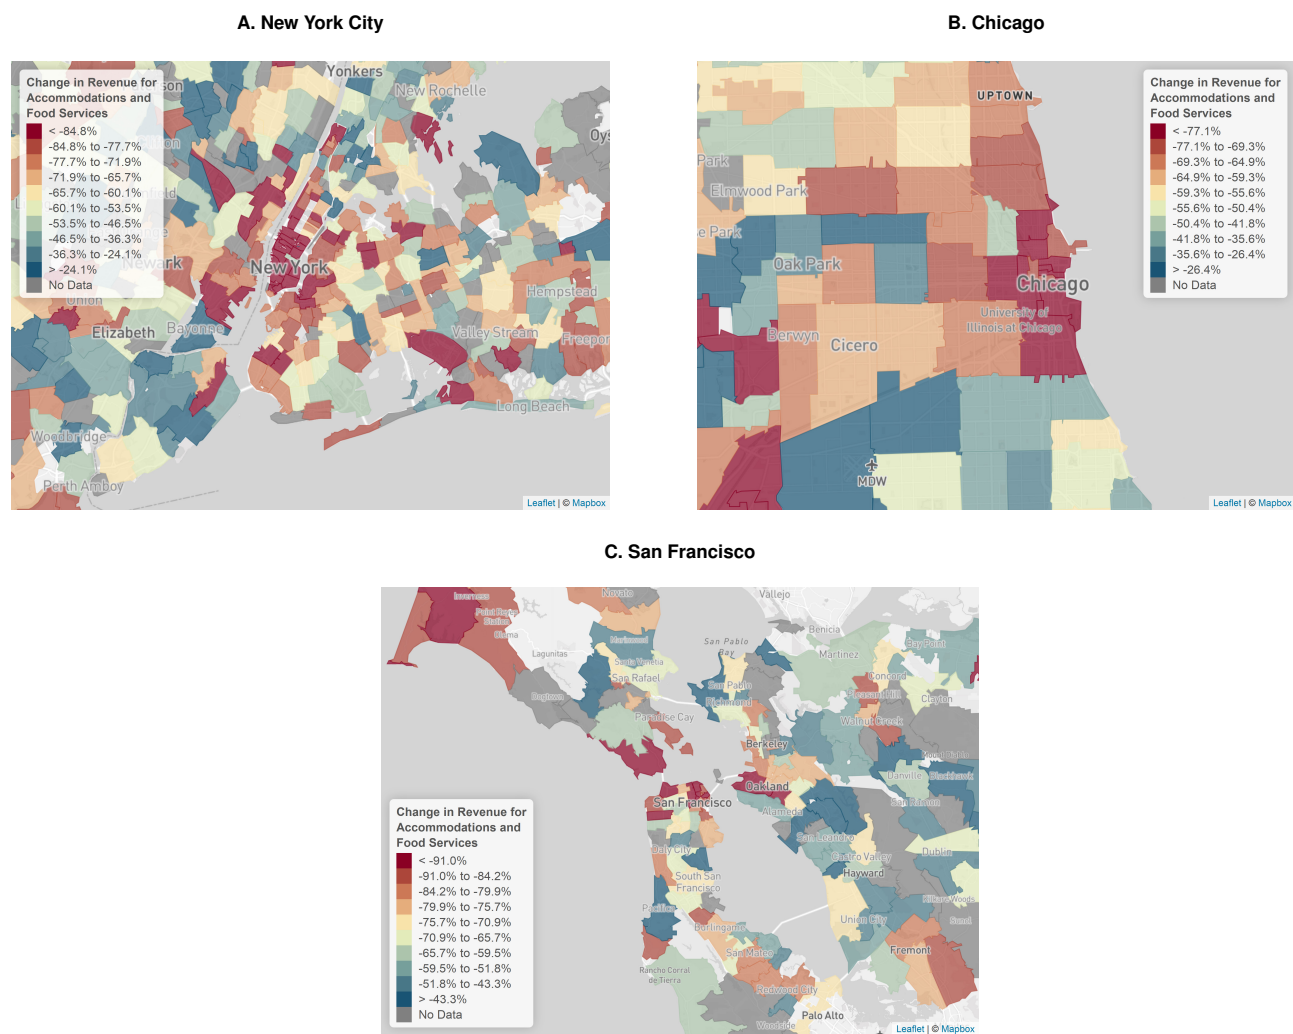
Notes: To construct this figure, we divide the log COVID cases into 20 bins, each of which contain 5% of the population, and plot the mean value of the log of COVID cases and change of spending variables within each bin, controlling for state fixed effects and median-household income. COVID cases and decline in spending are both measured during the two week period of April 1st to April 14th, and is benchmarked to the pre-period of January 8th to January 28th. Data source: Affinity Solutions.

APPENDIX FIGURE 6: Small Business Revenue Changes vs. Consumer Spending Changes



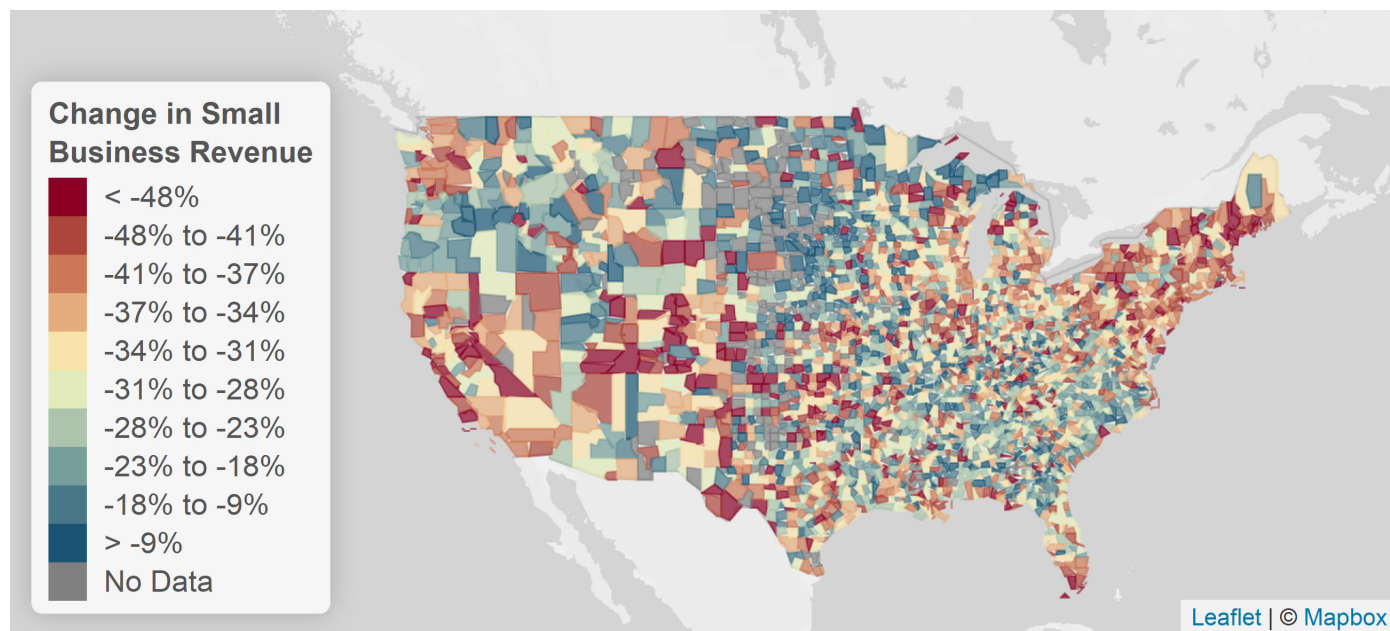
Notes: This figure compares the seven-day moving average of total consumer spending (from Affinity Solutions purchase data) and small business revenue (from Womply) normalized to the average January and February levels of each year. Following the sectors defined in the Monthly Retail Trade Survey (MRTS), Panel A restricts to specifically retail trade sectors (NAICS code 44-45) excluding motor vehicles (NAICS code 441) and gas (NAICS code 447), and Panel B restricts specifically to food services and accommodations (NAICS code 72). The bottom right corner of each panel displays the root mean squared error (RMSE) of a regression of indexed Affinity Solutions daily spending on indexed Womply daily spending. Data sources: Affinity Solutions, Womply

APPENDIX FIGURE 7: Changes in Small Business Revenues by ZIP code for Food and Accommodation Service Businesses



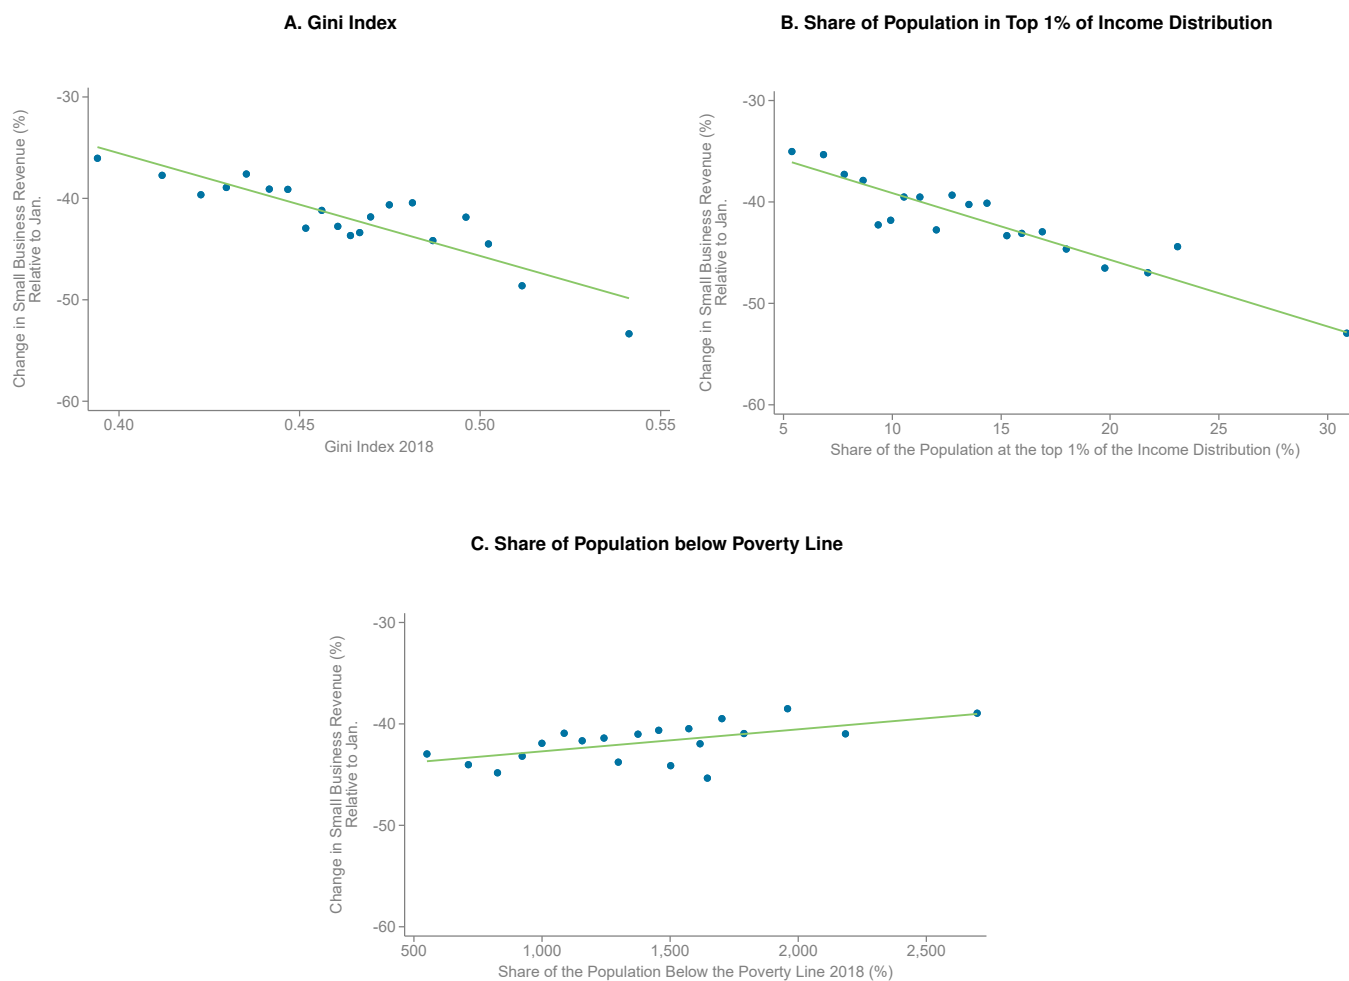
Notes: This Figure replicates Figure 4 within Food and Accommodation Service (NAICS 72) small businesses. For details, see notes to Figure 4. The signal variance to total variance ratios for the panels are 0.82 (New York), 0.89 (Chicago), and 0.68 (San Francisco). Data source: Womply

APPENDIX FIGURE 8: Changes in Small Business Revenues by County




Notes: This figure replicates Figure 4 for the entire United States instead of a single MSA, using counties instead of ZIP Codes. See notes to Figure 4 for details. The signal variance to total variance ratio for this map is 0.71. Data source: Womply

APPENDIX FIGURE 9: Womply Business Revenue vs. Poverty Share, Top 1% Share, and Gini by County




Notes: This Figure replicates Figure 5 but compares the declines with different measures of inequality. Panel A compares the within county Gini index against the declines. Panel B uses the share of the county with incomes at the top 1% of the income distribution. Panel C compares the declines with the share of the county population with incomes below the poverty line in the 2010 decennial census. See notes to Figure 5 for details. Data source: Womply

APPENDIX FIGURE 10: Womply Business Closures vs. Rent by ZIP



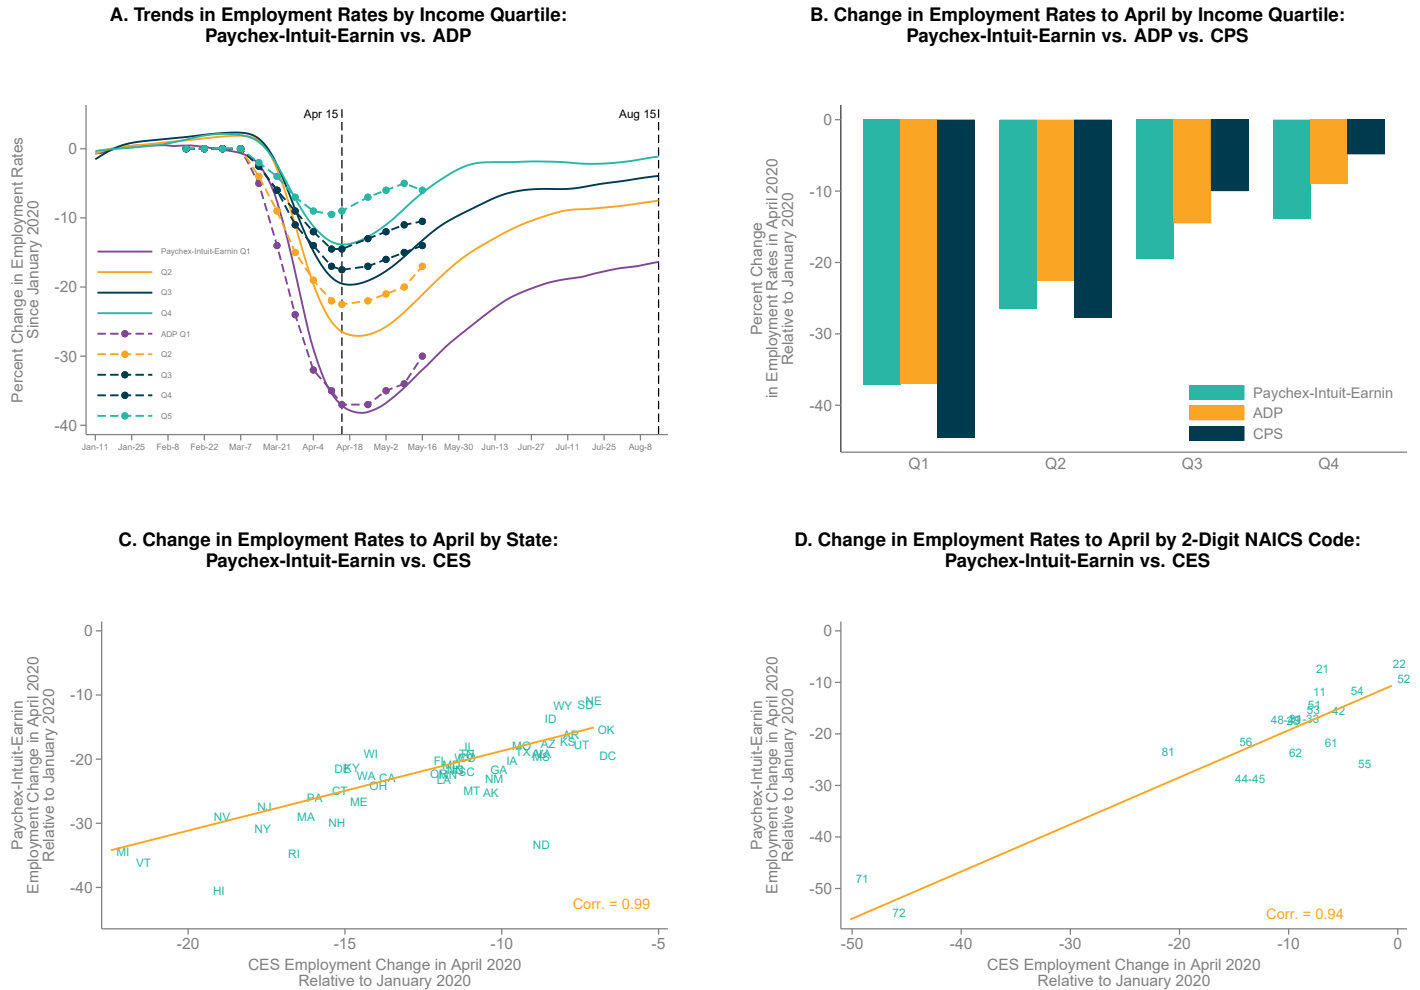
Notes: This figure replicates Panel C of Figure 5 but shows average changes in small businesses that remain open instead of changes in revenue. See notes to Figure 5 for details. Data source: Womply

APPENDIX FIGURE 11: Changes in Wages and Employment Over Time



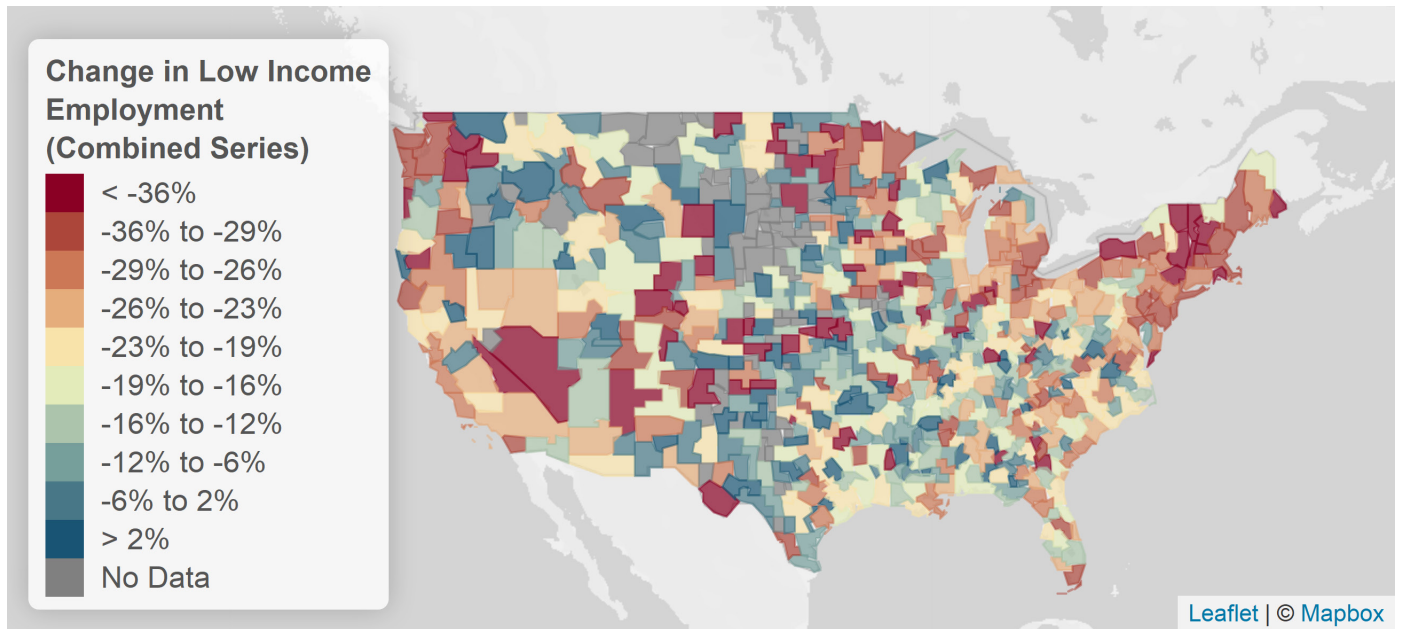
Notes: This figure compares changes in mean wages and employment relative to January 2020 within the Earnin dataset. We construct employment by calculating total workers employed on each day and then taking a 7-day moving average; we express the employment series as a percentage change relative to January 4-31 2020. We construct wages for Earnin by calculating the change between wages in week w and wages in week $(w + 1)$, for workers who are employed in both weeks. We then express our wage series by chaining together changes in wages since January 2020. Data source: Earnin

APPENDIX FIGURE 12: Paychex-Earnin-Intuit Combined Employment vs. ADP, CPS and CES Employment




Notes: This figure compares the Paychex-Earnin-Intuit combined employment series to the Current Population Survey (CPS), the Current Employment Statistics (CES), and estimates based on data from ADP by Cajner et al. (2020). Panel A shows employment trends in Paychex-Earnin-Intuit combined data (solid series) and ADP data (dotted series), cut by income quartile (Paychex-Earnin-Intuit combined employment) or income quintile (ADP data). The Paychex-Earnin-Intuit series is expressed as a percentage change relative to January 4-31 2020. The ADP series in Cajner et al. (2020) is expressed as a percentage change relative to February 15 2020. Panel B shows changes in employment to April 2020, cut by income level, in the Paychex-Earnin-Intuit combined, ADP, and CPS datasets. In the combined Paychex-Earnin-Intuit data, we express the change in employment relative to January 4-31 2020. The ADP series in Cajner et al. (2020) is expressed as a percentage change relative to February 15 2020. The CPS series is expressed as a percentage change relative to January 2020. Panel C shows a scatterplot of changes in employment in Paychex-Earnin-Intuit combined data between January 4-31 and April 15 vs. changes in CES employment between January and April, by state. We exclude Hawaii and North Dakota, where Paychex-Earnin-Intuit data has poor coverage. Panel D shows a scatterplot of changes in employment in Paychex-Earnin-Intuit combined data between January 4-31 and April 15 vs. changes in CES employment between January and April, by two-digit NAICS code. In Panels C and D, the bottom right corner displays the correlation between changes in CES employment and changes in Paychex-Earnin-Intuit combined employment, weighted by state population (Panel C) and CES employment in each NAICS code (Panel D), respectively. Data sources: Paychex, Earnin, Intuit

APPENDIX FIGURE 13: Changes in Low-Income Employment Rates by CZ



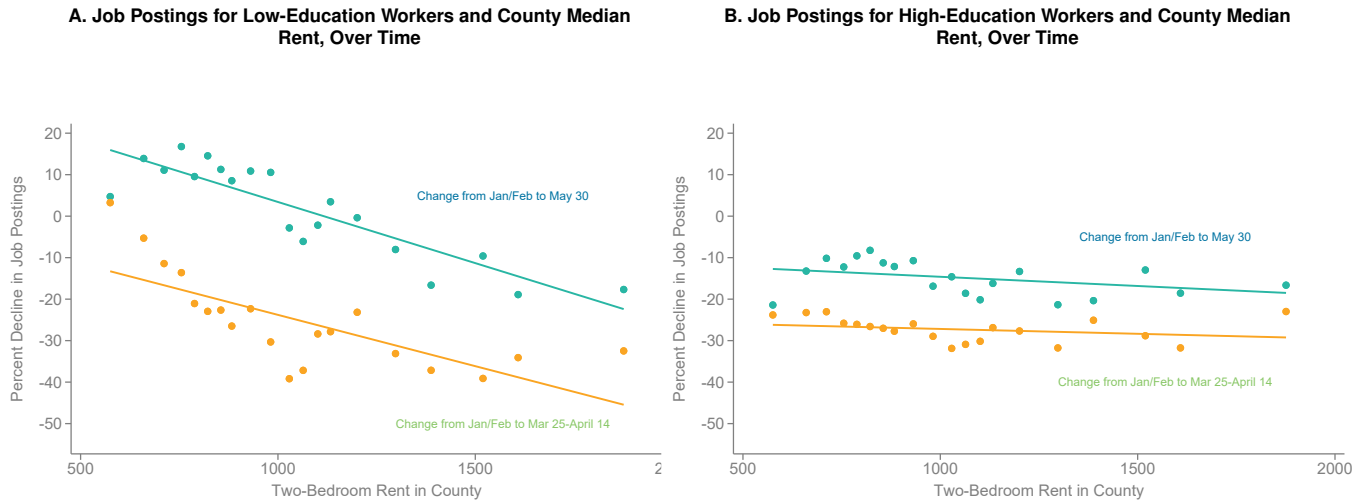
Notes: This figure replicates Figure 7 at the CZ level for the entire United States instead of a single city and its surrounding area, and using Paychex-Earnin-Intuit combined employment in the first quartile, rather than Earnin data on first-quartile employment. See notes to Figure 7 for details. For details on the construction of Paychex-Earnin-Intuit Data combined employment, see Appendix D. Data sources: Paychex, Earnin, Intuit

APPENDIX FIGURE 14: Changes in Employment by Firm Size




Notes: This figure displays a binned scatterplot of average percent declines in employment in the Earnin data at firms of different sizes. The decline is calculated by taking total employment by firm in a pre-period that spans from January 8, 2020 to January 31, 2020, and comparing to employment in a post-period that spans from June 1, 2020 to June 23, 2020. We estimate the size of firms by matching Earnin employer names and locations to employer names and locations in ReferenceUSA data; for details, see Appendix D. Data source: Earnin

APPENDIX FIGURE 15: Changes in Job Postings vs. Rent Over Time




Notes: This figure shows binned scatterplots of the relationship between median rent and changes in job postings between a pre-period of January 8 - March 10 and the periods March 25 - April 14 or the period May 30-June 5. The change in job postings is computed using Burning Glass data. Median two-bedroom rent is computed using the 2014-2018 ACS at the county level. Panel C presents a binned scatterplot of the relationship between the percentage change in job postings for workers with minimal or some education and median 2 bedroom rent. Panel D presents a binned scatterplot of the relationship between the percentage change in job postings for workers with moderate, considerable or extensive education and median 2 bedroom rent. Data source: Burning Glass Technologies

APPENDIX FIGURE 16: Legislated Stay-at-Home Orders and Non-Essential Business Closures




Notes: This figure shows percent change in seasonally-adjusted consumer spending in the Affinity Solutions data, pooling together states that closed non-essential business early (between March 19th and March 24th), states that closed non-essential businesses late (between March 30th and April 6th), and those that never closed. Data source: Affinity Solutions

APPENDIX FIGURE 17: Daily Treasury Statements and IRS Transactions Among Earnin Users




Notes: This figure displays daily total spending on Stimulus Checks (IRS Individual Tax Rebates) from the Daily Treasury Statment in Panel A. Panel B displays the total dollar amount of IRS transactions for Earnin users. Data source: Earnin

APPENDIX FIGURE 18: Impact of Stimulus on the Composition of Consumer Spending



Notes: See notes of Figure 2 Panel B. The pre-stimulus, post-COVID period is defined as March 25th-April 14th. The post-stimulus period is defined as April 29th to May 5th. The total recovery is computed use the post-stimulus period and the average weekly spending in the pre-stimulus period. This figure disaggregates spending by Merchant Category Codes (MCCs), grouping together similar MCCs. We define durable goods as the following MCC groups: motor vehicles, sporting goods and hobby, home improvement centers, consumer electronics, and telecommunications equipment. Non-durable goods include wholesale trade, agriculture, forestry and hunting, general merchandise, apparel and accessories, health and personal care stores, and grocery stores. Remote services include utilities, professional/scientific services, public administration, administration and waste services, information, construction, education, and finance and insurance. In-person services include real estate and leasing, recreation, health care services, transportation and warehousing services, and accommodation and food, as well as barber shops, spas, and assorted other services. Data source: Affinity Solutions


APPENDIX FIGURE 19: Histograms of PPP Eligibility Firm Size Cutoffs for Firms with 300 to 700 Employees in ReferenceUSA




Notes: This figure plots a histogram of the firm size cutoffs for PPP eligibility in the set of firms in ReferenceUSA. In the ReferenceUSA data, we take the establishment-size-weighted distribution of PPP employee-based eligibility thresholds, which are based on parent company size. We exclude NAICS 72. Data source: ReferenceUSA

APPENDIX FIGURE 20: Effect of Paycheck Protection Program on Employment (Earnin and Kronos Data)


A. Change in Employment by PPP Eligibility, All Industries Excl. NAICS 72 (Earnin Data)




B. Change in Employment by Firm Size, All Industries Excl. NAICS 72 (Earnin Data)



C. Change in Employment by PPP Eligibility, All Industries Excl. NAICS 72 (Kronos Data)




D. Change in Employment by Firm Size, All Industries Excl. NAICS 72 (Kronos Data)



Notes: This figure replicates Figure 15 using Earnin data (Panels A and B) and Kronos data (Panels C and D), rather than combined Paychex-Earnin data. For details, see notes to Figure 15. Data sources: Earnin, Kronos

APPENDIX FIGURE 21: Out-Of-Sample Fit of Advance Series



Notes: This figure compares the out-of-sample fit for different time periods from a model predicting employment in the top and bottom wage quartile in the combined Paychex-Earnin-Intuit series. See Figure 7 for details on how the model is constructed. Panel A compares out-of-sample fit in the testing period of May 25 through June 25. Panel B compares the out-of-sample fit in the testing period of June 15 through July 15. The Root Mean Square Error (RMSE) for the prediction model across the top, middle, and bottom quartiles in the first testing period is 1.030, while the RMSE across the top, middle, and bottom quartiles in the second testing period is 0.893. Data sources: Paychex, Earnin, Intuit, Kronos