Real-Time Economics: A New Public Platform to Analyze the Impacts of COVID-19 and Macroeconomic Policies Using Private Sector Data*

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

August 2020

Abstract

We build a publicly available platform 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 physical 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 most of their low-income employees, leading to a surge in unemployment claims in affluent areas. 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 very 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 modestly, by approximately 2.5%, implying a cost per job saved in excess of $300,000. 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 can be constructed from private sector data to 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 Stablization 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, 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 Shalinee 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 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, Cailtin Kupsc, Sarah LaBauve, Maddie Marino, Kai Matheson, Kate Musen, Danny Onorato, Sarah Oppenheimer, Trina Ott, Lynn Overmann, Max Pienkny, Jeremiah Prince, Daniel Reuter, Peter Ruhm, Emanuel Schertz, Kamelia Stavreva, James Stratton, 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 detailed 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 how the coronavirus pandemic (COVID-19) affected the U.S. economy and the causal effects of stabilization policies in near-real-time – within three weeks of the economic 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 core 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 navigate this tradeoff, one must construct statistics that are sufficiently aggregated and masked to mitigate privacy concerns

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.

yet sufficiently granular to support many research and policy analyses that would be infeasible with current publicly available data.[2]

We combine several methods to navigate this tradeoff and to comply with legal regulations governing the disclosure of material non-public information: reporting only changes since January 2020 (rather than raw levels), omitting small cells, and pooling data from multiple companies. 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 mitigate 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 typically less than seven days after the relevant transactions occur.

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 reduction in GDP 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 in particular 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

2. An alternative approach to navigating this tradeoff – pursued by several recent studies summarized at the end of this section – is to use 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. We find that the time series patterns in the Homebase data differ significantly from representative statistics on employment for the sectors it covers (although the qualitative patterns are generally similar). We therefore turn to other sources of employment data to produce publicly available series that more closely track representative benchmarks and are now available for future research.

by *high-income* households. High-income households both cut spending more heavily when the COVID shock first hit in mid-March and then increased spending more substantially afterwards. As of August 15, more than half of the total reduction in card spending since January had come from households in the top quartile of the income distribution; only 5% 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 reduced their spending by 17%, whereas low-income households reduced their spending by only 5% as of August 15.

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 physical 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

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).

the fact that many of the sectors in which spending fell most are non-tradable goods produced by 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: by the end of May, employment levels for workers in the top wage quartile were back to baseline pre-COVID levels, but still remained 14% below baseline for low-wage workers. 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.

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 reopenings, stimulus payments to households, and loans to small businesses.[6]

State-ordered 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 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

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.

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. In particular, employment rates at firms within fewer than 500 employees (who were eligible for PPP assistance) trended similarly to those at larger firms that were ineligible for PPP both before and after April 3, when the PPP program began. At the upper bound of the 95% confidence intervals, our estimates imply that the cost per job saved by the PPP exceeds $300,000, likely because the vast majority of PPP loans went to inframarginal firms that were not planning to lay off many workers. The results are consistent with those of 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), who report similar results using data from ADP. 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.

In sum, the primary barrier to economic activity in the COVID-19 pandemic was depressed consumer spending due to health concerns as opposed to government restrictions on economic activity, inadequate income among consumers, or a lack of liquidity for firms. These findings suggest that the only path to full economic recovery in a pandemic is to restore 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 pandemics affect 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 create 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 individual behavioral responses lead to cascades in market 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 within three weeks after the economic shock or policy change of interest (e.g., the deposit of stimulus checks).[7] This

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

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, we view the data assembled here as a first 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.

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

many questions studied by others with the aggregate statistics that we construct.

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. 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.[11]

In the rest of this section, we describe each of the series in turn, discussing the raw data sources,

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.

11. We construct metro area values for large metro areas using a county to metro area crosswalk described in the Appendix.

construction of key variables, and cross-sectional comparisons to publicly available benchmarks.[12] 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

We measure consumer spending using aggregated and anonymized consumer purchase data collected by Affinity Solutions Inc, a company that aggregates consumer credit and debit card spending information to support a variety of financial service products.

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. Industries are defined by grouping together similar merchant category codes. ZIP code income quartiles are constructed at the national level using Census data on population and median household income by ZIP. Cells with fewer than five unique card transactions are masked.

The raw data include several discontinuous breaks caused by entry or exit of credit card providers from the sample. We identify these breaks using data on the total number of active cards in the cell. We then estimate the discontinuous level shift in spending resulting from the break (using a standard regression discontinuity estimator). At the state level (including Washington, DC), we adjust the series within each cell by adding the RD estimate back to the raw data to obtain a smooth series. At the county-level, there is too much noise to implement a reliable correction, so we exclude counties that exhibit such breaks from the sample. 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.[13] Finally, we index the seasonally-adjusted series relative to pre-COVID-19 spending by dividing each day's value by the mean of the seasonally-adjusted seven-day moving average from January 8-28.

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. Affinity Solutions captures nearly 10% of debit and credit card spending in the U.S. To assess which categories of spending

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

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

are covered by the Affinity data, Appendix Figure 1 compares the spending distributions across sectors to spending captured in the nationally representative Quarterly Services Survey (QSS) and Monthly Retail Trade Survey (MRTS). Affinity has broad coverage across industries. However, as expected, it over-represents categories where credit and debit cards are used for purchases. In particular, accommodation and food services and clothing are a greater share of the Affinity data than financial services and motor vehicles. We therefore view Affinity as providing statistics that are representative of total card spending (but not total consumer spending). We assess whether Affinity captures changes in 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.[14] 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.

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

14. 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.

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.[15]

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 establishment size. 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 establishment 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). 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).

The Paychex data is based on payrolls processed over the preceding five weeks ending on the Thursday of the given week, to account for the fact that pay cycles vary across workers (weekly, bi-weekly, monthly). We shift the Paychex data back by a period of 35 days, to account for the fact that the raw data as of a given date reflects the preceding five weeks.

15. 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.

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 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 at the paycheck level with information on home ZIP, workplace ZIP, unemployment status, earnings, and hours worked over the past 28 days from January 2020 to present. 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 and are 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

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.

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 electronic time sheet. We obtain these data by county, industry, and firm size, where firm size is the total number of employees at a parent firm, (rather than at a single establishment). 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 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 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 employment in the COVID pandemic, but we do not include it in our primary employment indices because it does not track publicly available benchmarks (for the sectors it represents) as closely as our other data sources (see Section III.C below).

Combined Employment Series. To protect business privacy and maximize precision, we combine the data sources above to construct our primary employment series. Because Paychex covers all sectors and wage levels fairly comprehensively, we use it as the base for the combined employment series. We then use Earnin and Intuit to refine the series in cells represented by those datasets.

Earnin best represents workers in the bottom wage quartile (see Appendix Table 1 below). We therefore combine Earnin data with Paychex data to construct employment estimates for the bottom wage quartile. To do so, we first calculate total employment levels within each two-digit

NAICS code by geography cell by summing employment levels for bottom-wage-quartile Paychex workers and Earnin workers.[17] We then compute mean levels of employment for bottom-wage-quartile workers by geography by taking a weighted average of the NAICS-by-geography combined estimates, weighting by the January Paychex NAICS shares for bottom-wage-quartile workers in each geography.

Next, we combine Intuit with the Paychex+Earnin data (see Appendix Figure 12 below). Intuit provides us with overall national industry shares as of 2019, but does not release data broken down by wage level or industry. We therefore must effectively impute the Intuit data to wage-industry cells in order to combine it with the Paychex data. To do so, we assume that any differences in employment between Intuit and Paychex are constant (in percentage terms, relative to the January baseline) by industry and wage quartiles within a given geography and month. We reweight the Paychex data to match the national Intuit industry distribution and compute the percentage difference between the employment decline in the reweighted Paychex data and the Intuit data in each geography-month cell. We then apply this correction factor to each wage-industry cell in the Paychex data to obtain imputed values by wage and industry for the relevant industries covered by Intuit. For instance, if Intuit exhibits a 5% larger employment decline than the reweighted Paychex series in Manhattan in April, we would impute a value for each wage-by-industry cell covered in the Intuit data that 1.05 times the decline observed in Paychex for that cell. Finally, we take a weighted average of the Paychex data and the imputed Intuit data in each industry to compute the final combined series. We place the majority of the weight on Paychex, with greater weight on Intuit in sectors where it has greater coverage; the exact weights are undisclosed to protect privacy.

The preceding steps yield combined data at the industry by wage quartile for each geography (county, state, and national). We construct aggregate estimates across industries, wage quartiles, and overall by aggregating these estimates using Paychex January employment weights.[18] We report seven-day moving averages of these series, expressed as a percentage change relative to January 4-31. We construct a series for average total earnings analogously, used total earnings instead of total employment.

Employment Predictions Based on Time Sheet Data. The employment series constructed based

17. We convert the weekly Paychex data to daily measures of employment by assuming that employment is constant within each week.

18. In a few cases, Earnin and Intuit data do not provide coverage for a given geographical region or industry; we suppress such cells. We also suppress cells in which Paychex records less than an average of 100 total monthly employees in the second half of 2019 at the industry by geography or income quartile by geography level. When aggregating employment series to the geographical level without breakdowns by industry or wage quartile, however, we use data from all cells, without masking.

on payroll data is generally available only with a one month lag because people get paid after completing work over multiple prior weeks. To obtain more current estimates, we use the Kronos punches data to construct a forecast of employment rates for low-wage (bottom quartile) workers that is available with a lag of 7 days. The Kronos punches data (which we make publicly available in raw form) closely track the payroll-based employment data for the bottom wage-quartile over the weeks when both series are available, with a lead of approximately 7 days as expected given the payroll cycle (Appendix Figure 24). To create a prediction of employment rates based on the Kronos data, we regress the combined employment series for bottom-quartile workers on the Kronos series for the same week and three prior weeks at the relevant geography. We then use these regression coefficients combined with the most recent Kronos data to create a series of predicted employment rates for low-wage workers in each geography.

Comparisons to OES and QCEW. Appendix Table 7 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 1). 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 particularly good coverage of workers at the bottom of the wage distribution, who are a group of particular interest given their particularly 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 2, 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 $247 billion (an annualized rate of 5%) from the fourth quarter of 2019 to the first 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 $230 billion.[19] Government purchases did not change significantly, while net exports increased by $65 billion and private investment fell by $90 billion.[20] 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.[21]

Benchmarking. We begin by assessing whether the credit card data track patterns in corresponding spending categories in the national accounts. Figure 1b plots spending on retail services (excluding auto-related expenses) in the Affinity Solutions credit card data alongside the Monthly Retail Trade Survey (MRTS), one of the main inputs used to construct the national accounts. Both series are indexed to have a value of 1 in January 2020; each point shows the level of spending in a given month divided by spending in January 2020. Figure 1c replicates Figure 1b for spending on food services. In both cases, the credit/debit card spending series closely tracks the inputs that make up the national accounts. In particular, both series show a rapid drop in food services spending in March and April 2020 and a smaller drop in retail spending, along with an increase in May. Given that credit card spending data closely tracks the MRTS at the national level, we proceed to use it to disaggregate the national series in several ways to understand why consumer

19. GDP is released at a quarterly level in the U.S. The reduction in consumer spending occurred in the last two weeks of March (Figure 2 below); hence the first quarter GDP estimates capture about one-sixth of the reduction in spending due to the COVID shock.

20. Most of the reduction in private investment was driven by a reduction in inventories and equipment investment in the transportation sector, both of which are plausibly a response to reductions in current and anticipated consumer spending. The increase in net exports was driven primarily by a reduction in imports, with a large reduction in imports of travel and transporation services in particular, again reflecting a change in domestic consumer spending behavior.

21. 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.

spending fell so sharply.[22]

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.[23]

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.[24] The solid line shows data from January to May 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 \$5.4 billion per day by the end of March (a 31% 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 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. We estimate that as of mid-April, top-quartile households accounted for 39% 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 mid June, top-quartile

22. 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 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.

23. 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.

24. 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.

households accounted for over half of the total spending decline in the U.S. and were still spending 15% 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. Table 1 provides standard errors for the numbers on Figure 2a - the differences between high- and low-income households are highly significant. These differences are also highly significant comparing high- and low-income households in the same county.

Heterogeneity Across Sectors. Next, we disaggregate the change in total 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 COVID.

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

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

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.

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 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

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.

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 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 spending shocks that firms face. 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

28. Of course, these results only apply to the period we study: the first three months after the COVID shock hit the U.S. The recession may turn into a more traditional economic shock with Keynesian spillovers across a wider set of sectors and areas as time passes, in which case tools such as stimulus and liquidity could become much more impactful (Guerrieri et al. 2020).

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).

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 available here). There is substantial heterogeneity in revenue declines across areas. For example, average revenue declines range from -67% (or below) in the lowest-income-decile of ZIP codes to -15% (or above) in the top-income-decile in 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 3, 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]

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).

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

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]

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 (Appendix Table 3).

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

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.

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, although it exhibits a larger drop in employment at the trough. This is because our data capture variation at the daily and weekly level, whereas the survey-based data capture monthly variation.

Figure 6b examines how our series performs in matching national statistics on trends across sectors. For illustrative purposes, 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. For comparison, we also examine trends in food services employment based on timesheet data from Homebase, a dataset used to examine employment trends in the COVID recession in many studies. Homebase exhibits a much more rapid and much larger decline in employment than the other series and does not track national benchmarks very closely, perhaps because it reflects a selected group of clients or because data obtained from time clocks exhibits different patterns in the recession.

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 the vast majority of 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), we find that wage rates have remained unchanged through the COVID shock for workers who retained their jobs. 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 (Appendix Figure 11).

Heterogeneity by Wage Rates. Figure 7 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 very 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 – which is available with a shorter lag than payroll-based employment data – we construct a prediction of employment rates for low-wage workers up to August 23 as described in Section II.C (shown by the dashed line in Figure 7). The time sheet data suggest that the rate of recovery remained slower in August, and as a result employment levels for low-wage workers were likely to remain well below baseline levels even at the end of August.

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. This may be because high-wage workers were able to adapt to remote work and alternative tasks more rapidly or because low-wage workers were particularly like to be employed in sectors that experienced sharp losses in revenue (e.g., retail and food services). Next, we analyze why

employment rates for low-wage workers fell so much by returning to the geographic heterogeneity in spending changes and business revenue losses examined above.

Heterogeneity Across Areas. 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 work in high-rent areas than low-rent areas in all groups except workers employment at the very largest firms. 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. Because businesses located in high-rent areas lay off more workers and account for a greater share of employment to begin with, they account for a large share of the total loss in employment among low-income workers. 36% of the total loss in employment observed in the Earnin data comes from business located in the top-quartile of ZIP codes by rent; 11% comes from businesses located in the bottom quartile. The much flatter gradient among large firms is to be expected given that the largest firms' markets presumably extend well beyond where they are located, and hence local rents play less of a role in determining their revenues and employment outcomes.[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

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.

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.

of the gradient is attenuated as expected given the coarser geographic measure. Table 2 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. 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]

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.

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 current COVID-induced 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 have lost their jobs but are not actively searching for a new job in the midst of the pandemic. In the three prior recessionary periods, the labor force continued to grow by 0.3% to

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. 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 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 U.S.

0.8% annually.

We therefore focus on the decline in employment rates.

41. Unlike our analyses of private data, the publicly released unemployment claims data do not allow us disaggregagate 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.

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 hours of work 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 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 4, 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.

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.

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. 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.[42]

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

42. 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.

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 5 lists the control states we use for each date. We then calculate unweighted means of the outcome variables in the control 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 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).

Still, despite these significant effects, the reopenings account for a relatively small share of the 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, showing that early reopenings did not play an important role in explaining why some states had stronger economies during the beginning of the recovery than others.[43] These results are consistent with the findings of Lin and Meissner (2020), who use a state-border discontinuity design and find no impact of stay-at-home orders on job losses.

43. 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?[44] 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.[45] 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 36% 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 nearly 20 percentage points within a week. Spending among top-income-quartile households increased as well, but by only about 9 percentage points. This simple analysis suggests that the

44. 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.

45. 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.[46]

To estimate the causal effect of the stimulus payments more precisely, we use a regression discontinuity estimator with the daily spending data.[47] 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.[48] 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 12d 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).[49] 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

46. 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.

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

48. 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.

49. 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 as 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, payroll has recovered only slightly, which is a surprising contrast with the sharp recovery of small business revenues. 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. In a frictionless model where businesses and workers could costlessly reallocate their capital and labor to other sectors, this 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.

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 direct 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 levels (relative to pre-crisis levels), providing an incentive for small businesses to keep employees on payroll.

How effective was the PPP program in increasing employment, particularly among low-income workers? We study this question by 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.[50] 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]

We estimate the causal effects of the PPP 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. We compare the evolution of employment in eligible vs. ineligible firms in Earnin and Kronos data, where we are able to observe not just establishment employment size, but also parent employer size (which determines PPP eligibility).

Figure 15a plots the average change in employment rates (inferred from payroll deposits) relative to January for firms in the Earnin data employing 500 or fewer employees, which were eligible

50. 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 Reference USA data (Panel A) and employees in Earnin data (Panel B), in both cases 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.

51. We find no differences in employment trends below vs. above the 500 employee threshold in the food services sector as well (Appendix Figure 20), consistent with our results below.

for PPP loans, vs. firms employing 501-1,000 employees, which were largely ineligible for PPP loans. 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. Before April 3, trends in employment are similar among eligible vs. ineligible firms, showing that larger businesses are likely to provide a good counterfactual for employment trends one would have observed in smaller firms absent the PPP program. After April 3, we see that employment rises slightly in the treated (<500 employees) group relative to the control group, showing that the PPP program increased employment in smaller firms. However, the magnitude of this increase is quite modest, with a difference of approximately 2 percentage points as of mid-May; notably, even after PPP implementation, employment levels for low-wage workers remain well below baseline. Figure 15b plots the change in employment pre-COVID vs. post-COVID by firm size. 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. Figures 15c and 15d show that we observe very similar results in 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 Paycheck Protection Program.

In Table 5, we quantify the impacts of the PPP by estimating a set of regression models. Panel A shows difference-in-difference estimates of the effect of Paycheck Protection Program elibility in the Earnin data, clustering standard errors at the establishment ID level. Column 1 shows an estimated effect of PPP eligiblity of 1.67 percentage points relative to the January 2020 level of employment. Column 2 shows that we continue to find little effect of the PPP when we omit the smallest firms (<100 employees) in order to construct a sample with more comparable firms. Columns 3 and 4 show that the estimated effect of PPP eligibility is similarly low in both the lowest-rent (bottom quartile) ZIP codes and highest-rent (top quartile) ZIP codes. Panel B replicates these results in the Kronos data; as in the Earnin data, the point estimates for the effect of PPP eligibility are small.

One caveat in interpreting these results is that the measurement of firm size used in the Earnin sample, based on data provided by ReferenceUSA, may not correspond precisely to the firm sizes used by the Small Business Administration to determine PPP eligibility. For this reason, a simple difference-in-difference approach will lead to attenuated estimates of the effect of PPP eligibility. To account for this attenuation, we use publicly available data released by the SBA on PPP recipients

to assess the precision of the classification of eligibility in the Earnin data. The SBA has released identifiable data, including firm names and ZIP codes, of PPP recipients receiving over $150,000 in loans; these loans represent 72.8% of total PPP expenditure. Of the roughly 660,000 PPP recipients of these loans, we successfully merge around 450,000 (69%) to ReferenceUSA data; for details on this merge, see Appendix D. In this matched subset, we find that mean PPP expenditure on firms with 500 employers or fewer is around $3,480 per worker, whereas mean PPP expenditure on firms with over 500 workers is around $800 per worker. For this reason, raw estimates of the effect of PPP eligibility in the Earnin data will tend to be attenuated relative to the true causal effect of the PPP.

To account for this measurement error, we make two assumptions which permit us to recover the effect of PPP eligibility. First, we assume that the degree of misclassification of eligibility among identifiable PPP recipients matches the degree of misclassification of eligibility in the broader Earnin sample; this permits us to use the matched SBA/ReferenceUSA data to make inferences about misclassification of firm size more generally.[52] Second, we assume that the effect of the PPP on employment relative to January 2020 is linear in loan amount received per worker; intuitively, this permits us to use mean amount received per worker as a sufficient statistic for treatment intensity. Under these assumptions, we can rescale our raw (attenuated) estimate of the effect of PPP eligibility by multiplying by the ratio of true PPP spending per worker among eligible firms (approximately $\$521B/70M \approx \$7,440$ per worker) to the difference in PPP spending per worker among "eligible" vs. "ineligible" firms (approximately $\$3,480 - \$800 = \$2,600$ per worker in matched PPP loans, which we scale up to $\$5,175$ per worker when dividing by 0.69×0.728 to account for non-identifiable PPP loans and PPP loans that are not matched to the ReferenceUSA data, under our maintained assumption that the degree of misclassification of eligibility among identifiable PPP recipients matches the degree of misclassification of eligibility in the broader Earnin sample). The second row of Table 5 shows these rescaled estimates. The estimated effect of PPP eligibility remains small in magnitude, even after this rescaling.

In sum, although we cannot rule out a small positive employment effect of the PPP (of e.g., 4 pp on low-wage employment rates), it is clear that the program did not restore the vast majority of jobs that were lost following the COVID shock. At the upper bound of the 95% confidence interval, our rescaled baseline estimates in both the Earnin and Paychex imply that the cost per job saved

52. As identifiable PPP recipients receive larger loans, they may tend towards being larger firms, such that the SBA/ReferenceUSA data overstates the degree of misclassification. In these case, our estimates will represent an upper bound on the effect of PPP eligibility.

by the PPP exceeds $300,000. We conclude that the cost of saving jobs by providing forgiveable loans to small businesses is quite high.

Why did the PPP have relatively small effects on employment rates? One potential explanation is that the loans were taken by firms that intended not to layoff many employees to begin with, i.e. firms that were inframarginal recipients of loans. Consistent with this, 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.

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, sectores, 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.

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.*

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.*

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 (3): 571–630.

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.*

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.*

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.*

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."

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.*

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.*

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.*

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.

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.

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.*

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.

Granja, João, Christos Makridis, Constantine Yannelis, and Eric Zwick. 2020. "Did the Paycheck Protection Program Hit the Target?" *NBER Working Paper No. 27095.*

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.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (7): 1737–72.

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.*

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 Lalé, and Lien Ta. 2020. "The Impact of COVID-19 on U.S. Employment and Hours: Real-Time Estimates with Homebase Data."

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.

Lin, Zhixian, and Christopher M Meissner. 2020. "Health vs. Wealth? Public Health Policies and the Economy During Covid-19." *NBER Working Paper No. 27099.*

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.*

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.*

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.*

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127 (5): 2505–2558.

Online Appendix

A Key Dates and Geographic Definitions

In this appendix, we provide additional details about the key dates in the COVID-19 crisis as well as geographic definitions 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 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 from the New York Times stay at home order data, available here.
- Stay-at-home order ends: Sourced and verified from the New York Times reopening data, available here, and hand-collection where needed. Defined as the date at which the state government lifted or eased the executive action 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 from the New York Times reopening data, available here. 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 6. 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.

B Consumer Spending Series Construction

This appendix provides further details on our construction of consumer spending series in Affinity data.

Construction of Consumer Spending Series. We receive data from Affinity Solutions at the county x ZCTA income quartile x two-digit NAICS code x week level. In the raw data we receive, cells with fewer than five unique cards transacted are masked. ZCTA income quartile and county are both determined by the cardholder's ZIP code of residence, rather than the merchant's address. We use these raw data to calculate daily total spending in each industry and for each income quartile, for each geographic level. We then take a seven-day lookback moving average of daily

spending. For both 2019 and 2020, we divide this moving average measure of daily spending by mean daily spending over the period January 4-31. We then seasonally adjust the series by dividing each calendar date's 2020 value by its corresponding value from 2019.

Treatment of Structural Breaks Arising from Changing Coverage. Affinity's coverage may change over time due to changes in the credit card processesors providing data to Affinity. 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. When conducting this test, we exclude the period from March 10 to March 31 to avoid misclassifying structural breaks arising due to the decline in consumer spending due to COVID-19 as structural breaks in Affinity data. We identify 950 counties with a structural break where the p-value for 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 $t+1$ using other counties in the same state. Where there is a change in coverage, we adjust the series to be in line with the lower level of coverage. For example, consider a county which has n cards in weeks $1, \ldots, t-1$, and $3n$cards in week t. In week $t-2$, the county is assigned its reported value of $n$cards. 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, if counties in the rest of the state had a 10% increase in the number of cards, we would impute the county with a structural break to have a level of $(1.10) \times (1.05n) = 1.155n$. 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 the $t+1$value in the county with the structural break to have a level of (0.92) * (1.155n) = 1.0626n. In subsequent weeks, state-level flucuations 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 occurs in 98 cases.

We implement a structural-break correction for three counties: Philadelphia County, Pennsylvania (FIPS code 42101); Washington, District of Columbia (11001); and 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.

Finally, 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.

Spurious Increases in Consumer Spending. There is a large spike in national consumer spending between January 15th, 2019 and January 17th, 2019 that appears to be anomalous: it does not arise in other data series, and is not repeated in 2020. 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 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.

Assignment of ZCTA Income Quartiles. We assign an income quartile for each ZCTA using 2014-2018 ACS estimates of median household income at the ZCTA level. We use population weights when defining quartile thresholds so that each income quartile has the same number of individuals. When displaying these data on the Economic Tracker, we display the first ZCTA income quartile as low-income spending, the pooled second and third ZCTA income quartiles as middle-income spending, and the fourth ZCTA income quartile as high-income spending.

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 Womply 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 A 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 A. 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.[53]

Delayed Processing of Payments. Due to differences in the speed at which data providers share their data with Womply, the most recent date as of a given data refresh is invariably incomplete. If left unattended, 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 until four days

53. 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.

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.

Earnin Analysis sample. 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 data, 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 assume that an individual has been employed for every day of their paycycle. To account for the delay in receipt of paychecks, we shift the Earnin series back by one week.

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 A to aggregate the series to the county, state and national level.

ZCTA-level low-income employment series. As ZCTA is not observed in Paychex and Intuit, we separately construct ZCTA-level employment using the Earnin data only. We construct our analysis sample as above. To account for the noisier data at the ZCTA-level, we norm the ZCTA-level changes relative to a pre-period of January 5 - March 7. We suppress estimates for ZCTAs with fewer than 100 worker-days observed over this period.

PPP Analysis sample. We use data from Earnin, Paychex and Kronos to examine the effects of the Paychex Protection Program. In each case, we form our analysis sample by dropping firms whose firm size is missing. In Paychex, where firm entry is observed, we additionally exclude firms that are new to the data in 2020. In Earnin data, parent employer size is computed by matching Earnin employers to ReferenceUSA firms using the process above. In Paychex data, parent employer size is computed by matching Paychex clients to Dun & Bradstreet firm records. In Kronos, parent firm size is defined using March 2020 employee counts.

Earnin stimulus data. 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.

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.[54] We also report daily state-level data on the number of tests performed per day per 100,000 people from the COVID Tracking Project.[55] 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.[56] We report these values as changes relative to the median value for the corresponding

54. 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.

55. 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.

56. 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.

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

Table 1: Changes in Consumer Spending by Sector and Income Quartile

	Total Spending	Total Spending	In-Person Services	Remote Services	Durable Goods
	(1)	(2)	(3)	(4)	(5)
Low-Income	-0.28	-0.41	-0.63	-0.30	-0.38
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Q2	-0.26	-0.39	-0.63	-0.29	-0.34
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Q3	-0.26	-0.35	-0.59	-0.26	-0.29
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
High-Income	-0.31	-0.35	-0.53	-0.26	-0.28
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
Dollars Q1	-$1.5 bil	-$2.2 bil	-$3.4 bil	-$1.6 bil	-$2.0 bil
	($49 mil)	($57 mil)	($103 mil)	($82 mil)	($109 mil)
Dollars Q4	-$4.0 bil	-$4.5 bil	-$6.8 bil	-$3.3 bil	-$3.6 bil
	($121 mil)	($136 mil)	($170 mil)	($196 mil)	($213 mil)
Frac Q1	0.15	0.17	0.16	0.17	0.18
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
Frac Q4	0.39	0.34	0.32	0.34	0.32
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)
N	5,032	5,032	5,032	2,450	3,325
County FE		X	X	X	X

Notes: This table shows changes in spending by sector and income quartile using Affinity solutions spending data (seasonally adjustment described in section II.A).

Table 2

Changes in Business Revenue and Employment in Small Businesses

Panel A: Changes in Business Revenue

Outcome:	% Change in Changes in Small Business Revenue				
	(1)	(2)	(3)	(4)	(5)
Median 2BR Rent (per thousand dollars)	-13.00	-7.69	-14.07	-10.99	-0.06
	(0.38)	(0.51)	(0.70)	(0.73)	(0.01)
Log of Density of High Wage Workers		-1.16		-2.11	-1.30
		(0.08)		(0.13)	(0.10)
County FEs			X	X	
State FEs					X
Level of Observation	ZCTA	ZCTA	ZCTA	ZCTA	County
Observations	16396	15907	16396	15907	2759

Panel B: Changes in Low-Wage Employment

Outcome	% Change in Changes in Low-Wage Employment				
	(1)	(2)	(3)	(4)	(5)
Median 2BR Rent (per thousand dollars)	-12.73	-9.56	-11.58	-10.00	-1.14
	(0.53)	(0.71)	(0.97)	(1.02)	(1.51)
Log of Density of High Wage Workers		-0.88		-1.25	-0.93
		(0.11)		(0.19)	(0.16)
County FEs			X	X	
State FEs					X
Level of Observation	ZCTA	ZCTA	ZCTA	ZCTA	County
Observations	13440	12885	13440	12885	2340

Notes: This table shows OLS regressions of average percentage changes in small business revenue by ZCTA (using Womply data) and small business low-wage employment (Earnin at the ZCTA level; Combined Employment Series at the County level) on average ZCTA 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 ZCTA 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	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dif-in-Dif Effect of Reopening:	1.04	0.18	1.52	2.75	0.73	1.73	3.67	5.01
	(0.93)	(1.34)	(0.66)	(0.99)	(0.51)	(1.40)	(1.35)	(1.77)
State-Week Observations:	208	318	200	264	152	248	248	324
Weeks on either side of reopening included:	2	3	2	3	2	2	2	3

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 (2), (4), and (8) 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 11B, 11C, and 11D respectively.

Table 4
Regression Discontinuity Estimates of Stimulus Payments on Spending

	(1) Q1 ZCTAs	(2) Q1 ZCTAs	(3) Q4 ZCTAs	(4) Q4 ZCTAs
Panel A: Impact of Stimulus Payments on Consumer Spending				
Outcome:	Spending			
RD Effect of Stimulus:	0.26	0.38	0.09	0.16
	(0.07)	(0.10)	(0.04)	(0.05)
Window:	April 1 - April 30	April 7 - April 21	April 1 - April 30	April 7 - April 21
Panel B: Impact of Stimulus Payments on Revenue				
Outcome:	Revenue			
RD Effect of Stimulus:	0.18	0.21	0.01	-0.08
	(0.10)	(0.17)	(0.06)	(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 ZCTAs in the bottom quartile of the distribution of ZCTA median household income (based on data from the 2014-2018 ACS), while Columns (3)-(4) look cardholders living in the quartile of ZCTAs 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 ZCTAs in the bottom quartile of the distribution of ZCTA 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 ZCTAs with the highest mediantwo 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 Paycheck Protection Program (PPP) on Employment

	Outcome Variable: Employment			
	Baseline Estimate (1)	Min. 100 Employees (2)	Low-Rent ZCTAs (3)	High-Rent ZCTAs (4)
A. Earnin				
DD Estimate	1.67	-0.54	2.46	1.36
	(0.74)	(0.84)	(1.82)	(1.36)
DD Estimate (Rescaled)	2.39	-0.77	3.52	1.94
	(1.06)	(1.20)	(2.60)	(1.94)
Specification				
Reweighted on 2-Digit NAICS	X	X	X	X
Firms with 100 or More Employees Only		X		
Low-Rent ZCTAs Only			X	
High-Rent ZCTAs Only				X
B. Kronos				
DD Estimate	.35	1.25	1.99	-2.01
	(1.39)	(1.45)	(2.21)	(2.87)
Specification				
Reweighted on 2-Digit NAICS				
Firms with 100 or More Employees Only		X		
Low-Rent ZCTAs Only			X	
High-Rent ZCTAs Only				X

Notes: This table shows difference-in-difference estimates of the effect of PPP eligibility on employment. The outcome variable is employment at the establishment ID level, expressed as a percentage change relative to a pre-period. All regressions are restricted to firms with a parent employer size of between 1 and 1,000 employees. Panel A shows regressions in Earnin data at the establishment ID x date level, clustered on establishment ID. In Earnin data, the pre-period is defined as January 4-31. The first row shows the coefficient on an interaction term for PPP eligibility and the date being after April 3 in a regression of change in employment on an indicator for PPP eligibility, an indicator for the date being after April 3, and the interaction between these variables. Row 2 rescales these estimates to account for mismeasurement of firm sizes; see section 4.3 for details. Column (1) shows the estimated effect of PPP eligibility, reweighting by two-digit NAICS to match the January 2020 composition of two-digit NAICS codes. Column (2) restricts the sample to firms with at least 100 employees. Columns (3) and (4) restrict the sample to establishments in low-rent (bottom quartile) and high-rent (top-quartile) ZCTAs, respectively. . Panel A replicates Panel A for regressions in Kronos data at the establishment ID x week level, clustered on establishment ID. As Kronos data is first observed in March, the pre-period in Panel B is the week of March 1.

Appendix Table 1
Hourly Wage Rates By Industry

		2019 BLS Wages			Median in Private Datasets			
NAICS Code	NAICS Description	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 - March10th). 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)
Panel A: Income		
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

	Zearn Users	US K-12 Students
Panel B: School Demographics		
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	-0.0089	-0.0121
	(0.0006)	(0.0012)	(0.0039)
Median Home 2BR Rent		-0.0065	
		(0.0017)	
Controls:			
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 ZCTA code (using data from Affinity Solutions) on average workplace ZCTA code median two-bedroom rent. Standard errors are reported in parentheses. Workplace ZCTA 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 11B-11D and Appendix Figure 20.

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

		QCEW Industry Shares		Industry Shares in Private Data Sets			
				Earnin		Paychex	Intuit
NAICS Code	NAICS Description	QCEW All Establishments (1)	QCEW Small Establishments (2)	(3)	Homebase (4)	(5)	(6)
11	Agriculture, Forestry, Fishing and Hunting	0.84	1.04	0.19		0.65	0.83
21	Mining, Quarrying, and Oil and Gas Extraction	0.55	0.43	0.18		0.22	0.08
22	Utilities	0.44	0.29	0.48		0.17	0.17
23	Construction	5.72	7.62	1.14		6.61	7.32
31-33	Manufacturing	10.27	5.16	6.08		8.18	2.24
42	Wholesale Trade	4.72	5.99	4.21		5.60	1.66
44-45	Retail Trade	12.48	14.06	24.08	11.28	7.82	4.65
48-49	Transportation and Warehousing	4.30	2.82	5.74	0.87	2.21	1.58
51	Information	2.29	1.64	3.88		1.52	0.94
52	Finance and Insurance	4.83	4.60	7.26		3.31	1.72
53	Real Estate and Rental and Leasing	1.71	2.90	1.93		3.05	1.85
54	Professional, Scientific, and Technical Services	7.63	8.97	4.69	2.78	12.34	11.84
55	Management of Companies and Enterprises	1.93	0.79	0.36		0.45	0.15
56	Administrative Support	7.25	5.30	4.66		7.06	5.00
61	Educational Services	2.39	1.53	4.45	3.62	2.45	1.18
62	Health Care and Social Assistance	16.16	13.16	15.50	5.34	14.92	5.71
71	Arts, Entertainment, and Recreation	1.78	1.64	1.65	2.07	2.31	1.30
72	Accommodation and Food Services	11.04	15.60	10.99	49.17	11.51	2.61
81	Other Services (except Public Administration)	3.57	6.21	2.52	1.83	8.79	5.96
99	Unclassified	0.11	0.24		23.04	0.83	43.2

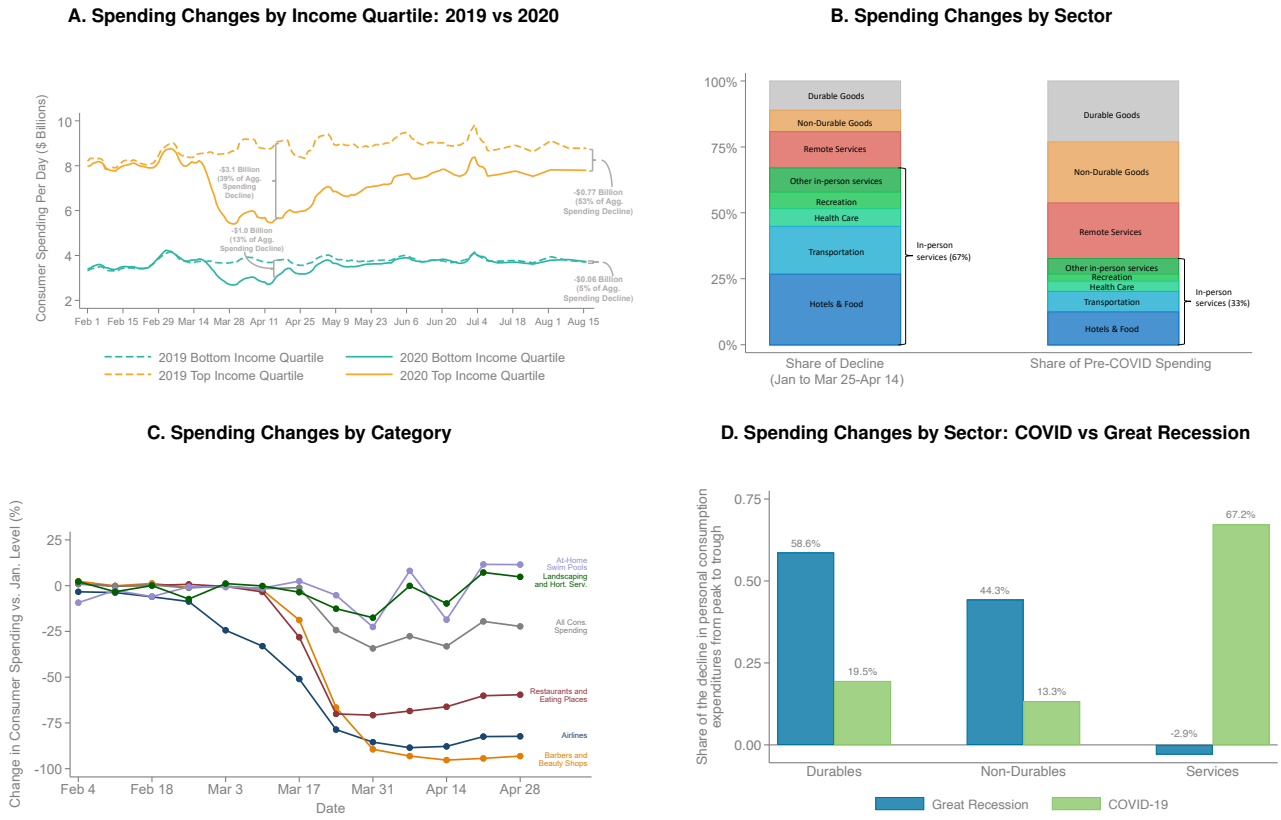
Notes: This table reports the NAICS two-digit industry mix for four private employment-based datasets compared with the Quarterly Census of Employment and Wages (QCEW), an administrative dataset covering the near-universe of firms in the United States. Columns (1) - (6) indicate the share of employees 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 Earnin, Homebase, and Paychex, we use January 2020 employment to do so. For Intuit, we constuct employment shares using January 2019 employment, because we do not observe 2020 Intuit data. We define employment in Earnin as the total number of worker-days in the month. 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 as the total number of employees recorded by Paychex clients 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



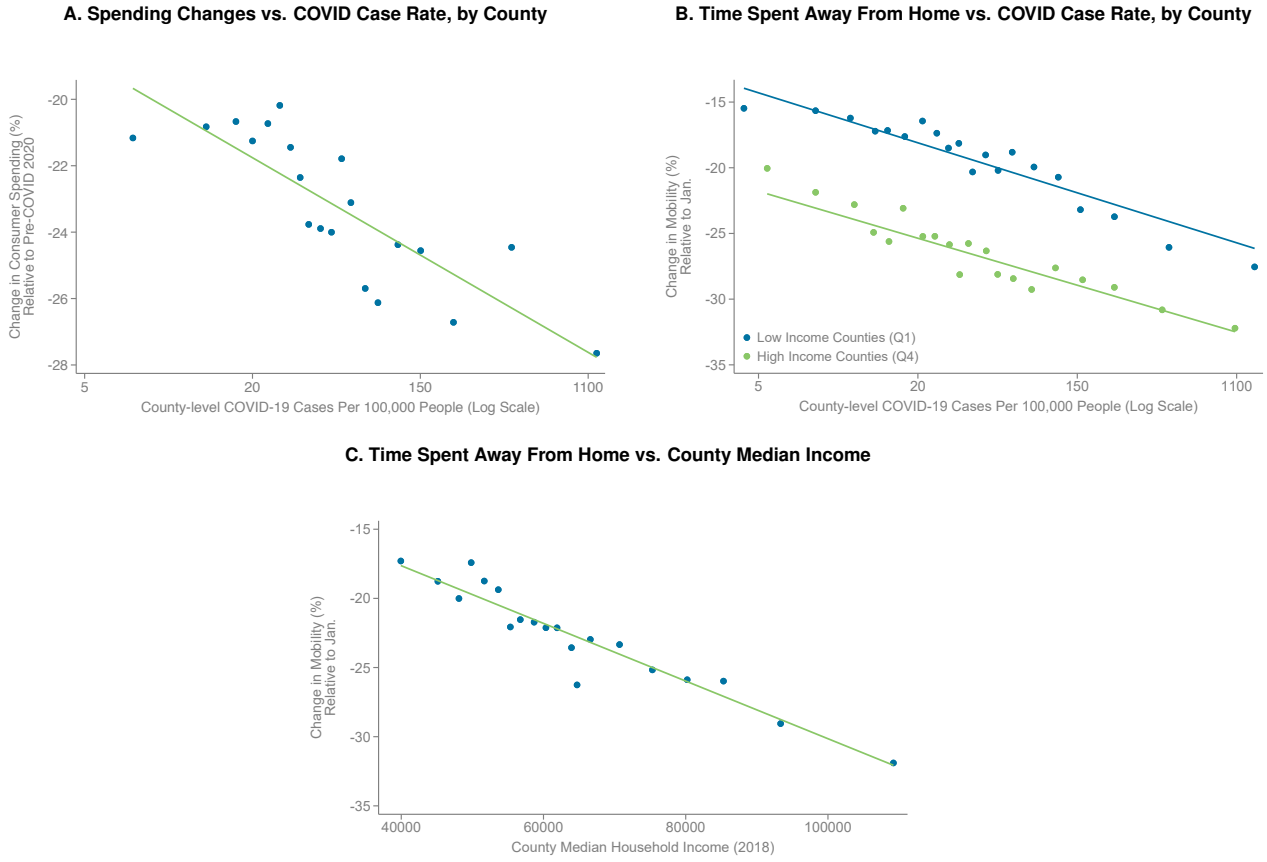
Notes: This figure compares changes in consumer spending in national income and product accounts (NIPA) data to measures of consumer spending recorded on debit and credit cards. Panel A summarizes NIPA data (Tables 1.1.2, 1.1.6 and 2.3.2), comparing Q4 2019 to Q1 2020. The first bar shows the seasonally adjusted change from Q4-2019 to Q1-2020 in real GDP in chained (2012) dollars (-$247.3B). In parentheses under the first bar is the compound annual growth rate corresponding to this change in real GDP (-5.0%). 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 (-$247.3B) 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 (-$138.2B). 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 (-$229.7B) 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 monthly spending in 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. Panel B restricts to specifically retail trade sectors (NAICS code 44-45) excluding motor vehicles (NAICS code 441) and gas (NAICS code 447). Panel C restricts to food services (NAICS code 722) in the MRTS and food services (NAICS code 722) as well as accommodations (NAICS code 721) in Affinity Solutions. Both series are normalized relative to January 2020 spending (Jan 1 - Jan 31). The bottom right corner of each panel displays the root mean squared error (RMSE) of a regression of indexed MRTS monthly spending on indexed Affinity Solutions monthly spending.

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 ZCTAs (ZIP Code tabulation areas defined by the Census Bureau) in the top and bottom quartiles of the distribution of ZCTA median household income. ZCTA median household income quartiles are constructed using population-weighted 2014-2018 ACS median household income. We scale the 2020 (2019) series 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 impute the value plotted for February 29, 2019 with the average of February 22, 2019 and March 7, 2019. Panel B disaggregates spending changes by sector. The first 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).

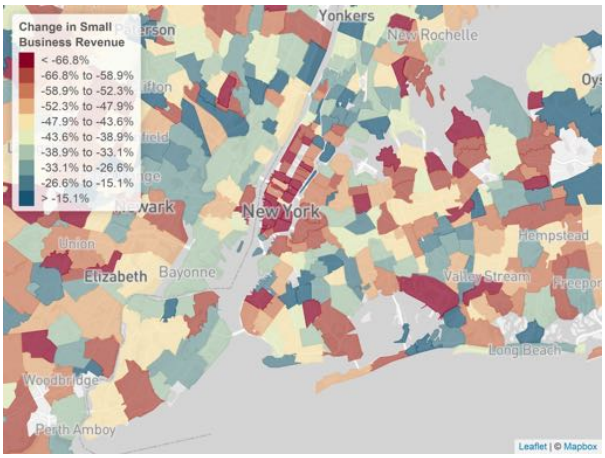
FIGURE 3: Association Between COVID-19 Incidence, Spending, and Time Outside Home



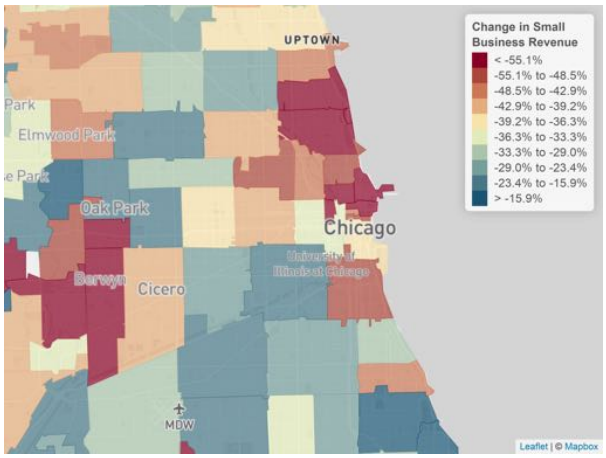
Notes: This figure presents three county-level binned scatter plots showing the association between COVID-19 incidence, spending, and time spent outside home. To construct each binned scatter plot, we divide the x-axis variable into twenty equal-sized bins weighting by the county's population, and plot 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 (using data from Affinity Solutions) in a county from the base period (January 8 - January 28) to the two-week period from April 1 - April 14 vs. the county's COVID case rate over the two week period from April 1 - April 14. Panel B presents a second binned scatter plot of the change in time spent outside the home in a 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. Low-income and high-income counties are defined as those with median household income in the bottom 25% and top 25% of all counties respectively, as measured in the 2012-2016 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.

FIGURE 4: Changes in Small Business Revenues by ZIP Code

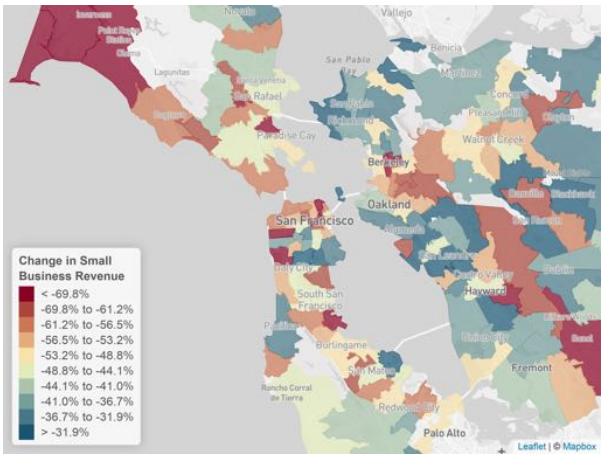
A. New York



B. Chicago

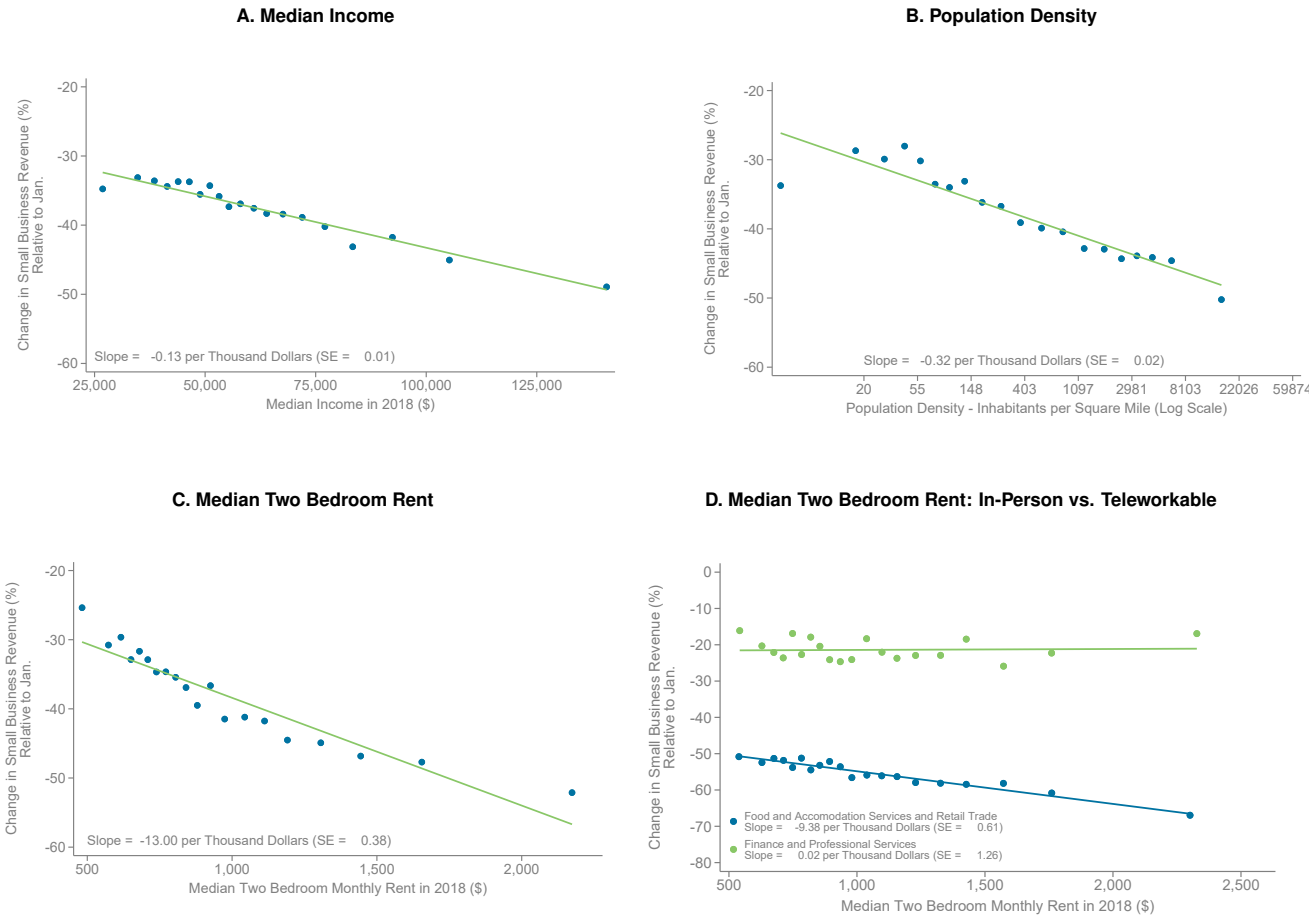


C. San Francisco



Notes: This figure plots percentage small business revenue declines during the COVID crisis by ZIP code in the MSAs corresponding to New York City, San Francisco, and Chicago. The change in revenue is determined in each ZCTA (ZIP code tabulation area) using data from Womply. First, normalized weekly revenue is calculated by dividing weekly revenue by the average weekly revenue from January 1 to 28 for both 2019 and 2020 and by dividng the 2020 proportion by the 2019 proportion for each week. Second, normalized revenue is regressed on a dummy variable indicating whether the week is before March 9, 2020 or after. The changes in revenue mapped in this figure are the coefficients on the dummy variables for each ZCTA. The signal variance to total variance ratios for the panels are 0.83 (New York), 0.89 (Chicago), and 0.73 (San Francisco). Panel A shows the New York-Newark-Jersey City, NY-NJ-PA MSA. Panel B shows the San Francisco-Oakland-Hayward, CA MSA. Panel C shows the Chicago-Naperville-Elgin, IL-IN-WI MSA. 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.

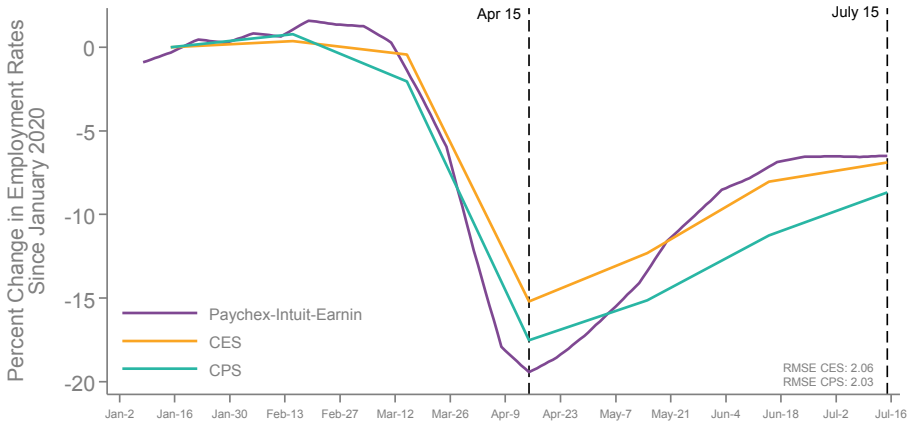
FIGURE 5: Changes in Small Business Revenues vs. ZIP Code Characteristics



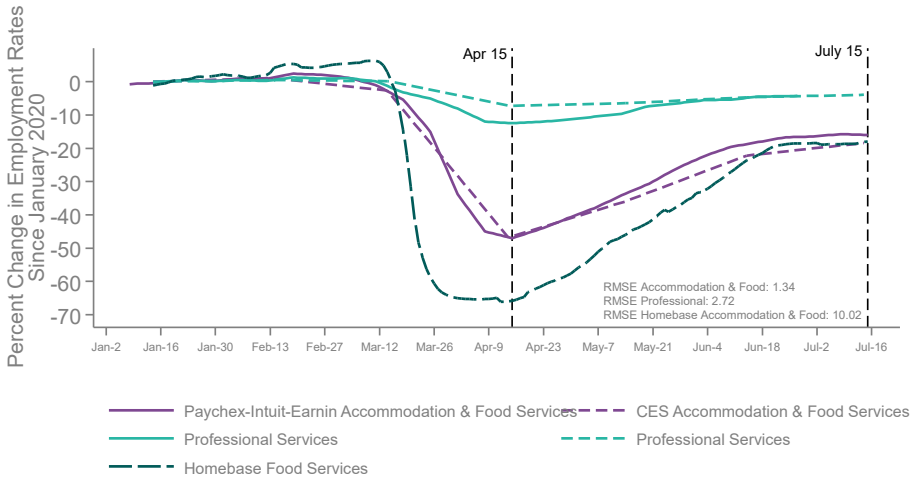
Notes: Panels A-C present binned scatter plots showing the relationship between changes in small business revenue using data from Womply vs. various ZCTA-level characteristics. Binned scatter plots are constructed as described in Figure 3. Changes in business revenue in each ZCTA are determined using average normalized weekly revenue, as defined in Figure 4, from March 25 to April 14. ZCTAs where changes are larger than 200% and where variance of normalized revenue exceeded 900 are excluded. To preserve the privacy of firms included in the data, ZCTA-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 ZCTA level, from the 2014-2018 ACS. The x variable in Panel B is the log number of inhabitants per square mile in the ZCTA the 2014-18 ACS. The x variable in Panel C is the median rent for a two-bedroom apartment in the ZCTA 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).

FIGURE 6: Changes in Employment Rates Over Time

A. All Industries

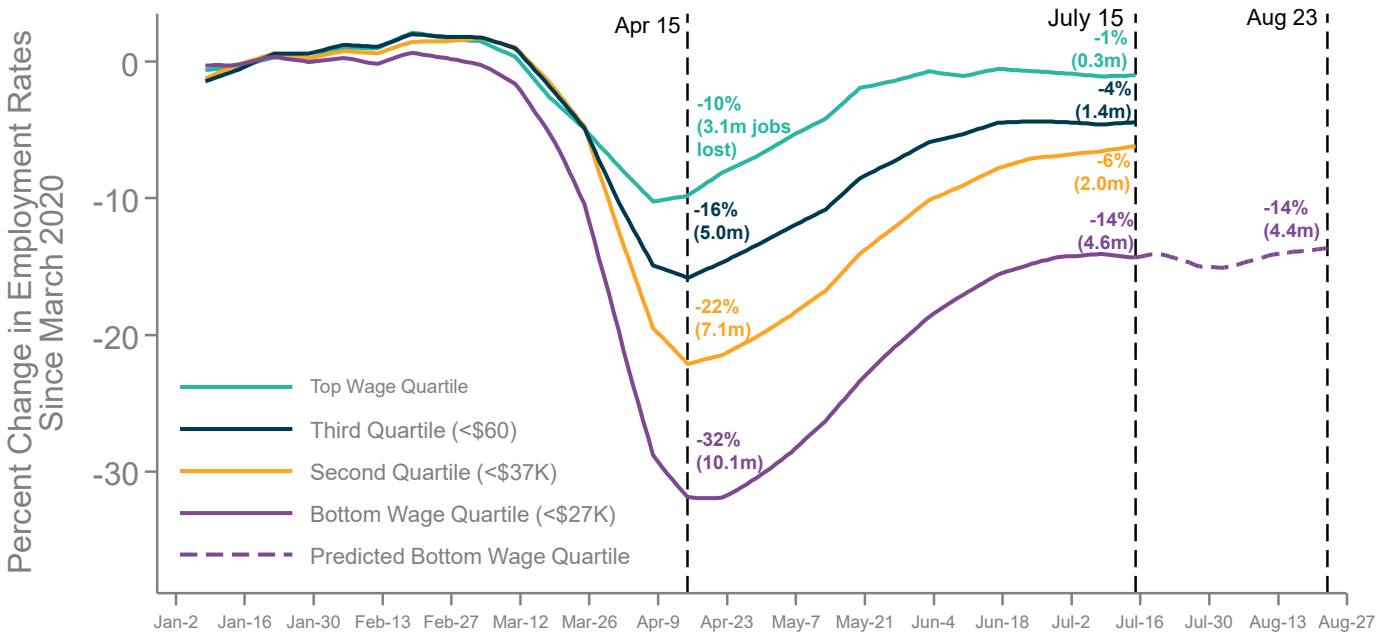


B. Accommodation and Food Services and Professional Services



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 the seven day moving average over the period, January 4 - January 24, 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

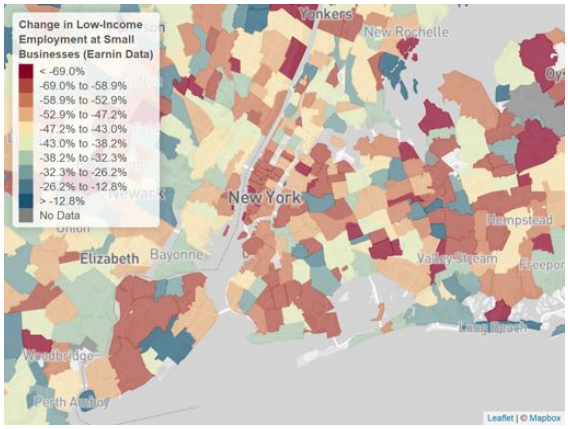
FIGURE 7: Changes in Employment Rates by Income Quartile



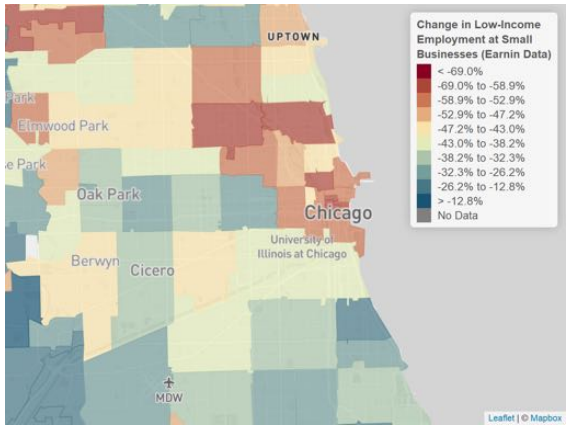
Notes: This figure shows employment changes over time by income quartile relative to January 2020. We combine Paychex, Earnin, and Intuit data to construct a daily private employment series. We then construct an employment index by averaging employment over the prior seven days and then norming to the average value of the seven day moving average over the period, January 4 - January 24, 2020. We use Kronos data to forecast the first income quartile of the combined series so that it extends to August 23. We regress the combined series on the Kronos series for the same date (t), as well as on the Kronos series for up to three observations prior $(t-7), (t-14), (t-21)$. We then use the resulting coefficients on Kronos and its lags to predict the combined series for all dates from July 16 through August 23. For details, see section 2.3. Data sources: Paychex, Intuit, Earnin, Kronos

FIGURE 8: Changes in Employment Rates by ZIP Code

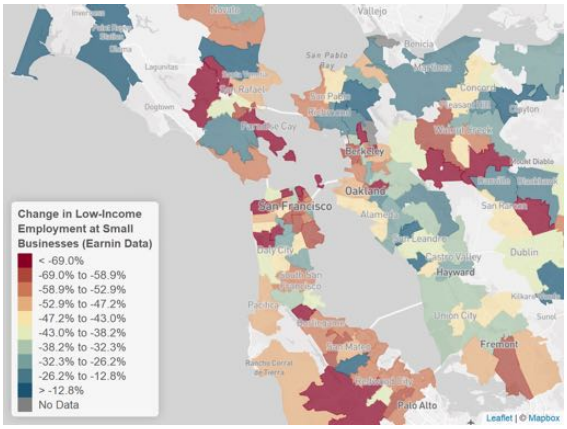
A. New York



B. Chicago

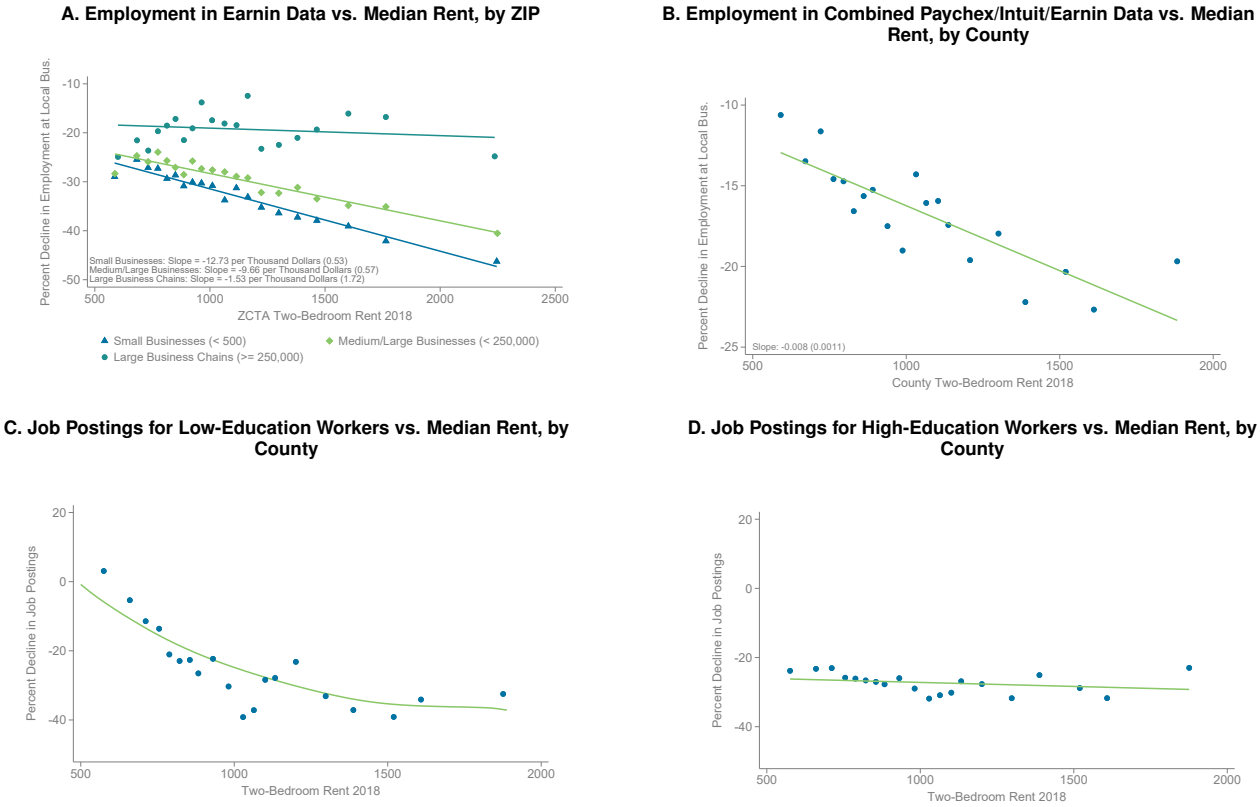


C. San Francisco



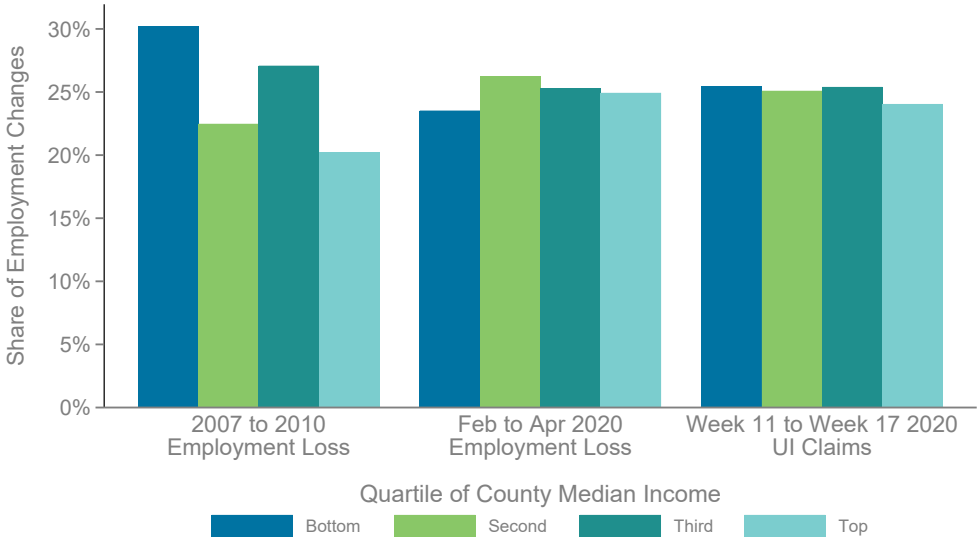
Notes: This figure replicates Figure 4 using changes in employment at small and medium-sized businesses. We restrict to small and medium-sized businesses by limiting our sample to Earnin users in firms with at most 500 employees. For users whose employer cannot be matched to ReferenceUSA data on firm 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 third decile of Earnin employers is roughly 200 employees. We first construct weekly measure of employment by taking the percent change in total employment (inferred from payroll deposits) at the ZCTA level from the period of January 5th to March 7th, 2020 to each week; for details on construction of the employment series, see Appendix D. We then define the percent change in employment in each zip code during COVID-19 as the coefficient of weekly percent change in employment on an indicator for the COVID-19 period (March 22nd - May 5th) from a weekly regression in each zip code. Panel A shows the New York-Newark-Jersey City, NY-NJ-PA MSA. Panel B shows the San Francisco-Oakland-Hayward, CA MSA. Panel C shows the Chicago-Naperville-Elgin, IL-IN-WI MSA. The signal to total-variance ratios are 0.79 in New York, 0.69 in Chicago, and 0.77 in 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.

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



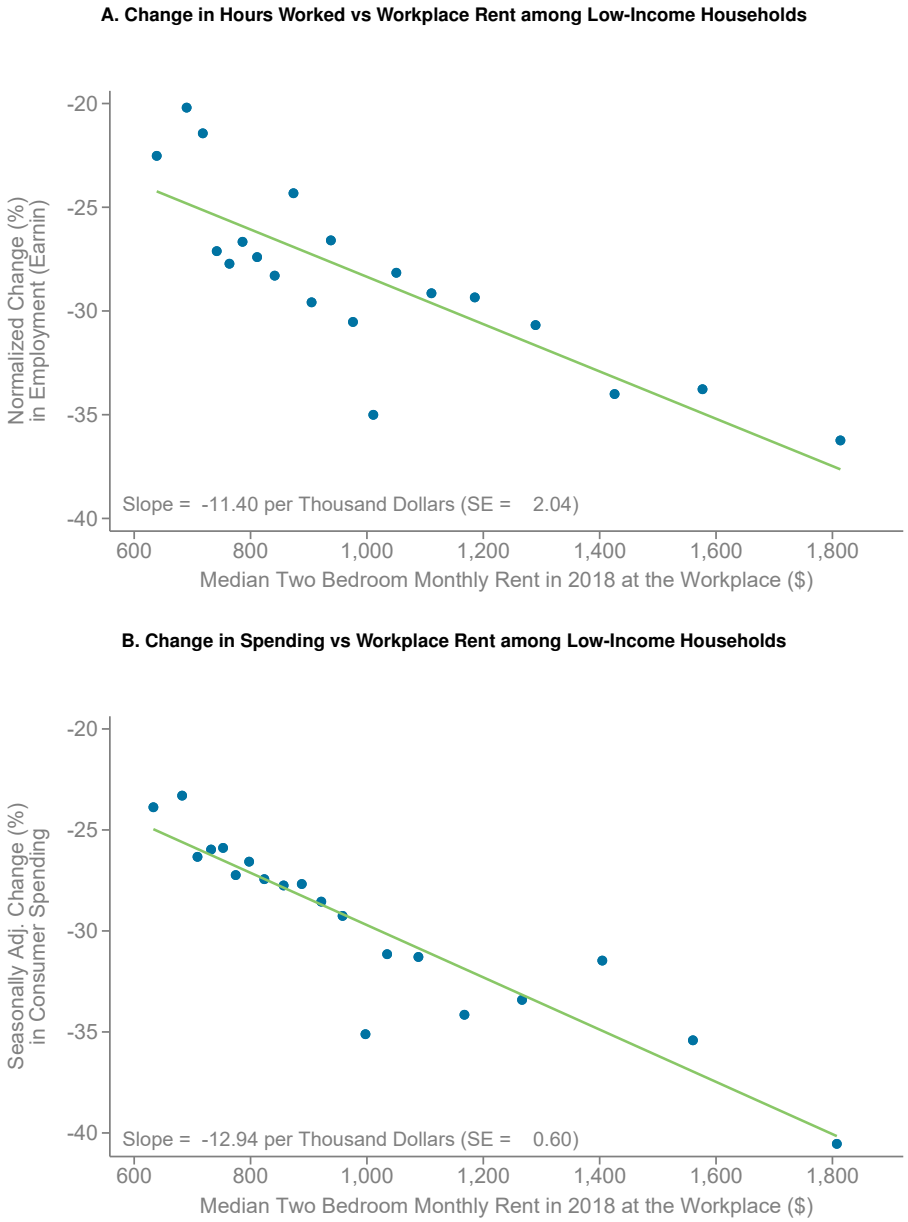
Notes: This figure shows binned scatterplots of the relationship between changes in employment rates and median rents (Panels A-B) and job postings and median rent (Panels C-D), by ZIP code. Binned scatter plots are constructed as indicated in Figure 3 by binning ZCTAs (Panel A) and counties (Panels 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/large firms (between 500 and 250,000 employees), and large business chains (250,000 or more employees). Both panels measure the percentage change in hours/employment from January 15-February 4th, 2020 to April 8-28th, 2020. 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. 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.

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



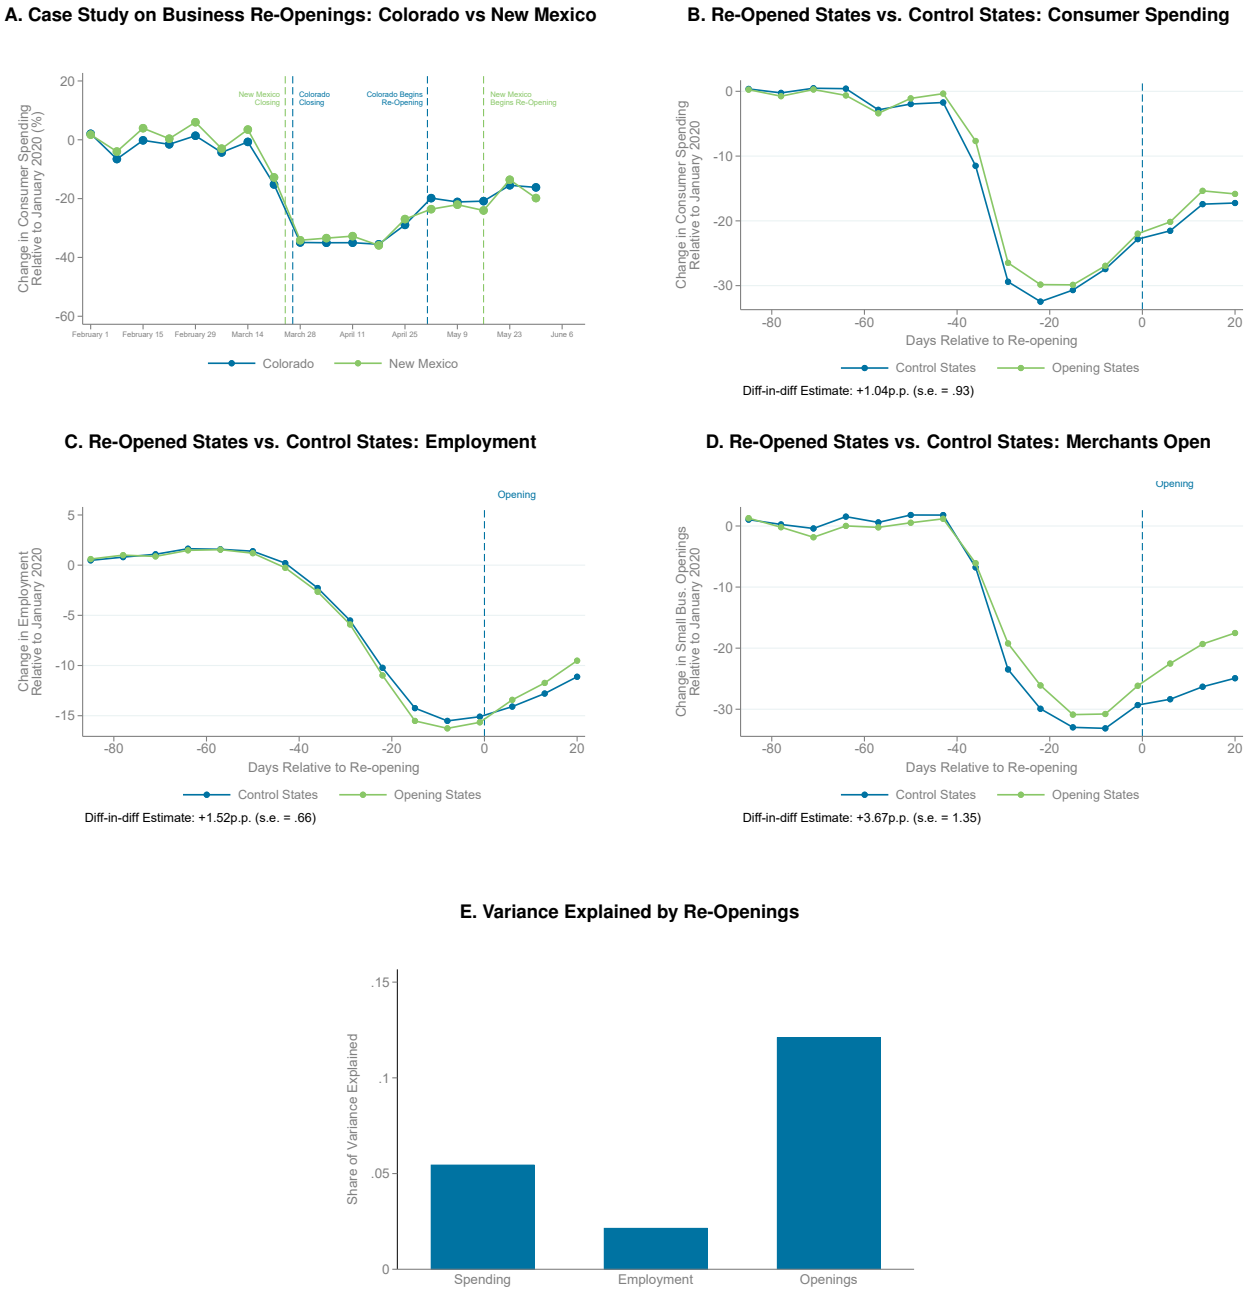
Notes: This figure displays the share of job losses occurring in low vs high income counties in the Great Recession and the COVID recession. To construct the first set of four bars, we first calculate national employment loss between 2007 and 2010 using data from the BLS. We then group counties into (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 Consumer Spending vs. Workplace Rent for Low-Income Households



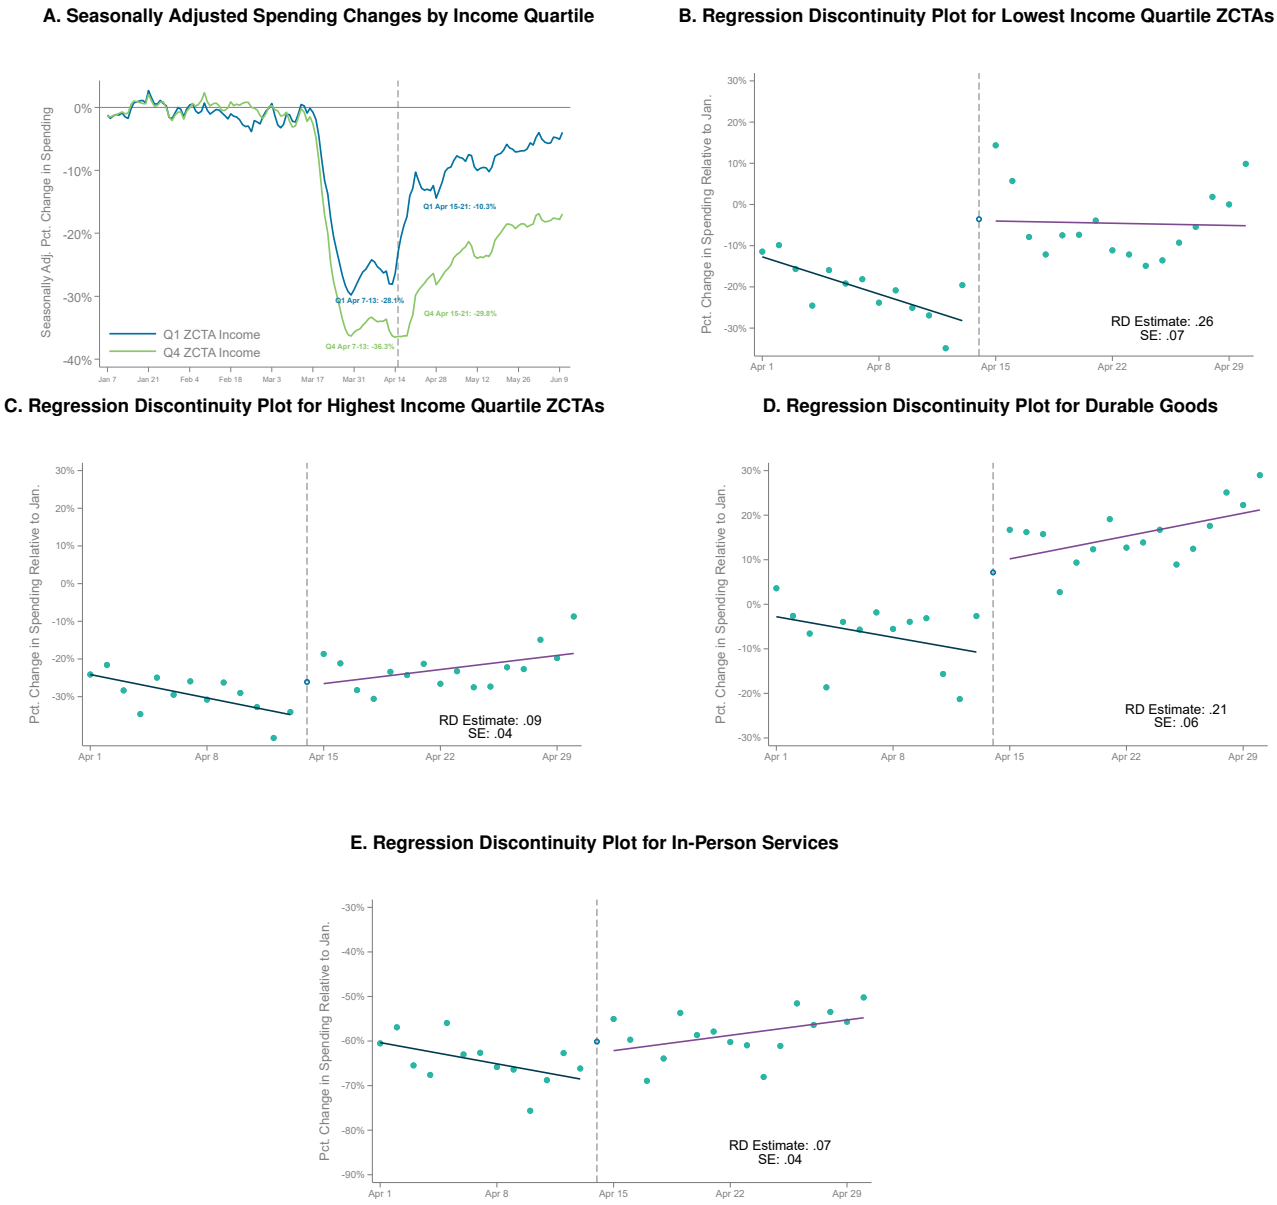
Notes: This figure plots changes in hours worked (Panel A) and consumer spending (Panel B) by ZCTA vs. the average median 2 bedroom rent in the *workplace* ZCTAs of individuals who live in a given ZCTA. The sample is restricted to individuals who live in ZCTAs in the bottom quartile of the household income distribution. We construct the average median 2 bedroom rent variable by combining data on the matrix of home residence by workplace ZCTAs 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 ZCTA of workplace in the LODES data and then collapse workplace rents to each home ZCTA, weighting by the number of jobs in each workplace ZCTA. 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 spending is computed from the period of Jan 5-Mar 7 to Mar 22-April 20, 2020.

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



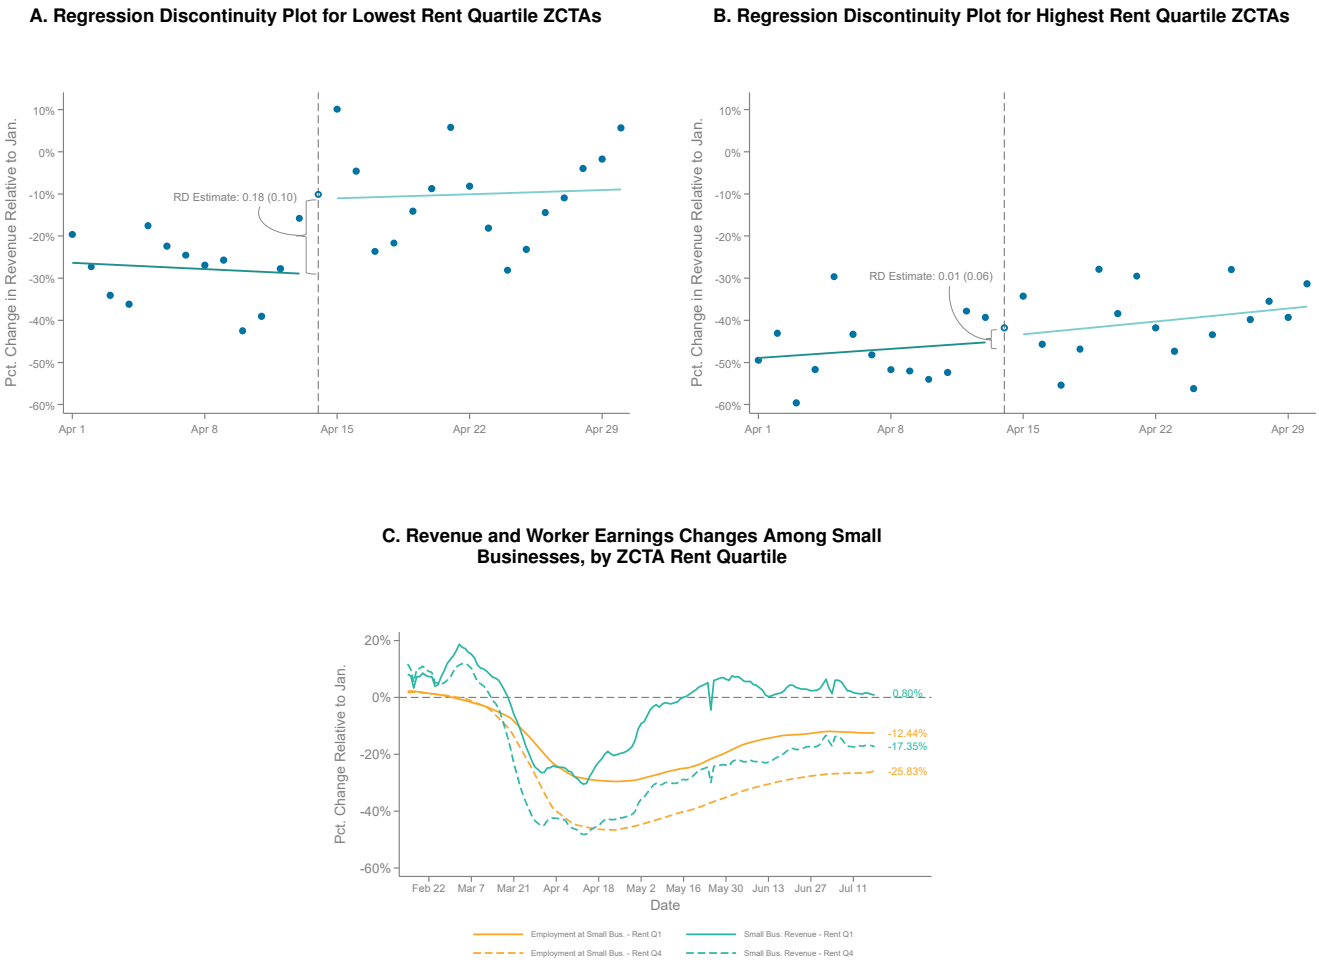
Notes: Panels A and B show percent change in consumer spending in the Affinity Solutions data. Consumer spending is normalized by its level over the period January 4-31, and seasonally adjusted using 2019 data, as described in Section II.A. Panel A shows the series for both New Mexico and Colorado; Colorado 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 Earnin, Intuit and Paychex 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.

FIGURE 13: Impact of Stimulus Payments on Consumer Spending



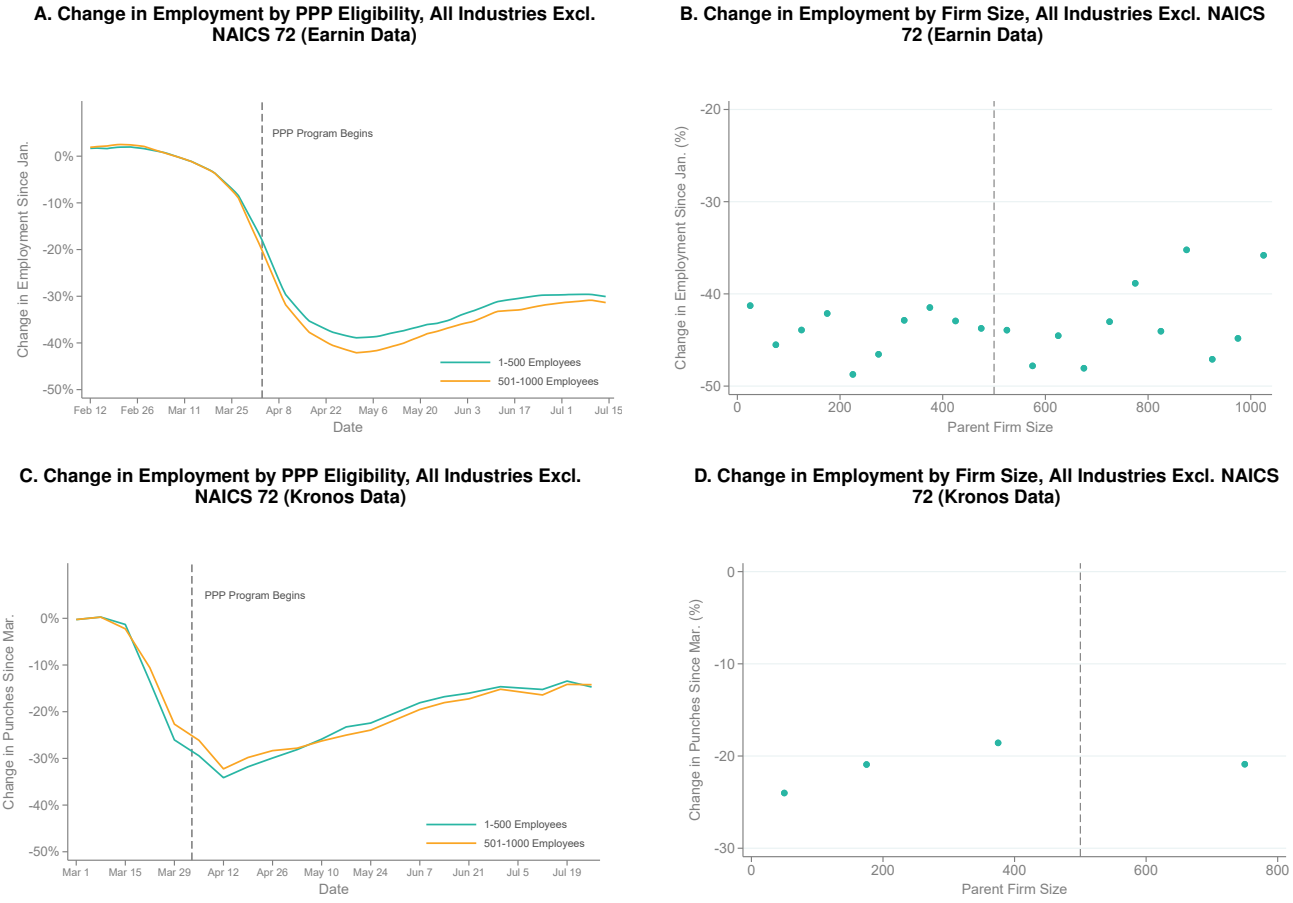
Notes: This figure studies the effect of the stimulus payments made on April 15, 2020 on credit and debit card spending using data from Affinity Solutions. Panel A plots the percent change in seasonally-adjusted consumer spending for cardholders living in ZCTAs in the bottom and top quartiles of the distribution of ZCTA 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. The points are 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. 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 ZCTAs in the lowest median household income quartile. Panel C replicates B for cardholders in highest income quartile ZCTAs. 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.

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



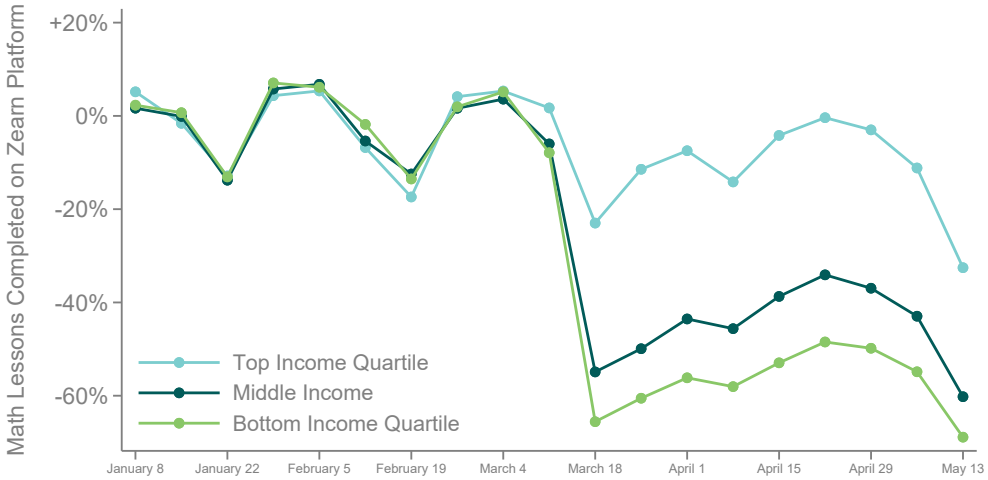
Notes: Panels A and B of this figure study the effect of the stimulus payments made on April 15, 2020 on small business revenues using data from Womply. These panels are constructed in exactly the same way as Panels B and C of Figure 13 except that (1) we use revenue instead of spending as the outcome and (2) we split ZCTAs 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 ZCTA rent-quartile. The employment series restricts to smaller businesses in the Earnin sample, defined by parent employer size being at most 500 employyes. 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.

FIGURE 15: Impact of Paycheck Protection Program on Employment



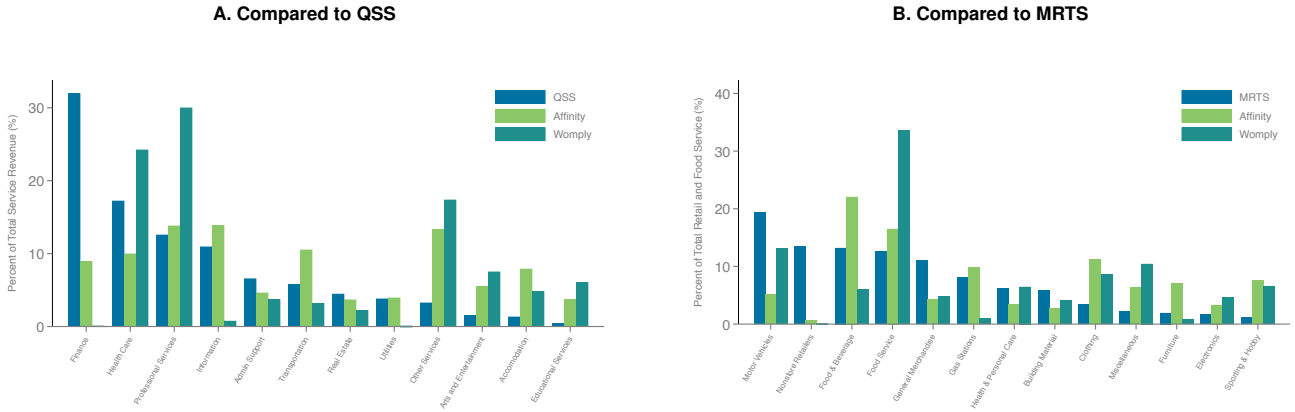
Notes: This figure shows the effects of the Paychex Protection Program on employmenr. Each panel excludes workers in the Accommodation and Food Services sector (NAICS 72). Panels A and B show the impact of the PPP in Earnin data. Panel A compares employment trends among firms with 1-500 employees, which were eligible for PPP loans, vs. firms with 501-1,000 employees, which were ineligible for PPP loans. Both eligible and ineligible groups are reweighted to match the 2-digit NAICS composition in January 2020 of firms with 1-1,000 employees. The percent change in employment respect to employment between January 4 and 31. The grey dashed line corresponds to April 3, 2020, the first day for enrollment in the Paycheck Protection Program (PPP). Panel B presents a binned scatterplot of changes in employment between the period January 4-31 and the period June 1-23 against 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. When taking both the mean change in employment and mean firm size, we reweight in each bin to match the 2-digit NAICS composition in January 2020 of firms with 1-1,000 employees. Panels C and D replicate Panels A and B respectively using Kronos data. As Kronos data is first observed in March we express changes relative to the March level. As industry is not observed in Kronos data, we do not reweight on industry. Data sources: Earnin, Kronos

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



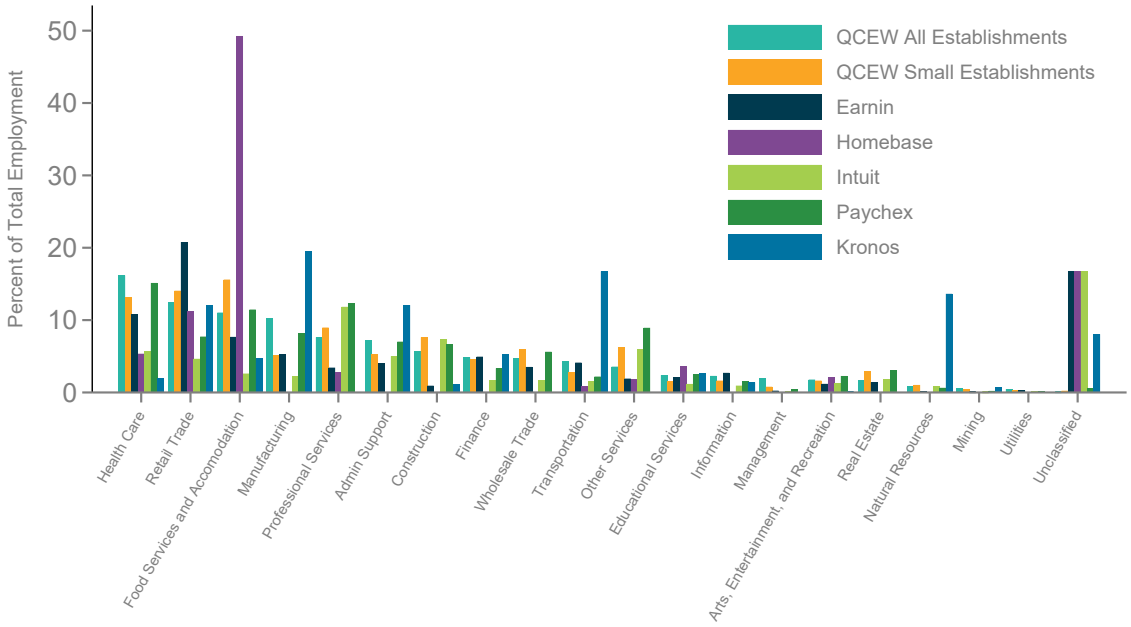
Notes: This figure plots educational progress on the Zearn online math platform for schools located in ZCTAs in the bottom, middle (quartiles 2 and 3), and top quartile of the distribution of median household income. Student progress is defined as the number of accomplishment badges earned in Zearn in 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 ZCTA level weighting by the average number of students using the platform at each school during the base period.

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



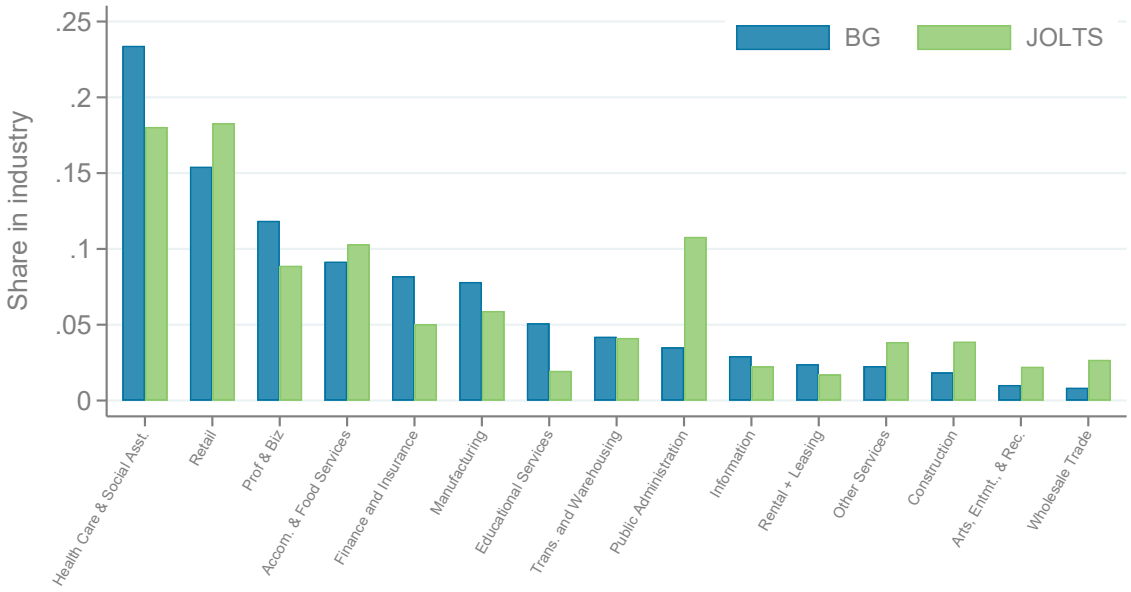
Notes: Panel A shows the NAICS two-digit industry mix for two private business credit card transaction datasets compared with the Quarterly Services Survey (QSS), a survey dataset providing timely estimates of revenue and expenses for selected service industries. Subsetting to the industries in the QSS, 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: Industry Shares of Employment Across Datasets



Notes: This figure shows the NAICS two-digit industry mix for two 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. Each bar represents the share of employees in the given dataset who work in the specified sector. 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. We construct employment shares for the private datasets, Earnin and Homebase, using January 2020 employment. We define employment in Earnin as the total number of worker-days in the month. We define employment in Homebase as the number of unique individuals working a positive number of hours in the month. Data sources: Earnin, HomeBase

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



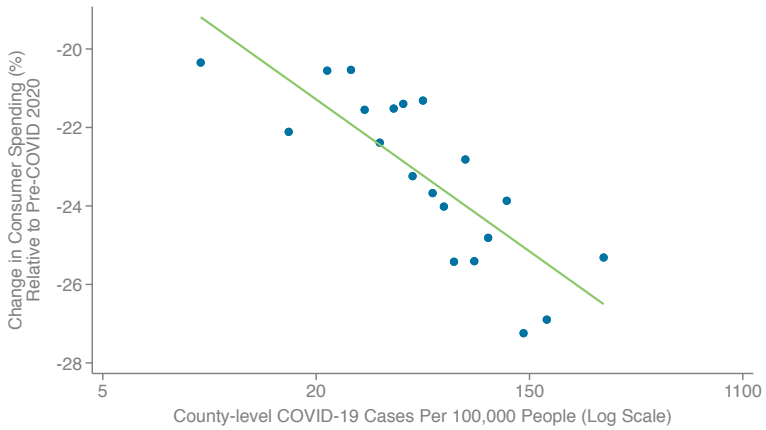
Notes: This Figure displays the NAICS two-digit industry mix of job postings in Burning Glass and job openings in JOLTS, the Job Openings and Labor Turnover Survey data provided by the U.S. Bureau of Labor Statistics, in January 2020. Data source: Burning Glass

APPENDIX FIGURE 4: Spending Changes by Sector and Income Quartile



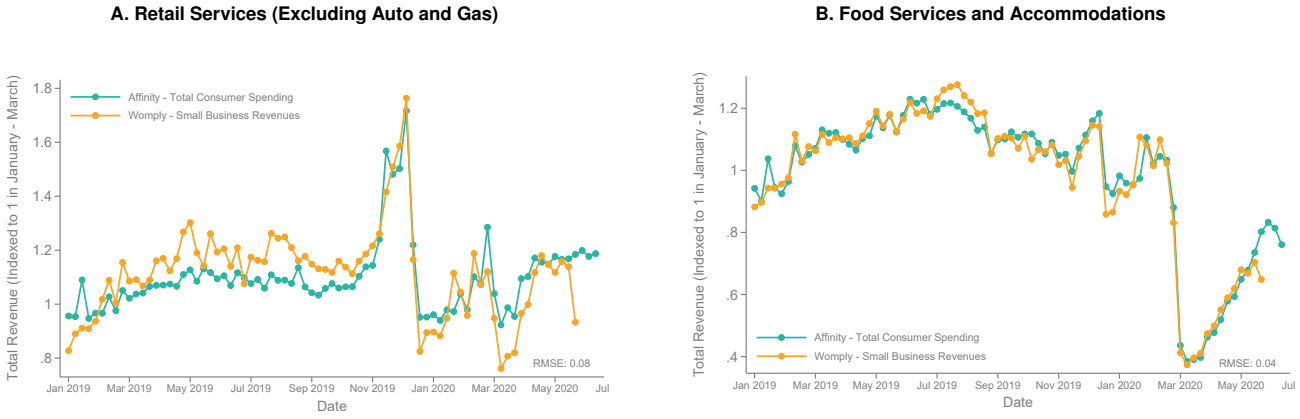
Notes: This figure displays the change in spending by sector for the four quartiles of ZCTA median household income (constructed using 2014-2018 ACS population and income estimates). These sectors were constructed by grouping together similar merchant category codes, not all merchant category codes were used in this plot. The change in spending displayed is (the log difference-in-difference of spending -1)*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 April14th, and is benchmarked to the pre-period of January 8th to January 28th. Data source: Affinity Solutions

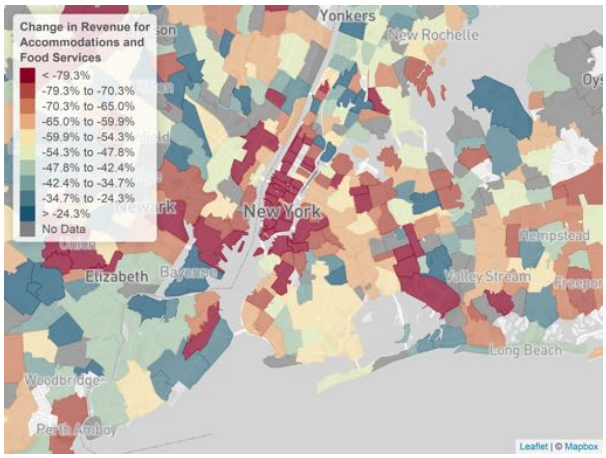
APPENDIX FIGURE 6: Small Business Revenue Changes vs. Local Income Distribution



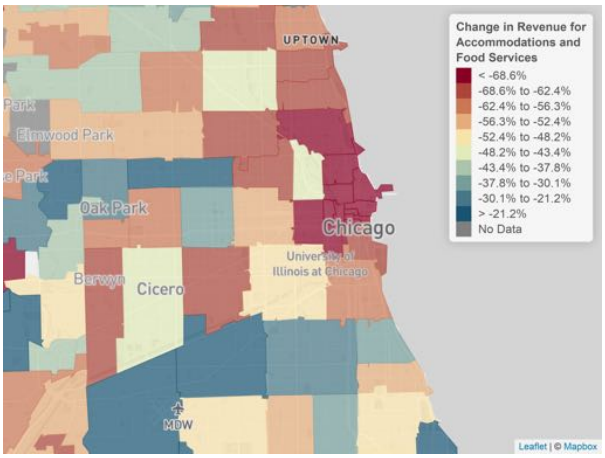
Notes: This figure compares weekly total consumer spending (from Affinity Solutions purchase data) and small business revenue (from Womply) normalized to the average pre-COVID levels of each year. The pre-COVID period is defined as January 8 - March 10 and we normalize within each calendar year to account for year fixed effects. 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 weekly spending on indexed Womply weekly spending. Data sources: Affinity Solutions, Womply

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

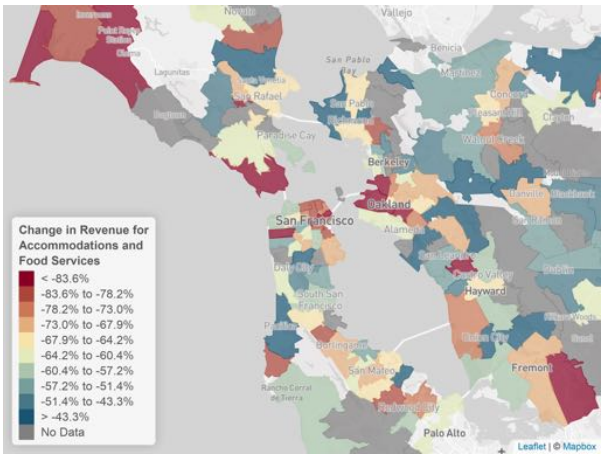
A. New York City



B. Chicago

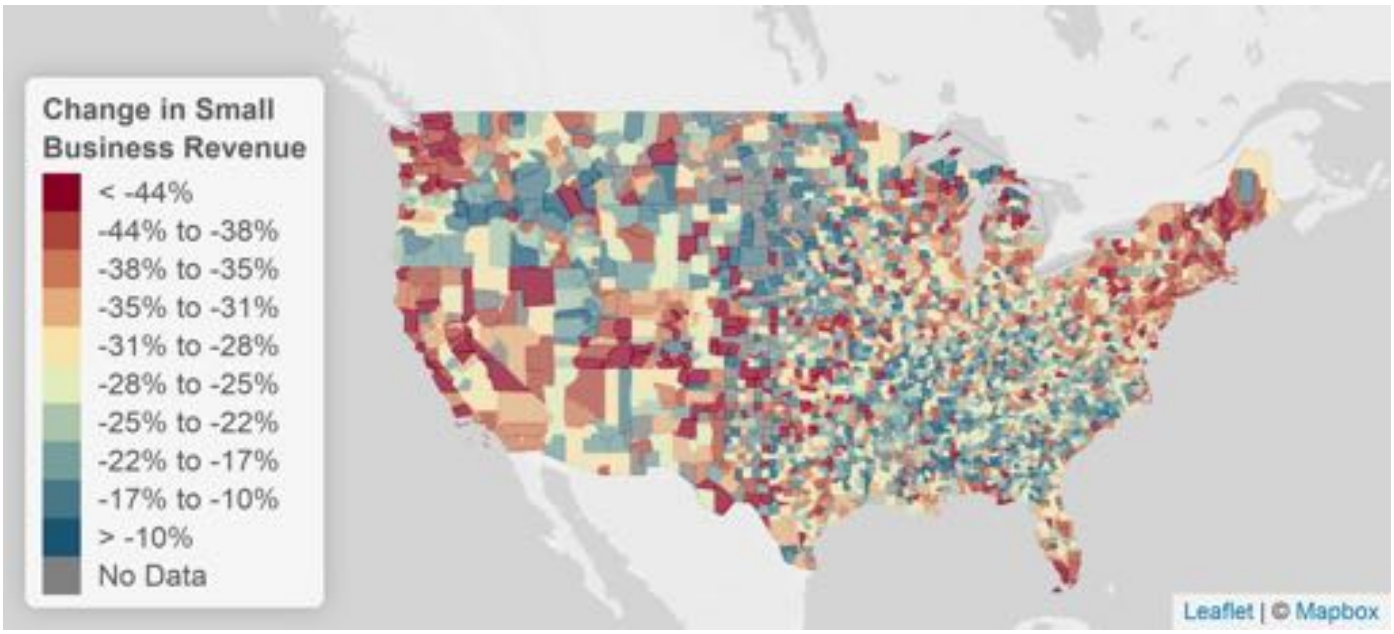


C. San Francisco



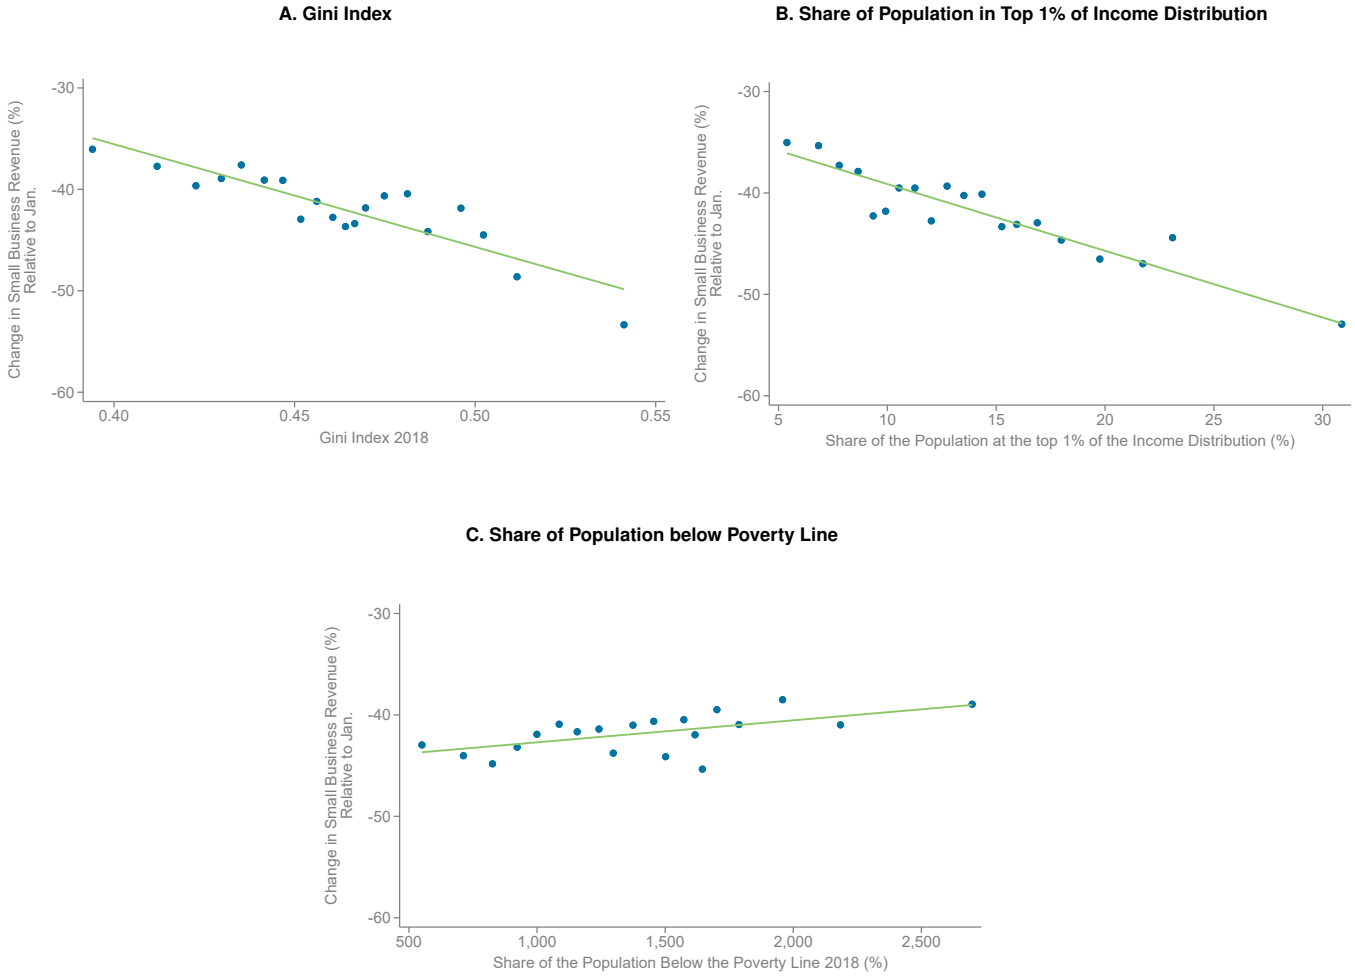
Notes: This Figure displays ZCTA-level maps of the MSAs corresponding to New York City, San Francisco, and Chicago, colored by their respective deciles of change in revenue for small businesses classified as NAICS 72 within each MSA. This figure was generated using the process described in the notes for 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



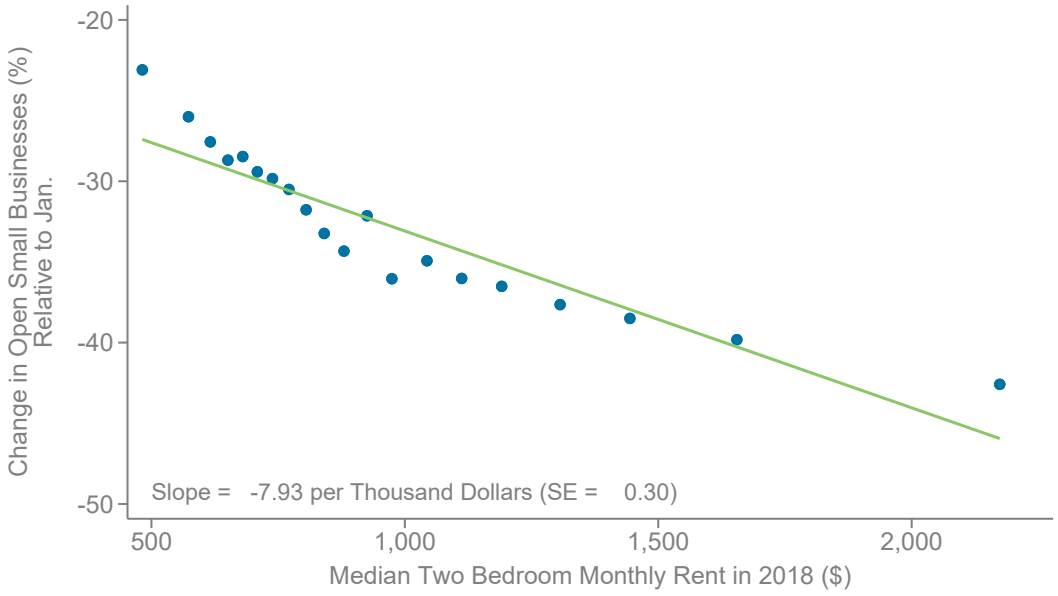
Notes: This figure replicates Figure 4 for the entire United States instead of a single MSA, using counties instead of ZCTAs. See notes to Figure 4 for details. The signal variance to total variance ratio for this map is 0.71. Data source: Womply

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



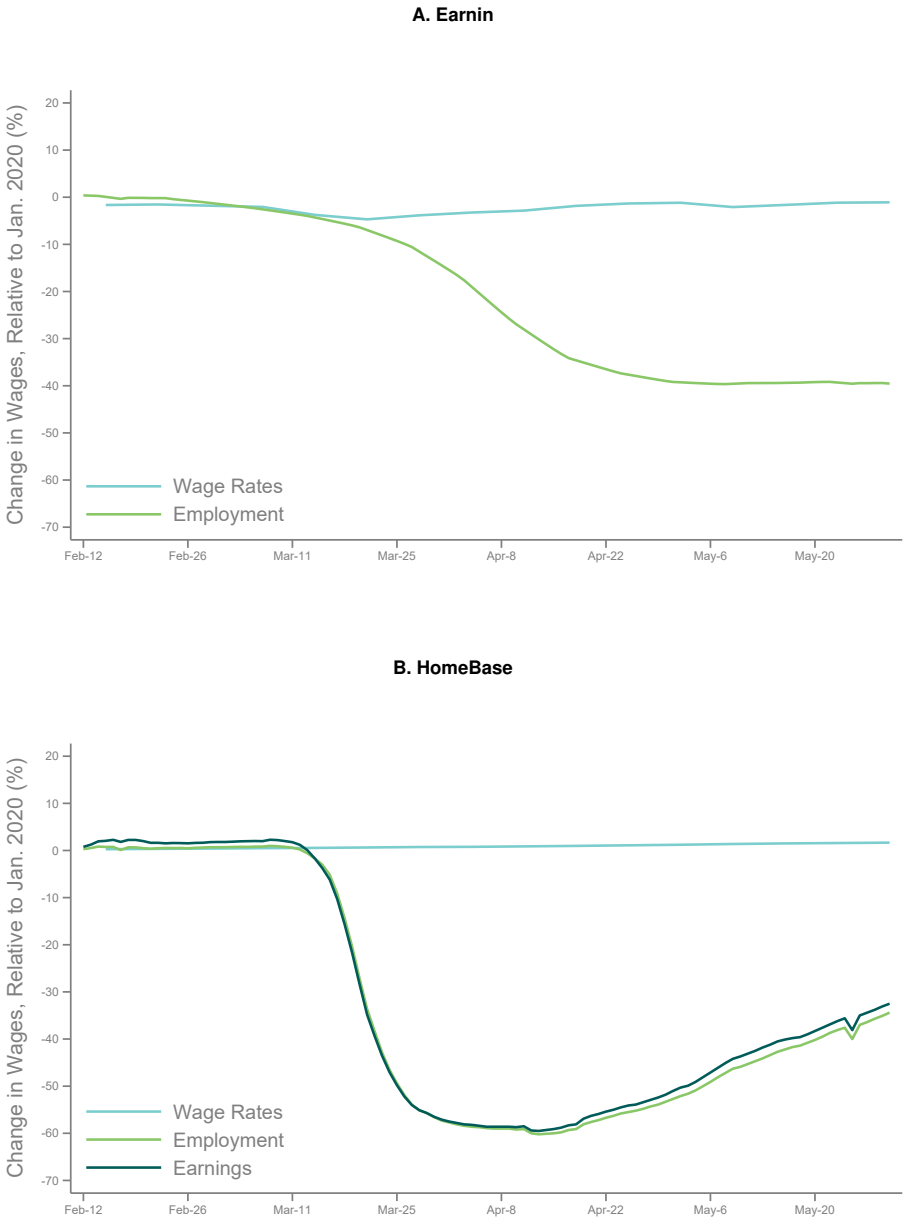
Notes: This Figure replicates Figure 5 but compares the declines with different measures of inequality. Panel A compares the within county Gini index against the declines. Panel B uses the share of the county with incomes at the top 1% of the income distribution. Panel C compares the declines with the 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



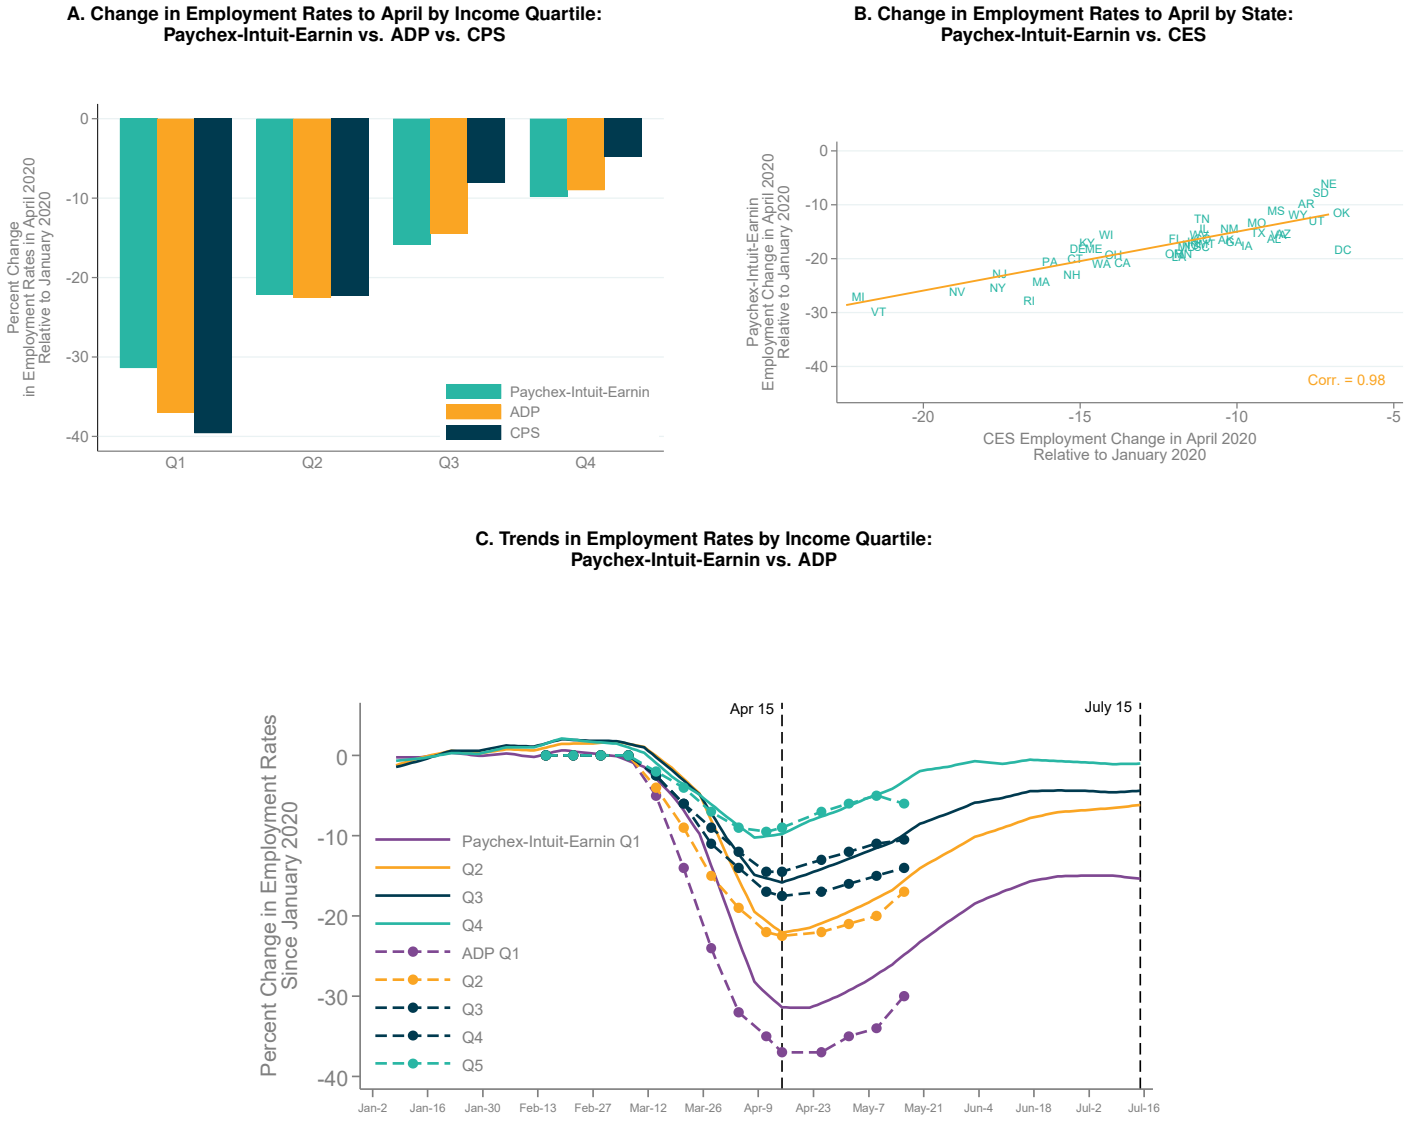
Notes: This figure replicates Panel C of Figure 5 but shows average changes in small businesses that remain open instead of changes in revenue. See notes to Figure 5 for details. Data source: Womply

APPENDIX FIGURE 11: Changes in Wages, Hours Worked and Earnings Over Time



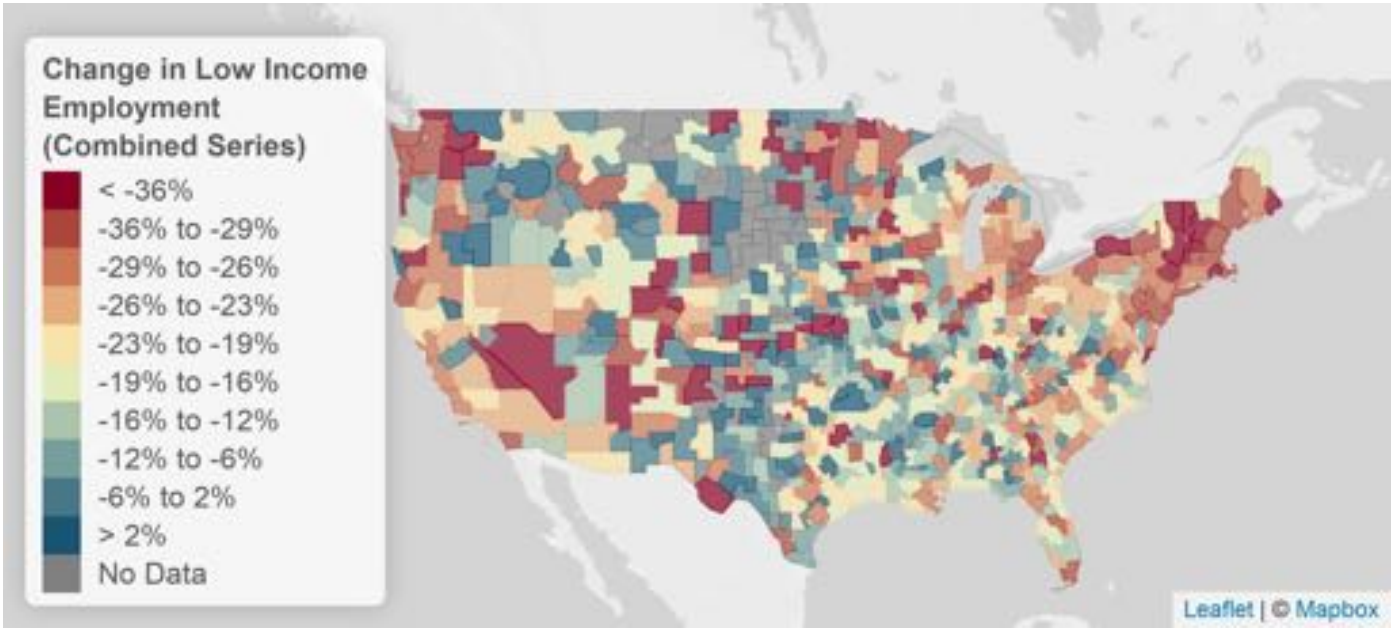
Notes: This figure compares changes in mean wages and employment relative to January 2020 within the Earnin (Panel A) and HomeBase (Panel B) datasets.We construct daily wages for both Earnin and HomeBase by calculating the mean wage on each day. In the HomeBase dataset, we condition on workers being employed by restricting the sample to workers who are observed working in every week of the series. We construct employment in the Earnin and Homebase data and earnings in the Homebase data by summing the total number of hours worked in each day and the total wages earned in each day, respectively. We then take the mean value of each series over the prior seven days and norm to the average value of the seven-day moving average over the period January 4 - January 31, 2020. Data sources: Earnin, HomeBase

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



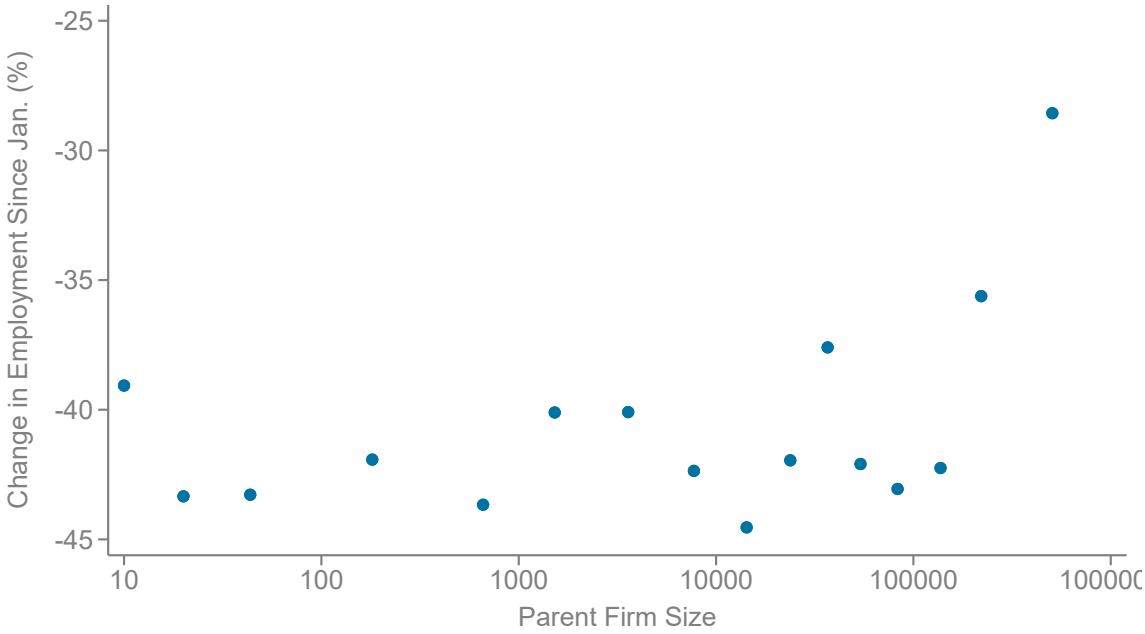
Notes: This figure compares the Paychex-Earnin-Intuit combined employment series to the Current Population Survey (CPS), the Current Employment Statistics (CES), and estimates based on data from ADP by Cajner et al. (2020). Data sources: Paychex, Earnin, Intuit.

APPENDIX FIGURE 13: Changes in Employment Rates by CZ



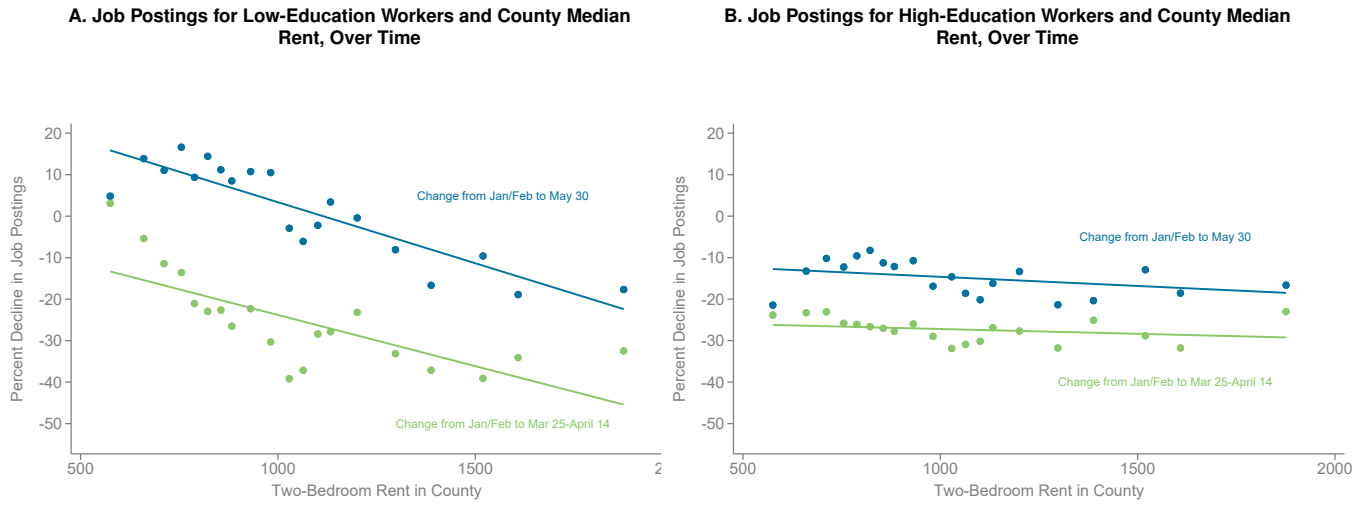
Notes: This figure replicates Figure 7 at the CZ level for the entire United States instead of a single city and its surrounding area, and using combined Paychex/Earnin/Intuit data on aggregate employment across all income quartiles, rather than Earnin data on low-income employment. See notes to Figure 7 for details. Data sources: Paychex, Earnin, Intuit

APPENDIX FIGURE 14: Changes in Total Employment by Firm Size



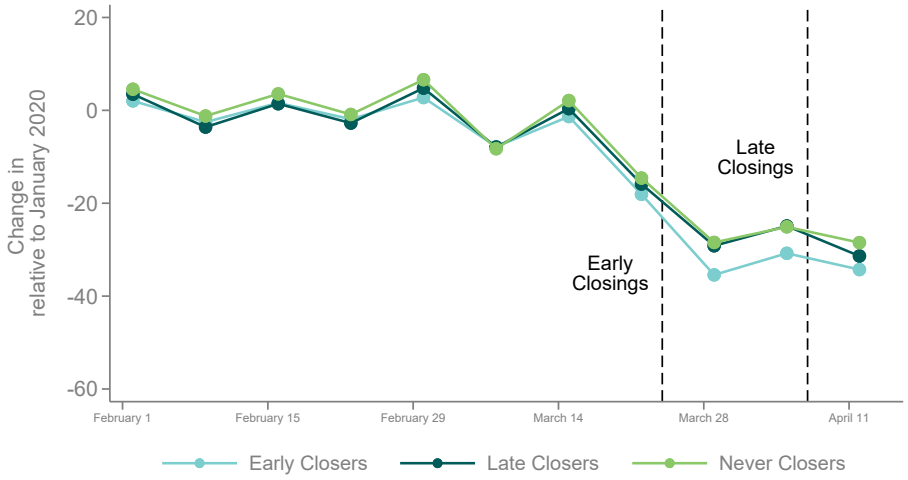
Notes: This figure displays a binned scatterplot of average percent declines in employment in the Earnin data at firms of different sizes. The decline is calculated by taking total employment by firm in a pre-period that spans from January 8, 2020 to January 31, 2020, and comparing to employment in a 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. Data source: Earnin

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



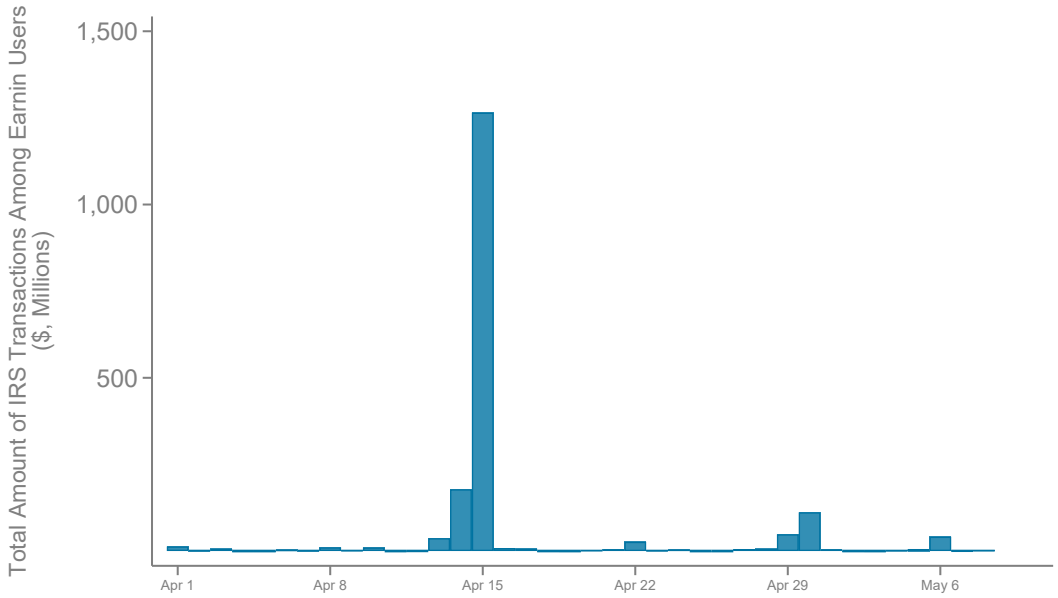
Notes: This figure shows binned scatterplots of the relationship between median rent and changes in job postings between a pre-period of January 8 - March 10 and the periods March 25 - April 14 or the period May 30-June 5. The change in job postings is computed using Burning Glass data. Median two-bedroom 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

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: IRS Transactions Among Earnin Users



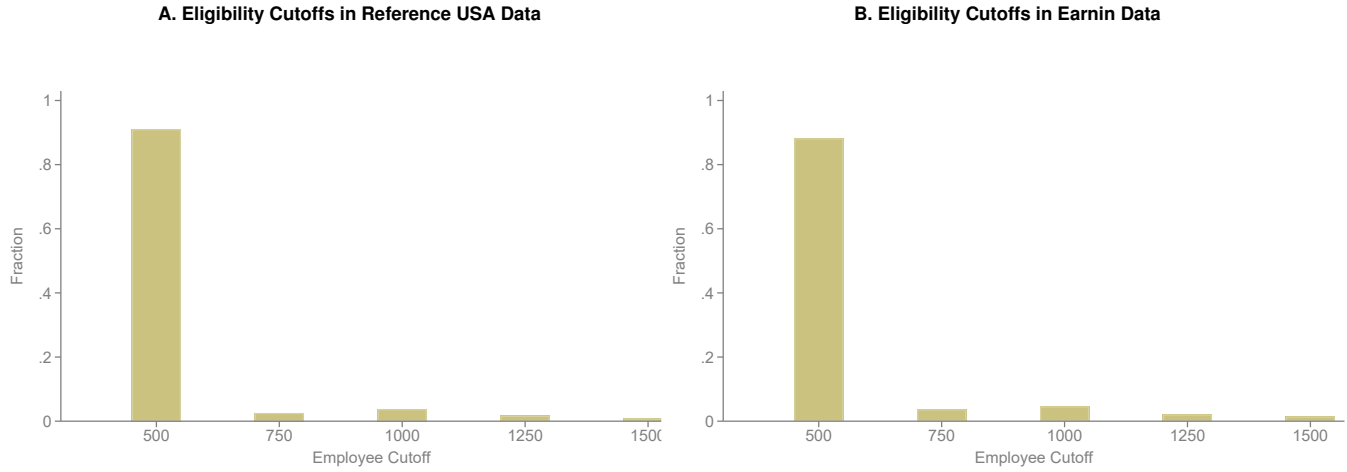
Notes: This figure displays the total dollar amount of IRS transactions for Earnin users. Data source: Earnin

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



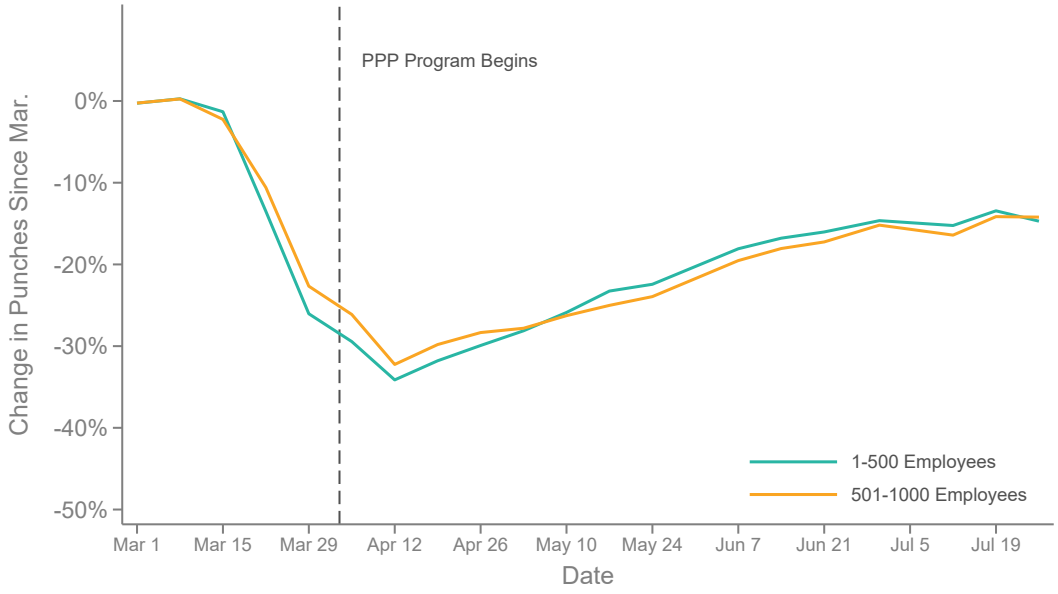
Notes: See notes of Figure 2 Panel B. The pre-stimulus, post-COVID period is defined as March 25th-April 14th. The post-stimulus period is defined as April 29th to May 5th. The total recovery is computed use the post-stimulus period and the average weekly spending in the pre-stimulus period. This figure 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



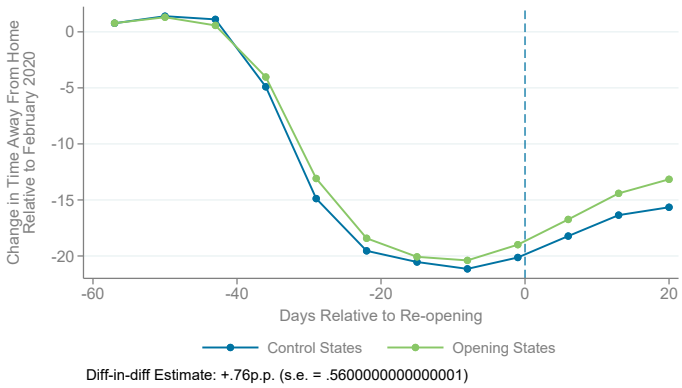
Notes: This figure plots a histogram of the firm size cutoffs for PPP eligibility in the set of firms in Reference USA and the set of firms in the Earnin sample. In the reference USA data, we take the establishment-size-weighted distribution of PPP employee-based eligibility thresholds, which are based on parent company size (except in the case of NAICS 72, which is not included here). In the Earnin sample, we assign a firm size threshold for which the individual's firm would be eligible for PPP loans.

APPENDIX FIGURE 20: Impact of Paycheck Protection Program on Employment: Kronos Data



Notes: This figure replicates Figure 14 using Kronos data on employment. As Kronos data is not observed until March, As industry is not observed in Kronos data, this figure does not reweight on industry. See Figure 14 for details. Data source: Kronos

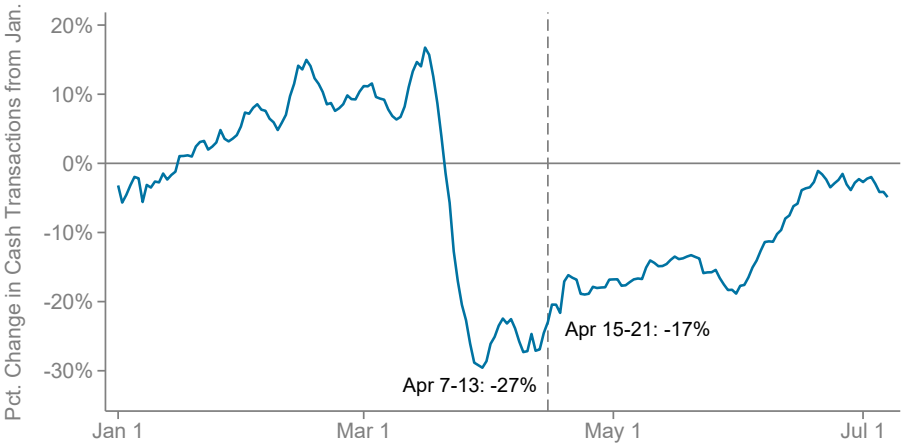
APPENDIX FIGURE 21: Causal Effects of Re-Openings on Mobility: Event Study



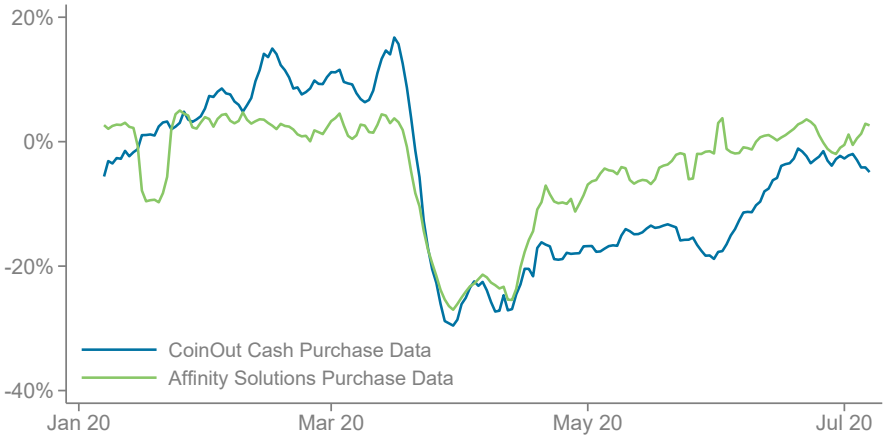
Notes: This figure shows changes in mobility from Google data. Like Figure 11B-D, this is an event study of states that partially reopened non-essential businesses between April 20th and May 4th, compared to a matched control group. We construct the control group separately for states on each opening day and then stack the resulting event studies to align the events. 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.

APPENDIX FIGURE 22: CoinOut Transaction Data Series

A. Change in Cash Transactions Per Day from January 2020



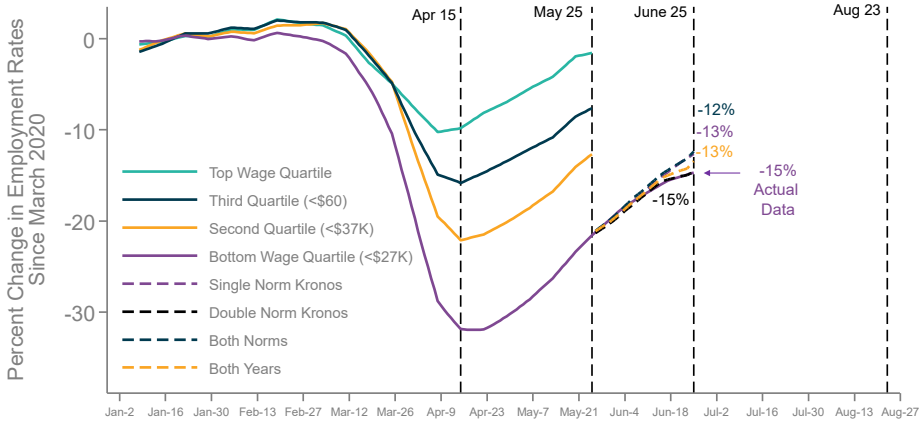
B. Cash Spending in CoinOut Transactions Data vs. Low-Income Spending in Affinity Solutions Purchase Data



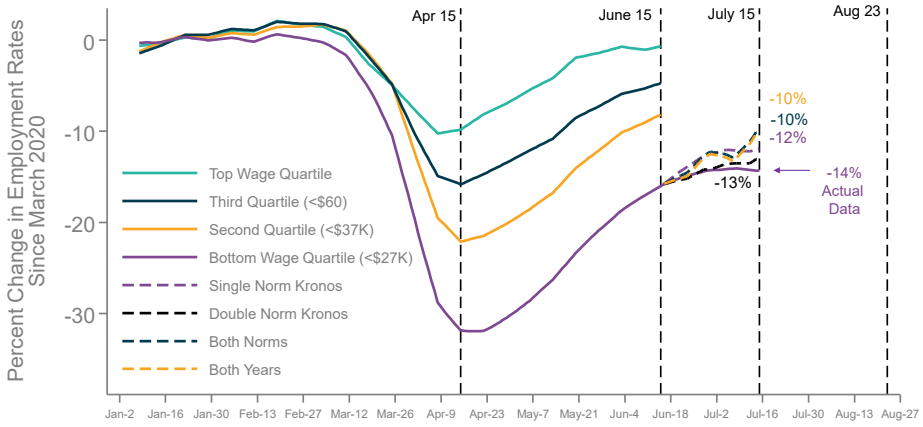
Notes: This figure shows the change over time in daily cash transactions as measured by receipts uploaded to the CoinOut mobile app. Panel A plots the percent change in the seven-day moving average of number of cash transactions uploaded to CoinOut per day. Panel B reports the seven-day moving average of daily cash transactions uploaded to CoinOut compared with the seasonally adjusted seven-day moving average of daily spending for the lowest population-weighted ZCTA median household income quartile from Affinity Solutions. Both series are normalized relative to January 2020 spending (Jan 1 - Jan 31). Data sources: CoinOut, Affinity Solutions.

APPENDIX FIGURE 23: Out-Of-Sample Fit of Advance Series

A. Testing Period: May 25 - June 25

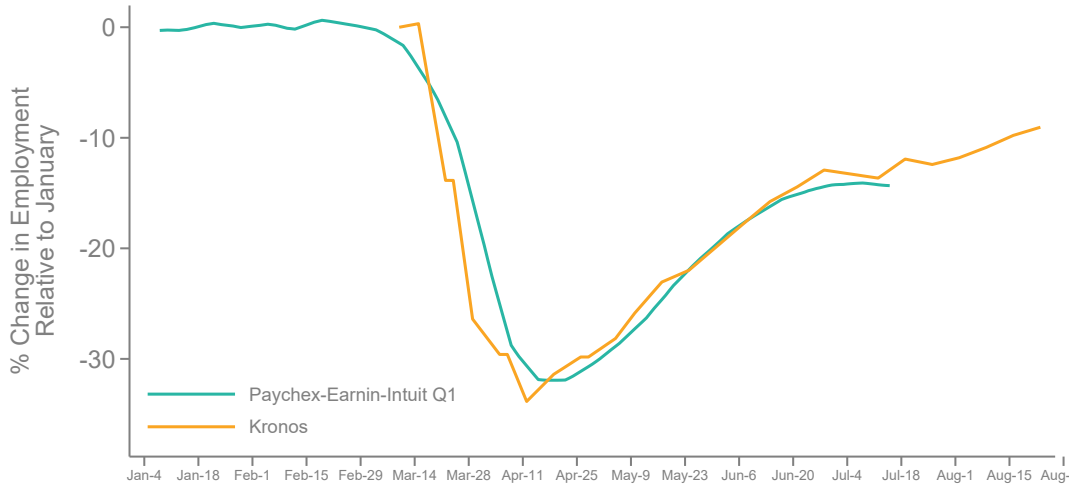


B. Testing Period: June 15 - July 15



Notes: This figure compares the out-of-sample fit of different models predicting employment in the bottom wage quartile in the combined Paychex-Earnin-Intuit series. See Figure 7 for details on how the model is constructed. Single norm Kronos is a model using changes in employment over time in the Kronos data. Double norm Kronos is a model using the de-seasonalized Kronos data. Both norms is a model using both seasonal and de-seasonalized Kronos data. Both years is a model using changes in employment in the 2019 and 2020 Kronos data. 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. Data sources: Paychex, Earnin, Intuit, Kronos.

APPENDIX FIGURE 24: Paychex-Earnin-Intuit Combined Employment for Bottom Quartile vs. Kronos Employment



Notes: This figure shows the change over time in employment as measured by the Kronos punches data compared the vs. the bottom wage quartile in the combined Paychex-Earnin-Intuit series. Data source: Paychex, Earnin, Intuit, Kronos.