Online Appendix

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

Raj Chetty, Harvard University and NBER
John N. Friedman, Brown University and NBER
Michael Stepner, University of Toronto
and the Opportunity Insights Team†

July 2023

†The Opportunity Insights Economic Tracker Team as of July 2023 has consisted of Hamidah Alatas, Camille Baker, Harvey Barnhard, Matt Bell, Gregory Bruich, Tina Chelidze, Lucas Chu, Westley Cineus, Sebi Devlin-Foltz, Michael Droste, Dhruv Gaur, Federico Gonzalez, Rayshauna Gray, Abigail Hiller, Matthew Jacob, Tyler Jacobson, Margaret Kallus, Fiona Kastel, Laura Kincaide, Caitlin Kupsc, Sarah LaBauve, Lucía Lamas, Maddie Marino, Kai Matheson, Jared Miller, Christian Mott, Kate Musen, Danny Onorato, Sarah Oppenheimer, Trina Ott, Lynn Overmann, Max Pienkny, Jeremiah Prince, Sebastian Puerta, Daniel Reuter, Peter Ruhm, Tom Rutter, Emanuel Schertz, Shannon Felton Spence, Krista Stapleford, Kamelia Stavreva, Ceci Steyn, James Stratton, Clare Suter, Elizabeth Thach, Nicolaj Thor, Amanda Wahlers, Kristen Watkins, Alanna Williams, David Williams, Chase Williamson, Shady Yassin, Ruby Zhang, and Austin Zheng.

A Automated Data Processing Pipeline

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

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

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

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

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

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

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

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

B Consumer Spending Series Construction

B.1 Structure of Initial Data

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

B.2 Internal Data Processing

Adjusting for sharp changes in the client base

The raw Affinity data have discontinuous breaks caused by entry or exit of card providers from the sample. We identify these sudden changes systematically by regressing the number of weekly county-level transacting cards on the date, then implementing a Supremum Wald test for a structural break at an unknown break point. This test is implemented separately for the pre-pandemic period before March 11, 2020 and the post-pandemic period after that date. We apply a correction to structural breaks where the p-value of the test is less than 5×10^{-10}.

For counties with a break below this threshold in only one of the two periods, we correct our estimates as follows. We first compute the state-level week-to-week percent change, 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 the county-level percent change with the state-level percent change in weeks $t-1$, t, and $t+1$, as we cannot ascertain the precise date when the structural break occurred (e.g., it may have occurred on the 2nd day of week $t-1$ or the 6th day of week t).

For example, suppose a county has n active cards up until week t, when the number of cards in the county increases to $3n$. In week $t-2$, the county would have a level of n cards, its reported value. In week $t-1$, if counties in the rest of the state had a 5% increase in the number of cards, we would impute the county with a break to have a level of $1.05n$ cards. In week t, if counties in the rest of the state had a 10% increase in the number of cards, we would impute t to have a level of $(1.10) \times (1.05n) = 1.155n$. 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 $t+1$ to have a level of $(0.92) \times (1.155n) = 1.0626n$. Finally, if the county with a break had an increase of 3% in week $t+2$, this would be multiplied by the imputed level for week $t+1$ to produce a level of $(1.03) \times (1.0626n) = 1.0945n$ in week $t+2$. This final step would be repeated to update levels for all periods after week $t+2$.

When there is a change in coverage, we adjust the series to be in line with the lower level of coverage. In cases where the lower level of coverage was after the break, the above example would implemented in reverse to impute levels prior to the break.

We omit counties with multiple structural breaks (identified either by the Supremum Wald test in both periods or manually) from our series. We also add manually identified structural breaks which are not identified by the test, for example because the break phases in over more than one week. We do not remove any counties where the structural break occurred between March 10 and 31, 2020 because the consumer spending response to the COVID-19 was so strong that in many places it could be classified as a structural break. Additionally, since holiday spending spikes are also sometimes classified as a structural break, we manually verify any breaks detected in the week of Thanksgiving or the weeks immediately before, during or after Christmas.

Adjusting for gradual changes in the client base

In our processing, we aim to isolate and remove variation in spending driven by the number of debit/credit cards in our sample changing due to steady growth or shrinkage in a payment provider's customer base. However, we also aim to incorporate extensive margin variation in spending driven by the number of debit/credit cards in our sample changing due to changes in the utilization of cards, which can reflect underlying economic conditions. To account for both, we first estimate a two-state switching model of the relationship between Personal Consumption Expenditure as measured by the U.S. Bureau of Economic Analysis and the level of credit and debit spending and number of cards in usage we received from Affinity Solutions. Our estimates imply a state change between February and August 2020: the number of cards in use is predictive of Personal Consumption Expenditures during this period, but statistically insignificant outside this period. This is consistent with a pattern where people exhibited extensive margin changes in card spending during the initial waves of COVID-19, but outside of that period changes in the number of cards being used reflect changes in the number of people in-sample as opposed to changes in the propensity to consume of a given set of people.

We therefore construct an estimate of spending that reflects total spending between February and August 2020, and spending per card outside this period. For each location (l), industry (i) and time (t) we compute:

$$\text{adjusted\_spending}_{l,i,t} = \begin{cases} \frac{spending_{l,i,t}}{cards_{l,\cdot,t}} \times \overline{cards_{l,\cdot,t' \in \text{Jan 2020}}} & \text{if } t \leq \text{January 2020} \\ \\ \frac{spending_{l,i,t}}{cards_{l,\cdot,t}} \times \overline{cards_{l,\cdot,t' \in \text{Jan 2020}}} \times \frac{\overline{cards_{l,\cdot,t' \geq t | t' \in \text{Feb 2020}}}}{\overline{cards_{l,\cdot,t' \in \text{Feb 2020}}}} & \text{if } t \in \text{February 2020} \\ \\ \frac{spending_{l,i,t}}{cards_{l,\cdot,t}} \times \overline{cards_{l,\cdot,t' \in \text{Jan 2020}}} \times \frac{cards_{l,\cdot,t}}{\overline{cards_{l,\cdot,t' \in \text{Feb 2020}}}} & \text{if } t \in [\text{March 2020, July 2020}] \\ \\ \frac{spending_{l,i,t}}{cards_{l,\cdot,t}} \times \overline{cards_{l,\cdot,t' \in \text{Jan 2020}}} \times \frac{\overline{cards_{l,\cdot,t' \leq t | t' \in \text{Aug 2020}}}}{\overline{cards_{l,\cdot,t' \in \text{Feb 2020}}}} & \text{if } t \in \text{August 2020} \\ \\ \frac{spending_{l,i,t}}{cards_{l,\cdot,t}} \times \overline{cards_{l,\cdot,t' \in \text{Jan 2020}}} \times \frac{\overline{cards_{l,t' \in \text{Aug 2020}}}}{\overline{cards_{l,\cdot,t' \in \text{Feb 2020}}}} & \text{if } t \geq \text{September 2020} \end{cases}$$

We use this adjusted spending series throughout our analysis, except in Section IV.A, where we always use total spending (thereby including both intensive and extensive margin changes) in order to ensure consistency in comparisons of our estimates across different stimulus rounds.

Addressing spurious changes

There is an large spike in consumer spending between January 15 and 17, 2019 that is not found in other data series. This spike in national consumer spending is likely driven by the early release of February 2019 SNAP benefits due to an impending government shutdown.[38] In order to avoid contaminating our seasonal adjustment with this one-time shock, we replace each impacted day with the average spending on $t-7$, $t+7$, and $t+14$, where t is the impacted day.

[38]Described in "Billions in food stamp payments to come early because of shutdown", Politico article on January 11, 2019: https://www.politico.com/story/2019/01/11/shutdown-food-stamp-scramble-benifits-1081210

Control for seasonal fluctuations in spending

We seasonally adjust the data by calculating, for each week and day, the year-on-year change relative to the 2019 value. We norm February 29, 2020 (a Saturday) relative to the average of February 23 and March 2, 2019 (both Saturdays). Labor Day in 2019 fell one week earlier than in 2020, so we adjust the week of Labor Day, as well as the two weeks before, based on the same week in 2019 relative to Labor Day rather than the week number in the year. We then calculate the change relative to the January index period: 2019 data is indexed relative to January 7 to February 3, 2019, data in 2020 onward is indexed relative to January 6 to February 2, 2020. We then seasonally adjust by dividing by the indexed 2019 value, which represents the difference between the change since January 2020 compared to the change since January 2019.

B.3 Masking and Publication

Cells with fewer than 5 card transactions are masked. Additionally, because the Supremum Wald test cannot identify breaks in the most recent few weeks of data, we mask particularly large changes. If a week-on-week change is > 4x the median change of that series, and $< 1/3$ of all series have a large change defined in this way, the value of the change is replaced with the preceding week multiplied by the national average change. We place this 1/3 restriction to avoid overmasking holiday spending, when almost all series experience large changes.

Definition of Categories of Goods

In parts of our analysis, we distinguish between four categories of goods and services: durable goods; non-durable goods; remote services; and in-person services.

- We define durable goods as the following MCC groups: Building materials, garden equipment, and supplies; electronics and appliances; furniture and home furnishings; motor vehicles and parts; sporting goods, hobbies, musical instruments, and book stores; and telecommunications.

- We define nondurable goods as the following MCC groups: General merchandise; wholesale trade; clothing and clothing accessories; health and personal care stores; food and beverage stores; misc store retailers; and gas stations.

- We define remote services as the following MCC groups: Utilities, construction, and manufacturing; professional, scientific, and technical services; public administration; administrative and support and waste management and remediation services; information; education; finance; and nonstore retailers.

- We define in-person services as the following MCC groups: Rental and leasing; repair and maintenance; and personal and laundry services.

Substitution from Card to Cash Spending

A potential concern with our card-based estimates of spending changes is bias from substitution out of cash purchases, which account for 6.3% of consumer spending in the United States (Greene and Stavins 2020). For instance, if individuals sought to use more contactless methods to pay or began placing more orders online, trends in card spending might exhibit excess volatility relative to overall spending. To assess the importance of such substitution, we examine cash purchases using receipts data from CoinOut, a company that allows individuals to receive rewards by uploading photos of their receipts to a mobile app. We focus on grocery spending in the card data because cash spending in CoinOut is concentrated in certain sectors such as groceries; unfortunately, we are unable to disaggregate the CoinOut data by sector or align sectoral definitions more precisely across the datasets.

Appendix Figure VI.B plots week on week changes in aggregate cash purchases in the CoinOut data vs. aggregate card spending at grocery stores over time. The time trends are very similar between the two series (with a correlation of 0.67 at the weekly level), showing a sharp spike in spending in late March 2020 (as households stocked up on groceries), followed by a more sustained increase in spending from the latter half of April 2020. These results suggest that households shifted spending similarly across both modes of payment.

B.4 Benchmarking

Comparison to QSS and MARTS. Total debit and credit card spending in the U.S. was $7.08 trillion in 2018 (Board of Governors of the Federal Reserve System 2019), approximately 50% of total personal consumption expenditures recorded in national accounts. Appendix Figure II compares the spending distributions across sectors in the Affinity data to spending captured in the nationally representative Quarterly Services Survey (QSS) and Advance Monthly Retail

Trade Survey (MARTS), which together cover 92% of the expenditure-weighted categories in the Affinity data. The Affinity series has broad coverage across industries, but over-represents categories in which credit and debit cards are used for purchases. In particular, accommodation and food services and clothing constitute a greater share of the card spending data than financial services and motor vehicles. We therefore view the Affinity series as providing statistics that are representative of total card spending, but not total consumer spending. We assess whether the Affinity series accurately captures changes in total card spending in the COVID recession in Section III.A.

C Small Business Revenue and Openings Series Construction

C.1 Structure of Initial Data

We receive total small business debit and credit revenue data from Womply, where the total revenue data is aggregated from settled credit card transactions that Womply receives from payment processing partners. Additionally, we receive the number of businesses open data from Womply where the number of businesses open is defined as the number of businesses making at least one transaction within a three day window. The primary raw data we receive and process is a 52-week "no-entry" panel of firms at the county by sector by week level, which is a repeated panel following the set of firms operating in each week t over the subsequent 52 weeks (with attrition from the sample but no entry during the panel). To construct ZIP-level estimates, we additionally use a cross-sectional sample of firms at the ZIP code by ZIP income quartile by sector by day level. We measure small business revenue as the sum of all credits (generally purchases) minus debits (generally returns). All transactions and derived data we receive are tied to the ZIP code or county containing the business. The sample is limited to small businesses as defined by the Small Business Administration.

C.2 Internal Data Processing

Adjusting for the evolving client base

In each calendar year, we follow the sample of businesses operating during the first week of the year (i.e. we start following a new panel each calendar year). No new businesses enter our panel during the calendar year. Businesses may exit because they stop operating or because the underlying payment processors ceased providing data.

We detect cases where a payment processor disappears by detecting sharp drops in businesses operating, at the national and the state level. We then adjust the series to identify and correct breaks in the data introduced by these merchant exits from the sample. To do so, we identify breaks in series from July 2020 to present as downward discontinuities of at least 2.5 percentage points from week $t-1$ that persist for at least two weeks—i.e. remain at least 2.5 percentage points below week $t-1$ in weeks t, $t+1$ and $t+2$. For weeks prior to July 2020 we manually identify breaks in weeks 21, 22, 27, 31 nationally and at week 25 for a subset of states (where week numbers refer to weeks of 2020).

For breaks identified prior to July 2020, business revenue was in a period of strong recovery so we assume "momentum" during the weeks with missing data. We calculate the average rate of change for the 4 preceding weeks to the discontinuity and use that value to impute the rate of change for the identified week of the discontinuity. For breaks identified from July 2020 to present we adjust a given merchant or sales series as follows. For a merchant series, for the week of a discontinuity we impute the value of the number of merchants as the value from week $t-1$ and adjust the following weeks accordingly resetting the number of merchant to the level prior to the identified exit from the sample. For a sales series, if week $t-1$ has a discontinuity we compute the change in sales between week $t-1$ and week $t-2$ to be zero and adjust the following weeks accordingly. If additionally week t has a discontinuity we correct week t following the same procedure and adjust the following weeks accordingly.

Produce the ZIP code series

As described in Appendix C.1, we do not receive any panel data disaggregated to the ZIP code level: we only observe cross-sectional ZIP level data. We therefore perform an additive adjustment on the cross-sectional ZIP level series so that the weighted sum of the processed ZIP series aligns with the county-level "same store" panel data. The resulting data obtains the levels from the county "no-entry" panel, and the within-county across-ZIP variation from the ZIP level cross-section.

Reducing the influence of outliers

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 also manually exclude some state x industry breakdowns that present extreme variation from our state and national level calculations, as well as a small number of counties that demonstrate extreme variation.

Controlling for seasonality

We seasonally adjust reported revenue by calculating the year on year change relative to the corresponding value in 2019. We calculate the change relative to the January index period: 2019 data is indexed relative to January 2019, data in 2020 onward is indexed relative to January 2020. We then seasonally adjust by dividing by the indexed 2019 value.

C.3 Masking and Publication

To preserve the privacy of firms in the data and to avoid displaying noisy estimates for small cells, we mask cells with less than $250,000 in total revenue during the base period of January 4 to 31, 2020. The ZIP-level data we receive from Womply adds merchants and an imputed revenue quantity such that every cell with 1 or 2 merchants has no fewer than 3 merchants. This imputation has the result of dampening the effect of any declines that would otherwise place the number of merchants in a cell at 1 or 2, lowering the effect of any increase from 1 or 2 merchants to 3 merchants, and enhancing the effect of any increase from 0 merchants to 1 or 2 merchants. We address the effects of this imputation by dropping any ZIP code that has imputed values for more than 25% of the weeks in our period of study.

C.4 Benchmarking

Comparison to QSS and MARTS. Appendix Figure II 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.

D Job Postings Series Construction

D.1 Structure of Initial Data

We receive a dataset containing the weekly number of unique new job postings, defined as those that have not already been posted within a 60 day window. This data is disaggregated by geography (county, state, national), by 2-digit NAICS code, and by ONET code. Jobs postings are sourced by Lightcast (formerly known as Burning Glass Technologies) from over 40,000 jobs boards worldwide.

D.2 Internal Data Processing

Reducing the influence of outliers

In order to avoid extreme outliers, we calculate a cutoff of one standard deviation above the 97th percentile of the state-level data for each variable and mask values that exceed this threshold.

D.3 Masking and Publication

We perform some imputations of the county by subgroup-level data to ensure the privacy of the underlying job postings dataset, as required by our data use agreement. In the 200 most-populated counties, total job posts by subgroup are reported directly without imputation. For total job posts by subgroup in smaller counties, we impute the number of job postings: we report the total county level postings by subgroup multiplied by the share of total state level postings of the corresponding subgroup. All state-level data and national data are reported without imputations.

D.4 Benchmarking

Comparison to JOLTS. Lightcast 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) show that the Lightcast data are reasonably well-aligned with government survey-based statistics on job openings and characterize the sample in detail. In Appendix Figure III, we compare the distribution of industries in the Lightcast data to nationally representative statistics from the Bureau of Labor Statistics' Job Openings and Labor Market Turnover Survey (JOLTS) in January 2020. In general, Lightcast is well aligned across industries with JOLTS, with the one exception that it under-covers government jobs. We therefore view Lightcast as a sample representative of private sector jobs in the U.S.

E Employment Series Construction

E.1 Structure of Initial Data

The employment series is constructed with data from three data providers: Paychex, Intuit, and Earnin.

Paychex

We obtain aggregated weekly data on total employment for each county by industry (two-digit NAICS), hourly wage quartile (defined below), firm size bin and pay frequency. Salaried employees' wages are translated to hourly wages by dividing weekly pay by 40 hours. To measure private sector employment, we exclude workers employed in public administration and those with an unclassified industry (which each represent 0.8% of workers as of January 2020). We restrict the sample to workers with weekly, bi-weekly, semi-monthly or monthly pay frequencies; these workers represent over 99.8% of employees in the Paychex data.

To classify workers by wage quartile while adjusting for wage growth, we construct moving wage quartile thresholds based on 100%, 150% and 250% of the federal poverty line (FPL). The FPL is defined as an annual income, which we convert into a full-time-equivalent hourly wage by diving by 2000 hours (50 weeks of work at 40 hours per week). In our benchmark period of January 2020, the thresholds are $13.10, $19.65 and $32.75. These thresholds group workers approximately into quartiles: in January 2020, the four bins in ascending order by wage contain 23.4%, 27.4%, 25.7%, and 23.5% of CPS respondents. Since the FPL is set annually at the beginning of each year, we estimate monthly thresholds within each year by multiplying the FPL by the growth in the monthly Consumer Price Index (CPI) since January. When the official FPL estimate is published at the beginning of the subsequent year, its growth does not exactly match the growth in CPI. To maintain consistency with the levels of the FPL, we revise the thresholds from the prior year when the FPL is released. We compute the difference between our CPI-based thresholds and the official FPL thresholds during the relevant year, divide this residual by 12, and add the rescaled residual to the projected FPL in each month.

We provide these moving wage quartile thresholds to Paychex, which are then used internally by Paychex to assign workers to an hourly wage quartile in the aggregated weekly dataset we receive. In January of each year, we provide revised quartile thresholds to Paychex and revise our estimates for the prior year—the revised estimates hold aggregate employment constant but reallocate some workers across quartiles.

Using moving thresholds to assign workers to wage quartiles generates occasional discontinuities in each quartile due to bunching in the wage distribution at integers, as discussed in Section II.B.4 and detailed in Appendix E.2. To correct these discontinuities, we receive weekly counts of the number of employees earning an integer hourly wage for each integer value between $13 and $25. These counts are received at the same level of disaggregation as the weekly data on employment: county by industry (two-digit NAICS), firm size bin and pay frequency.

Intuit

We obtain anonymized, aggregated data on month-on-month and year-on-year changes in total employment (the number of workers paid in the prior month) 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

We obtain anonymized data from Earnin at the paycheck level with information on industry (2-digit NAICS), firm size, ZIP code, unemployment status, wages, and earnings. The median worker in the Earnin sample in January 2020 has an hourly wage rate of $14.98, which falls at the 30th percentile of the national distribution of private sector non-farm workers in January 2020 CPS data. The interquartile range of wages in the Earnin sample is $9.36-$22.10 (corresponding to the 6th and 59th percentiles of the national distribution).

Industry and firm size are not directly measured in Earnin's administrative databases, so these variables were attached to the anonymized paycheck-level data by linking a list of employers with more than ten Earnin users to external databases. 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 match 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.

E.2 Internal Data Processing

Data is processed in the following ways for data received from the three data providers: Paychex, Intuit, and Earnin.

Paychex

Addressing anomalous spikes. We manually identify large anomalous spikes in the data and smooth these by interpolating values from adjacent weeks at the county x 2-digit NAICS code x hourly wage quartile x 2019 firm size bin x pay frequency level.

Adjusting for discontinuities when quartile thresholds cross integer wages. We adjust for occasional discontinuities in employment in each wage quartile due to bunching in the wage distribution. Our adjustment is mathematically equivalent to adding a uniform random variable distributed between $[-0.5, 0.5]$ to whole number wages, transforming the point mass of employees at the integer *wage* into a uniform distribution between $[wage - 0.5, wage + 0.5]$. Rather than adding the random variable *ex ante* using employee microdata before aggregating, which we did not access, we applied an equivalent procedure that is feasible *ex post* after receiving the aggregated data described above containing employee counts. As an example, suppose the threshold separating the first wage quartile (Q1) and the second wage quartile (Q2) is \$13.75 in a given month. Then adding the uniform random variable to the wages of employees earning \$14 will move 25% of the point mass below the threshold. Since we observe the counts on each side of the \$13.75 threshold and the count at the \$14 point mass, we add $0.25 \times$ (number of employees at \$14) to the Q1 count and subtract the same number from the Q2 count. We perform the analogous process for each threshold in each month, calculating the share of the point mass at the integer value within \$0.50 of the threshold that should be moved across the threshold, and adding and subtracting that value from the quartiles below and above the threshold accordingly.

Mapping paychecks to employment periods. To construct a series of employment as of each date, we construct a series of pay periods ending as of each date. We take a separate approach for paychecks following regular weekly cycles (i.e. weekly and bi-weekly paychecks) and for paychecks following a cycle based on fixed calendar dates (i.e. semi-monthly and monthly paychecks). For weekly and bi-weekly pay frequencies, we use data provided by Paychex on the distribution of the number of days between a worker's pay date and the last date in the worker's pay period (i.e., date at which payroll is processed – last date in pay period), for weekly and bi-weekly pay frequencies, to distribute paychecks to the last date of the corresponding pay period. We treat the distribution of (date at which payroll is processed – last date in pay period) as constant across geographies and NAICS codes. For monthly and semi-monthly pay frequencies, where cycles regularly occur on fixed calendar dates (e.g. the 15th and 30th of each month for semi-monthly paycycles), we assume that the last date within each pay period is the closest preceding calendar date that is the 15th or the 30th day of the month (semi-monthly paycycles) or the 30th day of the month (monthly paycycles). We then record a worker as being employed for the full duration of the pay cycle up until the last date in their pay period, under the assumption that workers are employed for each day during their pay period.

Adjusting for an evolving client base. We take steps to adjust for the entry and exit of Paychex clients from the sample. In each county x industry (two-digit NAICS code) x firm size x wage quartile cell, we compute the change in employment relative to January 4 to 31, 2020, and the change in employment relative to July 2020. Let Min Normed $\text{July}_{c,i,s,q}$ be the smallest value of indexed employment we observe at each date relative to its mean level over July 2020 and Max Normed $\text{January}_{c,i,s,q}$ be the largest value of indexed employment we observe at each date relative to its mean level over January 4 to 31, 2020.

For county x industry x firm size x wage quartile cells with at most 50 employees at all points between January 2020 and the end of the series, we reduce the weight we place on the series if we observe large changes in employment that indicate firm entry or exit. In particular, we compute the weight on the cell for county c, industry i, firm size s, and wage quartile q as:

$$\begin{aligned}\text{Weight}_{c,i,s,q} =& \max\Big\{1 - 1\{\text{Min Normed July}_{c,i,s,q} \leq 50\} \times (50 - \text{Min Normed July}_{c,i,s,q}) \times 0.02 \\ & - 1\{\text{Max Normed January}_{c,i,s,q} \geq 4000\} \times (\text{Max Normed January}_{c,i,s,q} - 4000) \times 0.001, \\ & 0\Big\}\end{aligned}$$

That is, we reduce the weight we place on the cell by two percentage points for each percentage point of decline we observe below 50 percentage points relative to July 2020. We further reduce the weight by 0.1 percentage points for each percentage point of growth we observe above 4000 percentage points relative to January 2020.

For county x industry x firm size x wage quartile cells with over 50 employees at any point between January 2020 and the end of the series, we reduce the weight we place on the series using more stringent restrictions, reflecting the fact that extreme growth rates are more anomalous in larger cells. In particular, we compute the weight on the cell for county c, industry i, firm size s, and wage quartile q as:

$$\begin{aligned}\text{Weight}_{c,i,s,q} =& \max\Big\{1 - 1\{\text{Min Normed July}_{c,i,s,q} \leq 50\} \times (50 - \text{Min Normed July}_{c,i,s,q}) \times 0.02 \\ & - 1\{\text{Max Normed January}_{c,i,s,q} \geq 600\} \times (\text{Max Normed January}_{c,i,s,q} - 600) \times 0.005, \\ & 0\Big\}\end{aligned}$$

That is, we reduce the weight we place on the cell by two percentage points for each percentage point of decline we observe below 50 percentage points relative to July 2020. We further reduce the weight by 0.5 percentage points for each percentage point of growth we observe above 600 percentage points relative to January 2020.

In addition, we address the especially concentrated entry and exit of firms from the sample at the end of each calendar year, where there is significant churn as firms renew their payroll processing contracts. Due to this seasonal pattern, the raw Paychex data display a downwards trend in employment at the end of each calendar year as some clients leave Paychex, followed by an upward trend in employment at the very beginning of each calendar year as new clients join. To avoid this source of error, we adjust for the end-of-year pattern in the Paychex data using data from the end of 2019. For each date between November 15, 2020 and January 10, 2021, using Paychex data on employment at the national level, we compute the change in employment relative to November 15, 2020 at the two-digit NAICS code x wage quartile level. We also compute the change in employment between the corresponding day in the previous year and November 15, 2019. We divide the change in employment relative to November 15, 2020 by the corresponding change in employment the previous year relative to November 15, 2019. To avoid a break in the series at January 10, 2021, we adjust the series from January 10 onwards using the adjusted level on January 10 and the unadjusted trend from January 10. We repeat this end-of-year adjustment during each subsequent year. We then apply the same adjustment to each two-digit NAICS code x wage quartile cell at the state, county and city levels.

Adjusting for minimum wage changes. Finally, we address mismeasurement of the change in employment in the first and second wage quartiles due to minimum wage increases in California, Massachusetts, and New York during our period of study. By increasing the minimum wage above the threshold between the first and second wage quartiles at the end of 2020 or start of 2021, these states mechanically shift workers out of the first wage quartile into the second wage quartile, leading to spurious changes in the quartile-specific series.[39] To address this issue, we make the following adjustments. We begin by calculating the percentage change in the number of employees in each state x industry (2-digit NAICS) cell from December 4, 2020 onwards for the first two wage quartiles combined. From December 4, 2020 onwards, for the first two wage quartiles in minimum wage change states, we impute the trend in each county x industry x wage quartile cell using the below-median income employment trend in their own state x industry cell. Since employment in the first wage quartile recovered less than employment in the second wage quartile in 2021, an imputation that aggregates the trends in the first and second wage quartiles tends to overstate the recovery in the first quartile and understate the recovery in the second quartile. As such, we further rescale the state x industry below-median income trend separately for each wage quartile, using the coefficient from the

[39] Arizona also increased its minimum wage during this period, but its minimum wage (\$12.15 in 2021; \$12.80 in 2022) remained well below the threshold between the first and second wage quartiles (\$13.25 in January 2021; \$13.88 in January 2022), so we do not adjust the employment series for this state.

regression of that quartile's employment change on the below-median income employment change in non-minimum wage-change states without a constant between December 4, 2020 and December 3, 2021. We then aggregate the adjusted employment counts to the relevant geographies (e.g. state; national) before calculating the change in employment since January 2020.

Paychex and Intuit: Combined Employment Series

We combine Paychex and Intuit data to construct our primary employment series. Our data sharing agreements requires us to produce combined estimates in our public releases of employment data.

Imputing local composition of Intuit data. Because Paychex is disaggregated by sector and covers all sectors and wage levels fairly comprehensively, we use it as the base for the combined employment series. We then use Intuit to refine the series in cells represented by those datasets. Intuit provides us with overall national industry shares as of 2019, but does not disaggregate the monthly employment data we receive by wage level or industry. We therefore impute the Intuit data to wage-industry cells before combining 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. For each date, we impute the number of Intuit employees in a given geography x industry x wage quartile cell as the number of Paychex employees in that cell reweighted by the national Intuit industry distribution. We then weight the geography-level changes in employment in the Intuit data by the imputed number of employees in its geography x industry x wage quartile cell.

Combining Paychex and Intuit data. We take a weighted average of the Paychex data and the imputed Intuit data 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. We report seven-day moving averages of these series, expressed as a percentage change relative to January 4 to 31, 2020.

Earnin

Sample Restrictions. We 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.

Finally, we omit the first and last 32 weeks of data for each user from our analysis to mitigate non-random entry and exit from the customer base. We observe that a person's probability of finding or losing a job is elevated near the time they entered or exited the sample—changes in employment likely induce people to sign up for the service or reconsider their need for the service. However, after omitting the first and last 32 weeks, the probability of a user finding or losing a job is no longer correlated with the number of weeks that elapse after entry into the sample or before exiting the sample.

Mapping paychecks to chained employment. We construct an employment series in the Earnin data from our analysis sample as follows. In the paycheck-level data, we observe the worker's paycycle frequency. As in the Paychex data, we use paycycle frequency to construct an employment series by assuming that workers are employed throughout the full duration of their paycycle. That is, we assume that a worker paid every two weeks has been fully employed for the two weeks prior to receiving their paycheck. We exclude workers with pay periods greater than 3 weeks, omitting approximately 1% of the sample. To account for the delay in receipt of paychecks, we shift the Earnin series back by one week. We then take the count of employed individuals across the Earnin sample as our measure of employment. We use that to calculate an indexed measure of week-over-week employment percentage change, for each geographic unit, indexed to their first week in the sample. Then, the average percentage change in the four weeks from January 4 to 31, 2020 is set as the reference value and indexed to zero. The employment series is then expressed as a change relative to January 4 to 31, 2020. We suppress estimates for ZIP codes with fewer than 10 paychecks observed over this period.

E.3 Masking and Publication

In the Paychex and Intuit combined series, we suppress cells in a few cases where Intuit data do not provide coverage for a given geographical region or industry. We also suppress cells in which Paychex records fewer than 150 monthly employees in January 4 to 31, 2020 at the geography x wage quartile or geography x industry (2-digit NAICS) x wage quartile level, depending on the series. When aggregating employment series to the geographical level without breakdowns by industry or wage quartile, however, we use data from all cells, without masking.

E.4 Accounting for Low-Wage Employment Changes due to Wage Growth

To calculate the share of low-wage employment changes due to wage growth, we decompose the level of the bottom-quartile-wage series from Figure V.A at the end of December 2021 into the part explained by real wage growth and not. We perform this decomposition by estimating ventile-specific wage growth rates in the CPS, calculated as the month-to-month change in ventiles, reweighting each monthly CPS sample to hold constant the distribution of industry, occupation, race, gender, age, education, region, and citizenship. This follows the procedure in Gould and Kandra (2022). Under the assumption that all remaining wage changes within these buckets reflect wage growth rather than changes in the nature of the jobs themselves, this procedure will estimate wage growth at each ventile of the wage distribution. There were extreme movements into and out of the labor force between January and July 2020, for which the reweighting may not adjust correctly. We thus employ our estimation procedure to estimate wage growth only between July 2020 and December 2021. We impute wage growth rates between January and July 2020 using a 2.9% annualized wage growth rate (which prevailed in 2019). Grigsby et al. (2021) suggest that this is likely an upper bound on the wage growth actually experienced during the early months of the pandemic, so that our overall wage growth adjustment will also be an upper bound on the changes in bottom-wage-quartile employment explained by wage growth.

With ventile-specific wage growth rates in hand, we then adjust the distribution of wages in the January 2020 CPS and recalculate total employment below the bottom-quartile-wage threshold for December 2021 ($13.79). We additionally add uniform noise between [-$0.50, $0.50] to whole number wages to smooth out spikes in the wage distribution at whole numbers, to mirror our treatment of these integer-wage spikes in the Paychex data (see Section II.B.4 for more details on this issue). Bottom-quartile-wage employment falls in this counterfactual by 12pp, reflecting the decline not explained by real wage growth; the remaining 7.7pp deficit in December 2021 is the share explained by real wage growth.

E.5 Benchmarking

Comparisons to QCEW and OES. Appendix Table III compares industry shares in each of the data sources above to nationally representative statistics from the Quarterly Census of Employment and Wages (QCEW). The Paychex-Intuit combined sample and the Earnin sample are broadly representative of the U.S. industry mix. Appendix Table IV shows that wage rates in our data sources are similar to nationally representative statistics from the BLS's Occupational Employment Statistics. Overall, our combined datasets appear to provide a representative portrait of private non-farm employment in the United States.

F Unemployment Claims Series Construction

F.1 Structure of Initial Data

We collect state-level unemployment initial claims data from the Department of Labor, Employment and Training Administration public data. We additionally collect data on state-level PUA, PUEC and continued claims from the Bureau of Labor Statistics public data. Finally, we gather county-level unemployment claims data from 22 state agencies who publish their state's data in various online formats, which we scrape and transform into a standard format.

F.2 Internal Data Processing

Harmonizing weekly and monthly estimates

In some cases, states only publish county-level data at the monthly frequency. To align with the weekly series of most UI data we collect, we impute weekly values for counties with monthly reports using the county-level monthly totals and the state-level distribution of weekly claims in that month, as published by the Department of Labor. This imputation implicitly assumes that each county within the state had the same relative distribution of UI claims within the month.

F.3 Masking and Publication

We apply no further masking beyond any masking applied by state and national agencies.

G COVID-19 Infections and Vaccinations Series Construction

G.1 Structure of Initial Data

We collect publicly available data on cases and deaths reported by the New York Times and the Centers for Disease Control and Prevention, publicly available data on hospitalizations from COVID-19 from the U.S. Department of Health and Human Services, publicly available data on the number of vaccine doses administered from the Centers for Disease Control and Prevention, and publicly available data on the number of COVID-19 tests from Johns Hopkins University.

G.2 Internal Data Processing

Reducing the influence of outliers

We manually review any spikes in cases, tests, or deaths that are larger than 25%. If news reports suggest that the spike is a reporting artifact, we smooth the data by imputing a value for the day of the spike using the growth rate in the outcome on the prior day.

G.3 Masking and Publication

We apply no further masking beyond any masking applied by state and national agencies.

H Time Outside Home Series Construction

H.1 Structure of Initial Data

We collect publicly available data on changes in time spent at various locations 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. Additionally we collect publicly available time use data from the American Time Use Survey.

H.2 Internal Data Processing

Generating measures of time at home and away from home

We use the American Time Use Survey to measure the mean time spent inside the home (excluding time asleep) and outside the home in January 2018 for each day of the week. We combine the data on time spent in and outside the home with Google's data on changes in time spent at and away from home. To do so we multiply the time spent inside the home in January 2018 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.

H.3 Masking and Publication

Google does not release data for geographies where their internal quality and privacy thresholds are not met.

I Educational Progress Series Construction

Zearn is a non-profit math curriculum publisher that combines in-person instruction with digital lessons. Zearn was used by approximately 970,000 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 use data from Zearn to measure educational progress during the pandemic.

I.1 Structure of Initial Data

We receive data from Zearn on the number of students using Zearn Math and student progress in Zearn Math as measured by the number of lessons on the platform completed by students in a given week for a given school.

I.2 Internal Data Processing

Reducing the effects of transitory outliers on the series

To reduce the effects of transitory outliers, 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.

Accounting for missing values

If a school is actively using Zearn at any point in a given semester, then for any week where we observe no reported values for both the number of students using Zearn Math and for student progress in Zearn Math, we impute these missing values as zeros.

I.3 Masking and Publication

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 preperiod of January 6 to February 7, 2020 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 which did not have at least 5 students using Zearn Math for at least one week during January 6 to February 7, 2020. After taking these steps, 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 to February 7, 2020, and we normalize relative to this base period to construct the indices we report.

I.4 Representativeness of Zearn Data

We assess the representativeness of the Zearn data in Appendix Table VIII by comparing the demographic characteristics of the schools for which we obtain 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, as measured in the American Community Survey. 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 students in the U.S.

J Dates and Geographic Definitions

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

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

- First U.S. COVID-19 Case: 2020-01-20
- National Emergency Declared: 2020-03-13
- CARES Act Signed in to Law: 2020-03-27

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 and verified from the New York Times reopening data, available here, and hand-collection from local news and government sources where needed.

- Stay-at-home order ends: Sourced and verified from the New York Times reopening data, available here, and hand-collected from local news and government sources where needed. Defined as the date at which the state government lifted or eased executive action or other policies instructing residents to stay home. We code "regional" and "statewide" expiry of stay-at-home orders separately. A "regional" expiration of a stay-at-home orders occurs when a stay-at-home order expires in one region within a state, but not everywhere within the state. A "statewide" expiration of a stay-at-home order occurs when a stay-at-home order first expired throughout a whole state, either due to a statewide change in policy, or due to the stay-at-home order in each county having expired.

- Partial business reopening: Sourced and verified from the New York Times reopening data, available here, and hand-collection from local news and government sources where needed. Defined as the date at which the state government allowed the first set of major industries to reopen (non-essential retail or manufacturing in nearly every case). Deviations from the New York Times reopening data are deliberate and usually involve our regional classification or our inclusion of manufacturing. A "regional" reopening occurs when businesses are allowed to reopen in one region within a state, but not everywhere within the state. A "statewide" reopening occurs when businesses are allowed to reopen throughout a whole state, either due to a statewide change in policy, or due to restrictions being eased in each individual county.

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

K Analysis of Economic Impact Payments

This appendix presents technical details for our analysis of the effects of the three rounds of economic impact (i.e., "stimulus") payments on consumer spending.

K.1 Construction of the Data

In order to prepare the spending data for analysis of the stimulus, we first aggregate the raw data to calculate total card spending on each date, in each ZIP code income quartile. See Section II.B.1 and Appendix B for more details on the source of the consumer spending series. Note that unlike in our baseline spending series, we use total spending (thereby including both intensive margin changes in spending per card and extensive margin changes in the number of cards) to ensure consistency in comparisons of our estimates across different stimulus rounds.

We then index spending in each ZIP code income quartile relative to spending between January 4 to 31, 2019. We include data from January 2020 in the "control" series, normed to January 2019, for the purpose of estimating the effects of the 2nd stimulus in January 2021.

We residualize indexed spending in each ZIP code income quartile with respect to day-of-week fixed effects calculated using 2019 data. For the April 2020 and March 2021 stimulus payments, we further residualize spending with respect to a linear pre-trend estimated in the pre-period data pooled across income quartiles; this procedure adjusts for economic trends related to the pandemic that differ from the prior year (and thus remain in the indexed data). We do not adjust for a linear pre-trend for the January 2021 stimulus payment due to the omission of the holiday period immediately before the stimulus payments.

Figure VII plots these data, differenced treatment minus control for each day relative to the stimulus events. Appendix Figure XXI depicts analogous plots without adjusting for pre-trends in any of the three stimulus payment periods. As discussed in Section IV.A, our analysis yields similar conclusions regardless of the linear adjustment for pre-trends.

K.2 Calculation of Effects of Stimulus Payments on Consumer Spending

We estimate the effects of each stimulus payment on indexed consumer spending using a difference-in-differences approach (Appendix Table VI). To capture the non-linear dynamics evident in the non-parametric figure, we estimate separate treatment effects for the first five days (Column 1) and from the 6th day onwards (Column 2). This window runs through the 25th day for the stimulus payments in April 2020 and March 2021, and through the 16th day for the stimulus payment in January 2021 (reflecting the data available at the time of our "real-time estimate"). Finally, we sum these effects to estimate the causal effect of the stimulus on spending the month after checks were sent out, under the identification assumption that daily trends in spending in 2020 would have matched those in 2019 (up to a linear difference in pre-trends for the first and third stimulus). We convert these estimates into a projected "first month" effect, measured in percentage points of baseline spending levels, by assuming the daily estimate from the second treatment coefficient within each stimulus continues through the 31st day (Column 3).

K.3 Calculating Spending Changes per Check Recipient

We map percentage point increases in spending into dollars per $1,200 check in three steps.

First, we translate the percentage point impact estimated from our difference-in-differences model (β) for each ZIP-income quartile into total additional dollars of spending. The rescaled total dollar effect $\tilde{\beta}$ is $\beta \times$ that ZIP-income quartile's share of total spending in the Affinity data in January 2019 $\times$ total card spending in NIPA Table 2.3.5 in January 2019 (following the method described in Appendix Figure V).

Second, we calculate the total amount of stimulus dollars for which residents of each ZIP-income quartile were eligible (S), which is the number of households in each set of ZIP codes times the eligibility rate. We calculate the number of households by household-type x income bin x ZIP code income quartile from the ACS 2014-2018, where household-type is the combination of single vs. married and number of children. We use total reported income in the ACS as a proxy for adjusted gross income (AGI), which determines eligibility for stimulus payments. We assume that mean household size is constant across income bins within each ZIP code for each household type (i.e. we assume the mean household size of married and unmarried households does not vary by income). This assumption permits us to combine these datasets to calculate the number of people in each income bin x household-type x ZIP code income quartile.

We then calculate the eligibility rate at the ZIP code level in three steps: (A) computing the share of people in single households who are eligible for stimulus payments; (B) computing the share of people in joint-filing households who are eligible for stimulus payments; and (C) combining these estimates using data on marriage rates to create overall ZIP code-level eligibility rates. We assign fractional eligiblity rates to reflect partial payments; for instance, a single household with no children eligible for a $600 check during the first stimulus (which paid $1,200 checks) would be assigned an eligibility rate of 50%.

(A) Eligibility among single households. Single households with incomes below $75,000 were eligible to receive the full stimulus amount. The precise structure of payments depends on

the number of children within the household; for simplicity, we use the structure of payments for families without children for all families. In practice, since high-income households with children faced more lenient eligibility requirements than we assume, our estimates represent an upper bound for consumption per stimulus recipient.

For the April 2020 stimulus, single households earning \$75,000-\$99,000 were eligible for partial stimulus payments. For the January 2021 and March 2021 stimulus payments, the partial eligibility income ranges were \$75,000-\$87,000 and \$75,000-\$80,000 respectively. However, the most granular income bins available in the ACS 2014-2018 data at a ZIP code level do not disaggregate between households earning \$75,000-\$99,999. As such, to assign eligibility rates for single households in this income bin, we assume that incomes are uniformly distributed within each ZIP code x income bin. We then assign these households an eligibility rate equal to $0.5 \times \frac{99{,}000-75{,}000}{99{,}999-75{,}000}$ for the April 2020 stimulus; $0.5 \times \frac{87{,}000-75{,}000}{99{,}999-75{,}000}$ for the January 2021 stimulus; and $0.5 \times \frac{80{,}000-75{,}000}{99{,}999-75{,}000}$ for the March 2021 stimulus.

We assign an eligibility rate of zero to single households earning above the partial eligibility income range.

(B) Eligibility among married households. We treat married households as joint-filers. Married households with incomes below \$150,000 were eligible to receive the full stimulus amount. As above, we do not account for the effects of number of children within the household on phase-out.

For the April 2020 stimulus, joint-filer households earning \$150,000-\$198,000 were eligible for partial stimulus payments. For the January 2021 and March 2021 stimulus payments, the partial eligibility income ranges were \$150,000-\$174,000 and \$150,000-\$160,000 respectively. The most granular income bins available in the ACS 2014-2018 data at a ZIP code level do not disaggregate between households earning \$150,000-\$199,999. As before, to assign eligibility rates for married households in this income bin, we assume that incomes are uniformly distributed within each ZIP code x income bin. We then assign these households an eligibility rate equal to $0.5 \times \frac{198{,}000-150{,}000}{199{,}999-150{,}000}$ for the April 2020 stimulus; $0.5 \times \frac{174{,}000-150{,}000}{199{,}999-150{,}000}$ for the January 2021 stimulus; and $0.5 \times \frac{160{,}000-150{,}000}{199{,}999-150{,}000}$ for the March 2021 stimulus.

We assign an eligibility rate of zero to married households earning above the partial eligibility income range.

(C) ZIP Code-level eligibility rates. Steps (A) and (B) allow us to assign eligibility rates to each income bin x ZIP code x household type (i.e. single vs. married). We then turn to calculating

the composition of household types within each income bin. To do so, we first regress marriage rates on median household income at the ZIP code level, weighting by population; we assume the ZIP code-level relationship between marriage rates and household income approximates the individual-level relationship between marriage rates and household income. We then assign a marriage rate equal to the fitted value of marriage rates at the midpoint of each household income bin. This allows us to estimate eligibility rates within each ZIP code x income bin. Finally, we calculate a population-weighted mean eligibility rate within each ZIP code income quartile. After that, we multiply this rate by $1,200 and by the population within each ZIP code income quartile to calculate total expected stimulus spending in each ZIP code income quartile.

Third, we adjust for the actual fraction of stimulus checks paid during the treatment period (f). To do so, we calculate the share of the stimulus payments distributed by the first week of the treatment period using Daily Treasury Statements. For instance, the total amount disbursed for stimulus payments under the CARES Act of April 2020 is $271.4 billion (Internal Revenue Service 2022). We use Daily Treasury Statements to calculate that total spending on stimulus payments prior to April 21, 2020 was roughly $174 billion. We therefore find that roughly 64% of payments had been distributed by April 21. A similar calculation for payments made under the COVID-related Tax Relief Act of 2020 yields estimates that roughly 91% of payments were distributed by January 10, 2021, and that 59% of payments made under the American Rescue Plan had occurred by March 23, 2021. We assume that these rates are constant across ZIP code income quartiles; it is possible that this fraction is higher for higher income households, due to the larger number of low-income households without bank information on file at the IRS, in which case we would overestimate the total stimulus payments in the lowest ZIP code income quartile, leading to an underestimate of the amount spent per recipient for these households.

Combining these three steps, our estimate of the dollars spent per stimulus payment is $\frac{\tilde{\beta}}{S \cdot f}$, reported in Appendix Table VI, Column 4. Our final estimate of the dollars spent per $1,200 received is $\frac{\$1,200}{\$X} \cdot \frac{\tilde{\beta}}{S \cdot f}$, where $X is the full stimulus payment amount for each stimulus ($1,200 for the April 2020 stimulus; $600 for the January 2021 stimulus; and $1,400 for the March 2021 stimulus). These estimates are reported in Figure VIII and Appendix Table VI, Column 5.

L Supplemental Policy Analyses

L.1 State-Ordered Reopenings

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 examine how these executive orders affected economic activity by exploiting variation across states in the timing of shutdowns and reopenings.

Throughout this section, we define the reopening date to be the day that a state *began* the reopening process (see Appendix J for details). In most states, reopening was a gradual process in which certain industries and types of businesses opened before others, but there was a lot of heterogeneity across states in the precise form that the reopening took. Our estimates should therefore be viewed as an assessment of the average impact of typical reopening efforts on aggregate economic activity; we defer a more detailed analysis of how different types of reopenings affected 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 2020 (New Mexico on March 24, Colorado on March 26). Colorado then partially reopened its economy, permitting retail and personal service businesses to open to the public, on May 1, 2020, while New Mexico did not reopen until two weeks later, on May 16. Appendix Figure XXVI.A plots consumer spending (using the Affinity Solutions data) in Colorado and New Mexico. Spending evolved nearly identically in these two states: in particular, there is no evidence that the earlier reopening in Colorado boosted spending during the two intervening weeks before New Mexico reopened.

Appendix Figure XXVI.B generalizes the case study in Appendix Figure XXVI.A by studying partial reopenings in the five states that issued such orders on or before April 27, 2020. For each reopening date (April 20, 24 and 27), we compare the trajectory of spending in treated states to a group of control states that had not reopened as of three weeks after the treated state reopened. We select multiple control states (listed in Appendix Table X) for each of the reopening dates by matching on pre-period indexed spending (relative to January) during the three weeks prior to reopening. Specifically, for each reopening date t, we estimate each state's rank (ranging from 0 to 1) in the distribution of mean indexed spending pooling weeks $t-1$,

$t-2$, and $t-3$. We then select control states with rank within 0.2 of the treated states' mean rank, separately for each reopening 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 three event studies together (redefining calendar time as time relative to the reopening date) to create Appendix Figure XXVI.B.

As in the case study of Colorado vs. New Mexico, the trajectories of spending in the treated states almost exactly mirror those in the control states. We formalize the estimate from this design using a difference-in-differences (DD) design that compares the two weeks before the reopening in the treated states and two weeks after. We estimate that reopenings led to a 1.27 percentage point increase in spending. This DD estimate also appears in Appendix Table XI, Column 1. Column 2 replicates that specification with a three-week analysis window; the DD estimate is virtually unchanged at 1.25 percentage points. Appendix Figure XXVI.B shows that we also find little impact of reopenings on employment (using the Paychex-Intuit data). Finally, Appendix Figure XXVI.B also shows (using data from Womply) that there was a 3.74 percentage point increase in the fraction of small businesses open after states allowed businesses to reopen – confirming that state orders did have some mechanical impact on the fraction of businesses that were open. However, this mechanical effect does not appear to translate to noticeable impacts on total employment or spending.

In line with these small treatment effect estimates, reopenings accounted for a relatively small share of the overall variation in economic conditions across states. To demonstrate this, we first calculate the actual variance in spending levels and other outcomes across states. We then counterfactually add our estimated effect of reopening to all states that were not yet open as of May 18, 2020, and recalculate the variance. Appendix Figure XXVI.C then plots 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 observed on May 18. These ratios are very low, showing that early reopenings did not play an important role in explaining why some states had stronger employment trajectories than others.[40] These results are consistent with the findings of other contemporaneous studies showing that little of the state-level variation in employment, job vacancies, or time spent outside home is related to state-level stay-at-home orders or business closures (Bartik et al. 2020; Forsythe et al. 2020;

[40] We emphasize that these results apply to *average* employment rates 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).

Lin and Meissner 2020; Goolsbee and Syverson 2021; Sears et al. 2023).

Why did these reopenings have so little immediate impact on economic activity? The evidence in Section III suggests that health concerns among consumers were the primary driver of the sharp decline in economic activity in March and April 2020. Consistent with that evidence, spending fell sharply in most states *before* formal state closures (Appendix Figure XXVII). If individuals' own health concerns are the core driver of reductions in spending during pandemics, governments may have limited capacity to mechanically restore economic activity through reopenings if those reopenings are not interpreted by consumers as a signal of reduced health risks.[41]

L.2 Paycheck Protection Program: Loans to Small Businesses

The Paycheck Protection Program (PPP) sought to reduce employment losses by providing financial support to small businesses. Congress appropriated nearly $350 billion for loans to small businesses in an initial tranche paid beginning on April 3, 2020, followed by another $175 billion in a second round beginning on April 27, 2020. The program offered loan forgiveness for businesses that maintained sufficiently high employment (relative to pre-crisis levels).

According to the House Committee on Small Business (2020), the stated primary purpose of the PPP was to encourage businesses to maintain employment even as they lost revenue. The Small Business Administration (2020) emphasized the employment impacts of the PPP as a key measure of the program's success, noting that the PPP "ensure[d] that over approximately 50 million hardworking Americans stay[ed] connected to their jobs" based on self-reports of the number of jobs retained by firms that received PPP assistance.

Here, we study the marginal impacts of the PPP on employment directly using payroll data from Paychex and Earnin, 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. One important exception to this rule was the food services industry, which was treated differently because of the prevalence of franchises. We therefore omit the food services sector from the analysis that follows.[42]

[41] In this vein, we stress that our research design only identifies the impacts of individual states opening earlier vs. later; if one state's actions impact behavior in other states (e.g., by shaping perceptions about health risks), the total impacts of shutdowns or reopenings at a national level could be larger. Moreover, these conclusions only apply to the initial stages of the pandemic that we study here. If health concerns diminish over time (e.g., due to quarantine fatigue), government restrictions could have larger effects on economic activity.

[42] According to SBA data on PPP receipt throughout the life of the program, 10.5% of total PPP loan volume

We estimate the causal effect of the PPP on employment rates at small businesses using a difference-in-differences research design, comparing trends in employment for firms below the 500 employee cutoff (the treated group) vs. those above the 500 employee cutoff (the control group) before vs. after April 3, 2020, when the PPP program began.[43] We do not condition on firm survival and simply count the total number of employees still working in each week at firms that initially had more than 500 vs. less than 500 employees. Our estimates thus take both the intensive (reductions in employees for surviving firms) and extensive (firm closure) margins into account.

Appendix Figure XXVIII.A plots the average change in employment rates (inferred from payroll deposits) relative to January 2020 for firms employing 100-499 employees, which were eligible for PPP loans, vs. firms employing 500-799 employees, which were generally ineligible for PPP loans, combining data from Paychex and Earnin.[44] To adjust for the fact that industry composition varies across firms of different sizes, we reweight by two-digit NAICS code so that the distribution of employees across industries in the below-500 and above-500 employee groups matches the overall distribution of employees across industries in January 2020. We further control for county x wage quartile x week fixed effects to account for the differential time patterns of employment rates by county and wage quartile shown in Section III.C.

Before April 3, 2020, trends in employment were similar among eligible vs. ineligible firms, showing that larger businesses provide a good counterfactual for employment trends one would have observed in smaller firms absent the PPP program (conditional on the reweighting and controls described above). After April 3, employment in the treated (< 500 employees) and control (≥ 500 employees) groups diverge and follow slightly different trajectories until August 2020, after which employment rates in the two groups are essentially identical again. These findings imply that the PPP program had little marginal impact on employment at small businesses under the identification assumption that employment trends in the two groups would have remained similar absent the PPP.

(7.1% of the total number of loans) was disbursed to firms in the food services sector (NAICS 72). The remaining exceptions to this rule affect relatively few workers: omitting food services, more than 90% of employees work at firms that face the 500 employee threshold for eligibility.

[43]Firms with more than 500 employees were still eligible for the Employee Retention Credit (ERC), which gave all firms that lost more than 50% of their revenue a tax credit worth up to $5,000 per employee if they did not take up the PPP. While data on ERC takeup are unavailable, fewer than 10% of CFOs of large firms report revenue losses larger than 25% (PwC 2020), suggesting that the vast majority of firms with more than 500 employees were not eligible for the ERC and hence serve as a valid counterfactual for employment in the absence of government assistance.

[44]We report estimates pooling Paychex and Earnin because our data use agreements do not permit us to report results based solely on Paychex data, and Intuit does not have coverage around the 500 employee cutoff.

Appendix Figure XXVIII.B plots the change in employment from January 4 to 31, 2020 to June 2020 by firm size bin. The decline in employment is quite similar across firm sizes, and is not markedly smaller for firms below the 500 employee eligibility threshold.[45]

In Appendix Table XII, we quantify the impacts of the PPP using OLS regressions of the form:

$$\text{Emp}_{scqit} = \alpha_{cqt} + \delta\text{Eligible}_s + \beta_{DD}\text{Eligible}_s \cdot \text{Post-PPP}_t + \varepsilon_{scqit}, \tag{1}$$

where Emp_{scqit} is the change in employment within each eligibility group $s \times$ county $c \times$ wage quartile $q \times$ 2-digit NAICS industry i cell on week t, relative to January 4 to 31, 2020; Eligible_s is an indicator variable for whether firm had fewer than 500 employees in the pre-COVID period; Post-PPP_t is an indicator variable for the date being on or after April 3, 2020; and α_{cqt} represents a county-wage quartile-week fixed effect. We estimate this regression on the sample of firms with 100-799 employees using data from March 11 to August 15, 2020. We focus on employment impacts up to August 15 because Appendix Figure XXVIII.A suggests that employment rates in the two groups converged after early August (extending the estimation window would only further reduce the estimated impacts of the PPP). We reweight by two-digit NAICS code so that the distribution of employees across industries in the below-500 and above-500 employee groups matches the overall distribution of employees across industries in January 2020. We cluster standard errors at the county-industry-eligibility group level to permit correlation in errors across firms and over time within counties and estimate the regression using OLS, weighting by the total number of employees in the cell from January 4 to 31, 2020.

Column 1 of Appendix Table XII presents the baseline estimate obtained from regression equation (1) of $\beta_{DD} = 2.48$ percentage points (s.e. = 2.71), an estimate that matches Appendix Figure XXVIII.A and is similar to that obtained in confidential ADP data in contemporaneous work by Autor et al. (2022a). The mean decline in employment among firms in the control group up to August 15, 2020 was 14.2 percentage points, implying that the PPP saved 17.5% of the jobs that would otherwise have been lost between April and August 2020. In Column 2, we reduce the bandwidth to focus more narrowly around the 500-employee size threshold; the

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

estimate is attenuated but not statistically distinguishable from that in Column 1.

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

Measurement Error in Firm Sizes. Our measures of firm size – which are based on employment levels in 2019 from Dun & Bradstreet data for the Paychex sample and ReferenceUSA in the Earnin sample – do not correspond perfectly to the measures used by the Small Business Administration to determine PPP eligibility. Such measurement error in firm size attenuates the estimates of β_{DD} obtained from (1) relative to the true causal effect of PPP eligibility because some of the firms classified as having more than 500 employees may have actually received PPP (and vice versa).

We estimate the degree of this attenuation bias by matching our data on firm sizes to data publicly released by the Small Business Administration (SBA) on a selected set of PPP recipients and assessing the extent to which firms are misclassified around the threshold. We restrict attention to firms receiving loans of at least $150,000, as the names and addresses of these firms are publicly available from the SBA. We first geocode addresses recorded in SBA and ReferenceUSA-Dun & Bradstreet data to obtain a latitude and longitude for each firm. We then compute the trigram similarities between firm names for all SBA and Dun & Bradstreet firms within twenty-five miles of another. We then select one "match" for each PPP recipient from the Dun & Bradstreet data for Paychex sample and ReferenceUSA for the Earnin sample, among the subset of firms within twenty-five miles. For firms with loans of above $150,000, exact loan size is not observed; we impute loan size as the midpoint of loan range. The SBA released firm names and ZIP codes of PPP recipients receiving over $150,000 in loans, which represent 72.8% of total PPP expenditure. Of the roughly 660,000 PPP recipients of these loans, we merge around 60% of firms and 62% of total expenditure to firm size data. In this matched subset, we find that mean PPP expenditure per worker is $2,303 for firms we classify as having 100-499 employees and $586 per worker for firms with 500-799 employees (excluding firms in the food services industry). Given that we match only 62% of the publicly available PPP expenditure to our data and the publicly available data covers only 73% of total PPP expenditure, this implies that firms measured as having 100-499 employees in our sample received $\frac{\$2{,}303}{0.62\times 0.73} = \$5{,}090$ of PPP assistance per worker, while firms with 500-799 employees received $\frac{\$586}{0.62\times 0.73} = \$1{,}290$ in PPP assistance per worker.[46] We calculate that PPP assistance to eligible firms with between 100 and 799 employees (excluding NAICS 72) is $\$5{,}092$ per worker on average.[47] Hence, firms with 500-799 workers in the ReferenceUSA-Dun & Bradstreet data (the control group) were effectively treated at an intensity of $\frac{\$1{,}290}{\$5{,}092} = 25.3\%$, whereas firms with 100-499 workers in the ReferenceUSA-Dun & Bradstreet data (the treatment group) were treated at an intensity of $\frac{\$5{,}090}{\$5{,}092} = 100\%$. Inflating our baseline reduced-form estimates by $\frac{1}{(1-0.253)} = 1.35$ yields estimates of the treatment effect of PPP eligibility adjusted for attenuation bias due to mismeasurement of firm size.

Under standard assumptions required to obtain a local average treatment effect in the presence of non-compliance – no direct effect of being classified as having more than 500 workers independent of the PPP and a monotonic treatment effect – we can estimate the LATE of the PPP on employment rates by multiplying the raw estimates reported in Appendix Table XII, Column 1 by 1.35 (Angrist, Imbens, and Rubin 1996). This gives us a final preferred point estimate for the effect of PPP eligibility on employment of 3.3 percentage points.

Costs Per Job Saved. Using Statistics of U.S. Businesses (SUSB) data, we calculate that approximately 62.4 million workers work at firms eligible for PPP assistance (53.7 million workers excluding those in the food services industry, NAICS 72). Thus 86.1% of total PPP expenditure was received by non-NAICS 72 firms. We then multiply this share by total PPP expenditure as of August 8, 2020 to reach an estimate of $486 billion in non-NAICS 72 firms. Under the assumption that the PPP's effects on firms with between 100 and 499 employees were the same in percentage point terms as the PPP's effects on all eligible firms, our baseline

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

[47] To compute this statistic, we first calculate the share of total loan amounts received by non-NAICS 72 firms in the publicly released SBA data. We begin by imputing precise loan amount as the midpoint of minimum and maximum of loan range, where precise loan amount is not released. We then calculate the share of loans in firms with firm size between 100 and 499, in NAICS codes other than NAICS 72, under the assumption that our merge rate is constant by firm size. Using this approach, we calculate that 13.1% of PPP loan spending was allocated to non-NAICS 72 firms with 100-499 employees. We then rescale the total PPP expenditure to the end of June 2020, $521 billion, by 0.131 to arrive at an estimate of $68.25 billion in PPP loan spending to non-NAICS 72 firms with 100-499 employees. Finally, we divide $68.25 billion by the number of workers at non-NAICS 72 firms with 100-499 employees to arrive at an estimate of loan spending per worker.

estimates in the combined Paychex-Earnin data (Appendix Table XII, Column 1), adjusted for attenuation bias, imply that the PPP saved 0.03×53.6M = 1.61 million jobs from April through August 15, 2020.[48] Given a total expenditure on the PPP program of $486 billion through August 8 (excluding firms in food services), this translates to an average cost per job saved (over the five months between April and August 2020) of $301,863. Even at the upper bound of the 95% confidence interval for employment impact, we estimate a cost per job saved of $86,201.

In order to compute net costs to government per job saved, we account for the fact that a reduction in job losses decreases UI spending. We evaluate replacement rates at the mean level of earnings for workers employed at PPP-eligible firms using the statutory rates in Ganong, Noel, and Vavra (2020, fig. 3a), which estimates that displaced workers received roughly 120% of weekly earnings for the seventeen weeks between the beginning of our treatment period (April 3, 2020) and the end of July 2020, and received roughly 40% of weekly earnings for the following two weeks until the end of our analysis window (August 15, 2020). Computing expenditure on UI given these replacement rates and mean earnings, we find that the effect of each job saved by the PPP on UI payments was $18,350 over our analysis period. Netting these savings out of the gross cost, we estimate a net cost to the government of $283,513 per job saved (and $67,851 at the upper bound of the 95% confidence interval for employment impact). For comparison, mean annual earnings for workers at PPP-eligible firms are only $45,000.

[48] If the treatment effect of the PPP program on food services were the same in percentage terms as in other sectors, we estimate the PPP saved a total of 1.87 million jobs.

Appendix Table I
Data Sources and Processing Steps

Data Series	Sources	Description of Processing	Overview of Processing Steps
Consumer Spending	Affinity Solutions	Appendix B	Adjust for sharp changes in the client base. Adjust for gradual changes in the client base. Address spurious changes. Control for seasonal fluctuations in spending.
Small Business Revenue and Openings	Womply	Appendix C	Adjust for the evolving client base. Produce the ZIP code series. Reduce the influence of outliers. Control for seasonality.
Job Postings	Lightcast	Appendix D	Reduce the influence of outliers.
Employment	Paychex Intuit Earnin	Appendix E	*Paychex* Address anomalous spikes. Adjust when quartile thresholds cross integer wages. Map paychecks to employment periods. Adjust for an evolving client base. Adjust for minimum wage changes. *Paychex and Intuit Combined* Impute local composition of Intuit data. Combine Paychex and Intuit data. *Earnin* Sample restrictions. Map paychecks to chained employment.
Unemployment Claims	Department of Labor, Employment and Training Administration Bureau of Labor Statistics State Departments of Labor	Appendix F	Harmonize weekly and monthly estimates.
COVID-19 Infections and Vaccinations	Centers for Disease Control and Prevention New York Times Johns Hopkins Coronavirus Resource Center U.S. Department of Health and Human Services	Appendix G	Reduce the influence of outliers.
Time Outside Home	Google COVID-19 Community Mobility Reports American Time Use Survey	Appendix H	Generate measures of time at home and away from home.
Educational Progress	Zearn	Appendix I	Reduce the effects of transitory outliers. Account for missing values.

Notes: This table provides an overview of the data series constructed in this paper and published at tracker.opportunityinsights.org. Further details are provided in Section II.B and in the appendices referenced in the table.

Appendix Table II
Distributions of Sample Sizes

Panel A: Spending, Small Business Revenue, Job Posts, Employment

	Distribution of Sample Across Local Areas				
	10th Percentile	25th Percentile	50th Percentile	75th Percentile	90th Percentile
Consumer Spending: Affinity					
January 2019 Spending (County level)	$ 19,657	$ 111,788	$ 454,612	$ 1,409,311	$ 3,750,812
Small Business Revenue: Womply					
January 2020 Revenue					
County level	$ 774,550	$ 1,994,485	$ 5,543,055	$ 12,800,000	$ 36,000,000
ZCTA level	$ 10,703	$ 18,577	$ 35,511	$ 65,762	$ 113,147
Job Postings: Lightcast					
January 2020 Job Posts (County level)	34	171	1,076	4,066	11,895
Employment: Paychex-Intuit					
January 2020 Jobs (County level)	549	3,146	28,390	111,533	219,660
Employment: Earnin					
January 2020 Jobs					
County level	92	615	3,373	12,154	35,570
ZCTA level	18	61	171	371	667

Panel B: Education

	Distribution of Sample Across the Most Populous Counties				
	Top 5	Top 10	Top 20	Top 50	Top 100
Education: Zearn (County level)					
Average number of Zearn students	14,820	10,190	7,410	6,790	3,240
Share of Zearn students	8 %	11 %	16 %	22 %	35 %
Share of US population	8 %	13 %	19 %	30 %	42 %

Notes: This table presents the distribution of cell sizes for data series we construct from private sector sources. For Affinity Solutions, we report the population-weighted percentiles of the county-level distribution of the average daily total spending during January 4 to 31, 2019. For Womply, we report population-weighted percentiles of the county-level and ZIP-level distributions of weekly revenue in January 4 to 31, 2020. So, for example, the $774,550 under the column "10th Percentile" means that the top 90% of the population lives in counties with at least $774,550 in average weekly small business revenue in the base period. For Lightcast, we report population-weighted percentiles of the county-level distribution of average weekly job posts in January 4 to 31, 2020. For Paychex-Intuit, we report population-weighted percentiles of the county-level distribution of employment in January 4 to 31, 2020. For Earnin, we report population-weighted percentiles of the county-level and ZIP-level distributions of employment in January 4 to 31, 2020. For Zearn, we report the average number of Zearn students and the share of Zearn students in the top 5; top 10; top 20; top 50; and top 100 counties by county population. We also report the share of the U.S. population in these counties for reference. Data sources: Affinity Solutions, Womply, Lightcast, Paychex, Intuit, Earnin, Zearn.

Appendix Table III
Industry Employment Shares Across Datasets

NAICS Code	NAICS Description	Industry Shares (%)			
		QCEW All Establishments	QCEW Small Establishments	Paychex + Intuit	Earnin
		(1)	(2)	(3)	(4)
11	Agriculture, Forestry, Fishing and Hunting	0.84	1.04	0.53	0.18
21	Mining, Quarrying, and Oil and Gas Extraction	0.55	0.43	0.15	0.21
22	Utilities	0.44	0.29	0.13	0.48
23	Construction	5.72	7.62	7.39	1.32
31-33	Manufacturing	10.27	5.16	8.05	6.33
42	Wholesale Trade	4.72	5.99	5.73	4.45
44-45	Retail Trade	12.48	14.06	7.81	23.98
48-49	Transportation and Warehousing	4.3	2.82	2.04	5.57
51	Information	2.29	1.64	1.45	4.13
52	Finance and Insurance	4.83	4.6	3.42	7.53
53	Real Estate and Rental and Leasing	1.71	2.9	3.2	1.94
54	Professional, Scientific, and Technical Services	7.63	8.97	14.16	4.91
55	Management of Companies and Enterprises	1.93	0.79	0.25	0.37
56	Administrative Support	7.25	5.3	7.07	4.91
61	Educational Services	2.39	1.53	2.36	3.85
62	Health Care and Social Assistance	16.16	13.16	15.41	16.72
71	Arts, Entertainment, and Recreation	1.78	1.64	1.87	1.57
72	Accommodation and Food Services	11.04	15.6	9.89	9.2
81	Other Services (except Public Administration)	3.57	6.21	9.09	2.34
99	Unclassified	0.11	0.24		

Notes: This table compares the industry composition (2-digit NAICS) of two payroll-based employment datasets to the Quarterly Census of Employment and Wages (QCEW), an administrative dataset covering the near-universe of firms in the United States. Each column displays the share of employees (in percentage terms) in the given dataset who work in the specified sector. Column (1) displays the industry composition of the QCEW in the first quarter of 2020. Column (2) replicates column (1) restricting to small establishments, defined as establishments with fewer than 50 employees. Column (3) shows the industry composition of Paychex-Intuit data in January 2020. To construct Column (3), we first separately calculate the number of employees in each 2-digit NAICS code in Paychex and Intuit as the number of worker-days in each 2-digit NAICS code. We then calculate combined Paychex-Intuit employment in each 2-digit NAICS code as a weighted sum of Paychex and Intuit, where the weights are undisclosed to meet privacy protection requirements. Column (4) displays the industry composition in the Earnin data in January 2020. Data sources: Paychex, Intuit, Earnin, QCEW.

Appendix Table IV
Hourly Wage Rates By Industry Across Datasets

NAICS Code	NAICS Description	Occupational Employment Statistics (OES) (1)	Paychex + Earnin (2)
11	Agriculture, Forestry, Fishing and Hunting	$16.35	$20.63
21	Mining, Quarrying, and Oil and Gas Extraction	$32.15	$32.89
22	Utilities	$39.80	$33.18
23	Construction	$27.87	$28.74
31-33	Manufacturing	$26.48	$25.44
42	Wholesale Trade	$28.85	$27.34
44-45	Retail Trade	$17.02	$21.07
48-49	Transportation and Warehousing	$24.33	$24.62
51	Information	$39.07	$32.78
52	Finance and Insurance	$36.73	$32.82
53	Real Estate and Rental and Leasing	$24.98	$25.66
54	Professional, Scientific, and Technical Services	$41.83	$34.37
55	Management of Companies and Enterprises	$42.59	$24.06
56	Administrative Support	$20.50	$23.51
61	Educational Services	$28.34	$24.78
62	Health Care and Social Assistance	$26.98	$25.47
71	Arts, Entertainment, and Recreation	$19.18	$22.47
72	Accommodation and Food Services	$13.65	$16.62
81	Other Services (except Public Administration)	$21.58	$22.50
	All	$25.72	$25.34
	Industry-Weighted Average of BLS Mean Wages		$26.00

Notes: This table compares mean wages in private sector datasets to mean wages in Occupational Employment Statistics (OES) data, within each two-digit NAICS code. Column (1) reports mean wages in each NAICS code in May 2019 OES data. We inflate these wages to 2020 dollars using the BLS Consumer Price Index. Column (2) reports mean wages in combined Paychex-Earnin data in January 2020. We first compute mean wages separately in Paychex and Earnin data as mean wages in January 2020, weighting by number of worker-days. In Paychex, wages are measured as pre-tax wages recorded by the employer. In Earnin, wages are post-tax wages recorded in payroll deposits. We then take a weighted mean of Paychex and Earnin wages within each industry, where the weights are not disclosed to meet business privacy requirements. The last row of Column (2) displays BLS mean wages, reweighted to match the 2-digit NAICS composition within the combined private sector dataset. Data sources: Paychex, Earnin, OES.

Appendix Table V
Consumer Spending on Debit and Credit Cards, by Income Quartile and Sector

	Mean Level of Card Spending Per Day ($ Billions)	Change Relative to Dec 30 2019-Jan 26 2020 ($ Billions)		
	Level as of Dec 30 2019-Jan 26 2020	Change as of April 14 2020	Change as of August 14 2020	Change as of December 31 2021
	(1)	(2)	(3)	(4)
Pooled Card Spending: All Income Quartiles, All Sectors	$23.6	-$7.3	-$0.9	+$3.6
Panel A: Card Spending, by ZIP Income Quartile				
Bottom Quartile	$3.5	-$0.9	+$0.1	+$0.8
	(15.0%)	(11.8%)	(-10.3%)	(21.7%)
Second Quartile	$5.3	-$1.5	-$0.1	+$1.0
	(22.3%)	(20.2%)	(11.3%)	(26.5%)
Third Quartile	$6.5	-$2.0	-$0.2	+$1.0
	(27.4%)	(27.3%)	(27.2%)	(28.3%)
Top Quartile	$8.3	-$3.0	-$0.6	+$0.8
	(35.3%)	(40.7%)	(71.8%)	(23.5%)
Panel B: Card Spending, by Sector				
Durable Goods	$3.7	-$0.5	+$0.6	+$0.8
	(15.6%)	(7.1%)	(-55.9%)	(21.1%)
Non-Durable Goods	$8.0	-$1.5	-$0.0	+$2.1
	(33.8%)	(20.5%)	(2.3%)	(55.2%)
Remote Services	$5.4	-$1.2	+$0.3	+$1.3
	(22.9%)	(15.7%)	(-29.9%)	(34.2%)
In-Person Services	$6.5	-$4.2	-$1.8	-$0.4
	(27.7%)	(56.7%)	(183.6%)	(-10.4%)
Panel C: In-Person Services Spending, by Sub-Sector				
Hotels & Food	$2.9	-$1.8	-$0.7	-$0.0
	(44.1%)	(43.1%)	(38.6%)	(5.0%)
Transportation	$1.0	-$0.8	-$0.6	-$0.5
	(15.9%)	(19.3%)	(34.6%)	(119.2%)
Health Care	$1.0	-$0.6	-$0.1	+$0.2
	(15.3%)	(13.8%)	(6.1%)	(-37.5%)
Recreation	$0.6	-$0.4	-$0.2	-$0.0
	(8.9%)	(9.8%)	(12.7%)	(12.4%)
Other In-Person Services	$1.0	-$0.6	-$0.1	-$0.0
	(15.8%)	(13.9%)	(8.0%)	(0.9%)

Notes: This table presents estimates of the changes in daily national consumer spending from the pre-pandemic baseline (measured as a daily average between January 6 to February 2, 2020) to April 14, 2020 (in Column 2), August 14, 2020 (in Column 3) and December 31, 2021 (in Column 4). We construct these estimates by combining statistics on total daily card spending in January 2020 (categories "Furnishings and durable household equipment", "Recreational goods and vehicles", "Other durable goods", "Food and beverages purchased for off-premises consumption", "Clothing and footwear", "Gasoline and other energy goods", "Other nondurable goods", "Transportation services", "Recreation services", "Food services and accommodations", "Financial services and insurance", and "Other services" from NIPA Table 2.3.5) with our consumer spending series from Affinity Solutions. See Section II.B.1 and Appendix B for more details on the construction of the consumer spending series. We estimate total daily spending (top row in Column 1) by dividing the NIPA monthly estimate by 31. We estimate the levels in the remaining rows of Column 1 by multiplying the spending share estimated in the Affinity data in January 2020, by total daily spending. We estimate the changes in Columns 2-4 in each row by multiplying the level in Column 1 by the change in the relevant Affinity series. The parentheses in each row report the share of the national spending decline (in Columns 2-4) or of the national spending total (in Column 1) accounted for by each subset. Panel A disaggregates spending by income quartile, measured at the ZIP code level using median household income from the 2014-2018 ACS. Panel B disaggregates spending by sector; Panel C disaggregates in-person service spending by sub-sector (see Appendix B.3 for sector and sub-sector definitions). Data sources: Affinity Solutions, NIPA.

Appendix Table VI
Effects of Stimulus Payments on Consumer Spending

	Dependent Variable: Change in Consumer Spending				
	First 5 Days Post-Stimulus (p.p. change)	After Day 5 Post-Stimulus (p.p. change)	Combined: 31 Days Post-Stimulus (p.p. change)	Combined: 31 Days Post-Stimulus ($ change)	Combined: 31 Days Post-Stimulus ($ change per $1200 stimulus)
	(1)	(2)	(3)	(4)	(5)
Panel A: April 2020 Stimulus Payments (Stimulus 1)					
Income Q1:	19.16	21.81	21.38	442.28	442.28
	(6.67)	(3.32)	(3.13)	(64.75)	(64.75)
Income Q2:	12.69	18.54	17.59	564.66	564.66
	(5.07)	(2.90)	(2.71)	(86.82)	(86.82)
Income Q3:	7.45	16.63	15.15	640.19	640.19
	(4.77)	(2.87)	(2.67)	(112.90)	(112.90)
Income Q4:	1.94	12.78	11.03	731.52	731.52
	(4.14)	(2.98)	(2.75)	(182.41)	(182.41)
Panel B: January 2021 Stimulus Payments (Stimulus 2)					
Income Q1:	14.56	4.83	6.40	93.56	187.12
	(5.56)	(2.42)	(2.39)	(34.94)	(69.87)
Income Q2:	9.46	3.69	4.62	105.41	210.83
	(4.28)	(2.20)	(2.13)	(48.58)	(97.15)
Income Q3:	6.36	2.13	2.81	85.28	170.56
	(3.85)	(1.95)	(1.91)	(57.89)	(115.78)
Income Q4:	3.39	-0.22	0.37	17.73	35.46
	(3.39)	(1.89)	(1.82)	(88.30)	(176.60)
Panel C: March 2021 Stimulus Payments (Stimulus 3)					
Income Q1:	26.90	11.54	14.02	317.29	271.97
	(6.66)	(3.87)	(3.78)	(85.47)	(73.26)
Income Q2:	15.98	5.64	7.31	259.24	222.21
	(4.36)	(3.20)	(3.05)	(108.12)	(92.67)
Income Q3:	8.59	2.76	3.70	175.11	150.09
	(3.41)	(3.12)	(2.93)	(138.89)	(119.05)
Income Q4:	2.89	0.57	0.95	72.65	62.27
	(2.67)	(3.33)	(3.08)	(236.80)	(202.98)

Notes: This table reports difference-in-differences estimates of the impacts of the three rounds of economic impact payments on consumer spending, separately for each ZIP income quartile. Panels A, B, and C respectively show estimates for the payments primarily made on April 15, 2020, January 4, 2021, and March 17, 2021. In Panels A and C, we use a window of 25 days before and after the stimulus payment date as the estimation sample. In Panel A, we exclude the partially treated date of April 14, 2020 and in Panel C, we exclude the partially treated dates of March 13 to 16, 2021. In Panel B, we use December 4 to 14, 2020 and January 4 to 19, 2021 as the estimation window; we exclude the intervening holiday period because of the high degree of volatility of spending during that period (Appendix Figure XXIII). We then regress daily consumer spending within the relevant window and the corresponding period starting in 2019 (after residualizing it on day-of-week fixed effects) on an indicator variable for the first five days of the post-period, an indicator variable for the rest of the post-period, and their interactions with an indicator for being in the treated group. In Panels A and C, we also adjust for a linear pre-trend in both the treatment and control series; we do not adjust for a linear pre-trend in Panel B due to the omission of the holiday period. The coefficients reported in columns 1 and 2 are those on the interaction of the two post-period indicators with the indicator for the treated year. Column 3 combines these two spending estimates to project the total percentage effect on spending for the first 31 days after the reform. Column 4 converts this combined percentage estimate into a total dollar estimate using base period daily average spending (see Appendix K for details) and column 5 rescales the dollar estimates to be per $1,200 of stimulus. Robust standard errors are reported in parentheses. Data source: Affinity Solutions.

Appendix Table VII
Association Between Changes in Consumer Spending and Workplace Rent by ZIP Code

Dep. Var.:	Change in Low-Income Consumer Spending (%)	
	(1)	(2)
Mean Workplace Two-Bedroom Rent (per thousand dollars)	-12.82	-13.55
	(1.85)	(7.05)
County Fixed Effects		X
N	8,974	8,974

Notes: This table presents results from regressions of changes in low-income consumer spending in the first month of the pandemic on the median rent of the workplaces of those low wage workers. We measure the dependent variable as the average value of our consumer spending index between March 25 and April 14, 2020 (see Section II.B.1 and Appendix B for details on the construction of this series). Unlike our baseline spending series, we compute these data for this table at the ZIP code level. We then construct the average workplace rent using the Census' LODES data (to measure the workplace ZIP codes for low-wage workers residing in each ZIP code) and the median rent data for each ZIP code from the 2014-2018 ACS. Our independent variable is the average of median workplace ZIP rents using the LODES workplace distribution as the weights for each residential ZIP code. We then restrict the sample for the regression to the residential ZIP codes in the (population-weighted) bottom quartile of median income. We scale the dependent and independent variables such that the coefficients represent the predicted change in percentage points for each $1000 increase in rent; for instance, the coefficient of -12.82 in Column (1) means that a $1000 increase in average workplace median rent for low-wage workers residing a given low-income ZIP code is associated with a 12.82 percentage point reduction in consumer spending in that ZIP code. Column (2) replicates the specification in Column (1) including county fixed effects. Standard errors are clustered at the county level and reported in parentheses. Data sources: Affinity Solutions, Census LODES, ACS.

APPENDIX TABLE VIII
Demographic Characteristics of Zearn Users

	Zearn Users (1)	U.S. Population (2)
Panel A: Income		
ZIP Median Household Income		
25th Percentile	43,750	45,191
Median	55,429	57,371
75th Percentile	71,967	76,080
Number of ZIP codes	6,749	33,229
Sample Population	971,843	326,274,368
Panel B: School Demographics		
Share of Black Students		
25th Percentile	1.7%	1.9%
Median	6.4%	6.8%
75th Percentile	22.1%	23.1%
Share of Hispanic Students		
25th Percentile	4.7%	5.2%
Median	11.9%	13.7%
75th Percentile	36.4%	37.0%
Share of Students Receiving FRPL		
25th Percentile	34.9%	31.5%
Median	56.6%	54.3%
75th Percentile	79.7%	79.0%
Number of Schools	18,437	88,459
Total Student Population of Schools	9,690,036	49,038,524
Panel C: Region		
Share of Students		
Midwest	23.9%	20.5%
Northeast	13.7%	15.6%
South	32.5%	39.6%
West	29.6%	23.7%

Notes: This table reports demographic characteristics for Zearn schools vs. the U.S. population. Panel A compares income characteristics of ZIP codes with Zearn coverage vs. all ZIP codes. We define Zearn to have coverage in a ZIP code if at least five students at schools in that ZIP code used Zearn between January 6 to February 7, 2020. Column 1 shows income characteristics of Zearn-covered ZIP codes. The first three rows in Panel A display the 25th, 50th, and 75th percentiles of ZIP-level median household income in Zearn-covered ZIP codes, as measured in the 2014-2018 ACS. The fourth and fifth rows of Panel A display the number of Zearn-covered ZIP codes, and the number of students using Zearn in those ZIP codes. Column 2 replicates Column 1 using all ZIP codes in the U.S. The fourth and fifth rows of Column 2 replicates Column 1 using all ZIP codes in the U.S. and counting the total population, respectively. Panel B presents the demographic composition of schools in the Zearn data (Column 1) and of all U.S. K-12 schools (Column 2), calculated using school-level data from the Common Core dataset as constructed by MDR Education, a private education data firm. The first three rows of Panel B show the 25th, 50th, and 75th percentiles of share of Black students in Zearn schools (Column 1) and in all US K-12 schools (Column 2). Rows 4-6 and 7-9 of Panel B replicate Rows 1-3 using the share of Hispanic students and the share of students receiving free or reduced-price lunch meals. Rows 10 and 11 of Panel B display the number of Zearn schools matched to the Common Core data and the number of students in those schools. Panel C compares the share of students by region in Zearn vs. the US population. Data sources: Zearn, ACS, Common Core.

Appendix Table IX
City to County Crosswalk

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

Notes: This table shows our metro area (city) to county crosswalk. We assigned metros to counties and verified 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 X
List of Re-Opening States and Control States for Event Studies

Date	States that Re-Opened	Consumer Spending Controls (1)	Employment Controls (2)	Small Businesses Open Controls (3)
April 20th 2020	South Carolina	Florida, Hawaii, Louisiana	California, Connecticut, Delaware, District Of Columbia, Florida, Indiana, Louisiana, New Mexico, Oregon, Washington, Wisconsin	Nebraska, South Dakota, Virginia
April 24th 2020	Alaska, Georgia	Delaware, Illinois, Massachusetts, Missouri, New Jersey, New York, Pennsylvania, South Dakota, Virginia, Wisconsin	California, Connecticut, Delaware, Indiana, Louisiana, Massachusetts, New Jersey, Pennsylvania, Washington	California, Connecticut, Delaware, Illinois, Louisiana, Maryland, Massachusetts, Missouri, New Jersey, Pennsylvania, Washington, Wisconsin
April 27th 2020	Minnesota, Mississippi	Illinois, Nebraska, New Jersey, New York, Pennsylvania, Virginia, Wisconsin	Delaware, District Of Columbia, Illinois, Maryland, New Mexico, Virginia, Wisconsin	New Mexico, South Dakota, Virginia

Notes: This table lists the treatment and control states for the analysis of state reopenings in Appendix Figure XXVI and Appendix Table XI. Column (1) displays the control states in the event study of consumer spending (as measured in the Affinity data) described in Appendix Figure XXVI and Appendix Table XI. Column (2) replicates Column (1) for employment (as measured in the Paychex-Intuit data). Column (3) replicates Column (1) for the number of small businesses open (as measured in the Womply data).

APPENDIX TABLE XI
Causal Effects of Re-Openings on Economic Activity

Dep. Var.:	Spending (%)		Employment (%)		Small Businesses Open (%)	
	(1)	(2)	(3)	(4)	(5)	(6)
DD Estimate of Effect of Reopening:	1.27 (0.48)	1.25 (0.60)	1.01 (0.42)	1.04 (0.47)	3.74 (1.29)	3.79 (1.63)
State-Week Observations:	100	150	128	192	92	138
Analysis Window (weeks on either side of reopening):	2	3	2	3	2	3
Mean Decline in Outcome (January to April 2020):	-30%		-22%		-29%	

Notes: This table estimates the effects of state reopenings on various outcomes using an event study design based on states that reopened non-essential businesses between April 20 to 27, 2020. Each state that reopens is matched to multiple control states (listed in Appendix Table X) that did not reopen within the subsequent 3 weeks but had similar trends of the outcome variable during the weeks preceding the reopening. We construct the control group separately for each re-opening day and then stack the resulting event studies to align the events. All estimates are from OLS regressions at the state x week level on an indicator variable for the state being a state that reopened, an indicator variable for the date being after the reopening date, and the interaction between these two variables. We report the coefficient and standard error on the interaction term, which we refer to as the difference-in-differences (DD) estimate of the effect of reopening. Standard errors are clustered at the state level and reported in parentheses. The dependent variable is rescaled to be in percentage terms such that, for example, the first row of Column (1) indicates that the difference-in-differences estimate for the effect of reopening on consumer spending over a two-week horizon is a 1.27 percentage point increase in consumer spending. The third row indicates the "Analysis Window" used in the regression: for example, the sample in column (1) is restricted to the two weeks before and after the date of reopening, whereas the sample in column (2) is restricted to the three weeks before and after the date of reopening. The last row shows the mean decline in the outcome variable across states from the period January 4 to 31, 2020 to the period March 25 to April 14, 2020 - except in Columns (1) and (2) where the reference period is January 6 to February 2, 2020. Columns (1) and (2) show the estimated effect of reopening on consumer spending using data from Affinity Solutions. Consumer spending is expressed as a percentage change relative to January 6 to February 2, 2020, and seasonally adjusted using 2019 data. Columns (3) and (4) replicate columns (1) and (2) using changes in employment as the dependent variable. Employment is calculated using Paychex-Intuit data and expressed as a percentage change relative to January 4 to 31, 2020. Columns (5) and (6) replicate columns (1) and (2) respectively using the number of small businesses open as the dependent variable, calculated using Womply data and expressed as a percentage change relative to January 4 to 31, 2020. Columns (1), (3), and (5) correspond to the specifications displayed in Appendix Figure XXVI.B. Data sources: Affinity Solutions, Paychex, Intuit, Womply.

Appendix Table XII
Causal Effect of the Paycheck Protection Program on Employment

Dep. Var.:	Change in Employment (%)	
	(1)	(2)
	Baseline Estimate (100-799 Employees)	Smaller Bandwidth (300-699 Employees)
DD Estimate	2.48	0.08
	(2.71)	(4.30)

Notes: This table reports difference-in-differences (DD) estimates of the effect of PPP eligibility (defined as the parent firm having fewer than 500 employees) on employment. The outcome variable is employment at the county x 2-digit NAICS x wage quartile x PPP eligibility x week level, excluding the Accommodations and Food Services Sector (NAICS 72), expressed as a percentage change relative to January 4 to 31, 2020. Both columns present regressions in combined Paychex-Earnin data. In the baseline estimate in column (1), we begin by restricting the sample to firms with 100-799 employees. We then reweight these cells so that employment shares by industry within each eligibility group match the overall employment shares by industry in January 4 to 31, 2020. Finally, we report estimates from an OLS regression of changes in employment on county x wage quartile x week fixed effects, an indicator for PPP eligibility (firm size < 500 employees), and an interaction term between PPP eligibility and an indicator for the date being after April 3, 2020 (the DD estimate). The sample for this regression is limited to weeks ending between March 11 and August 15, 2020. The DD estimate is the coefficient on the interaction term for PPP eligibility and the date being after April 3. We cluster standard errors (reported in parentheses) at the county x industry x eligibility level, and winsorize the dependent variable at the 99th percentile. We use reweighted employment in January 4 to 31, 2020, as regression weights. Column (2) replicates Column (1), restricting to firms with between 300 and 699 employees. Data sources: Paychex, Earnin.

Appendix Figure I
Correction of Structural Breaks in Spending, Small Business Revenue, and Employment

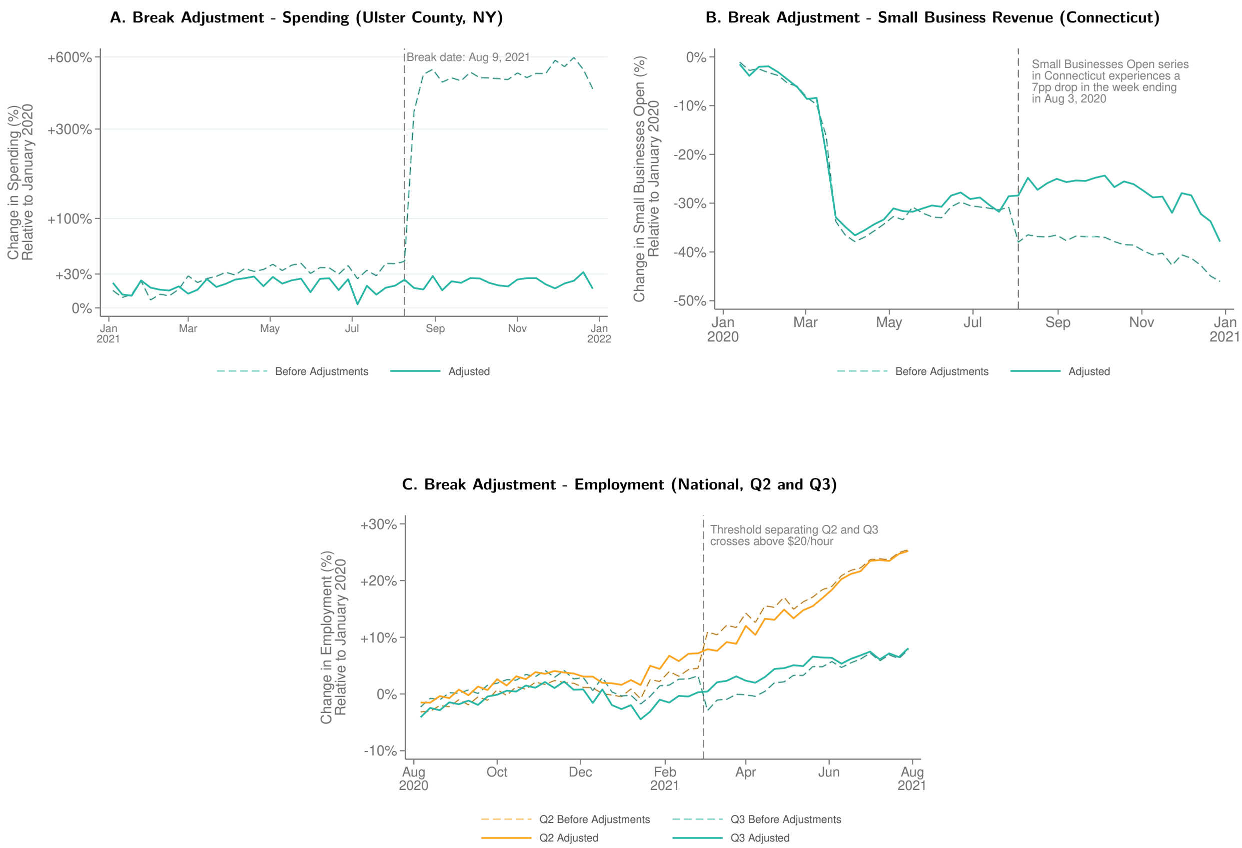

Notes: This figure presents examples of the data adjustment procedures described in Section II.B for the consumer spending series (Panel A), the small business revenue series (Panel B) and the combined Paychex-Intuit employment series (Panel C). Data sources: Affinity Solutions, Womply, Paychex, Intuit.

Appendix Figure II
Industry Shares of Consumer Spending and Business Revenues Across Datasets

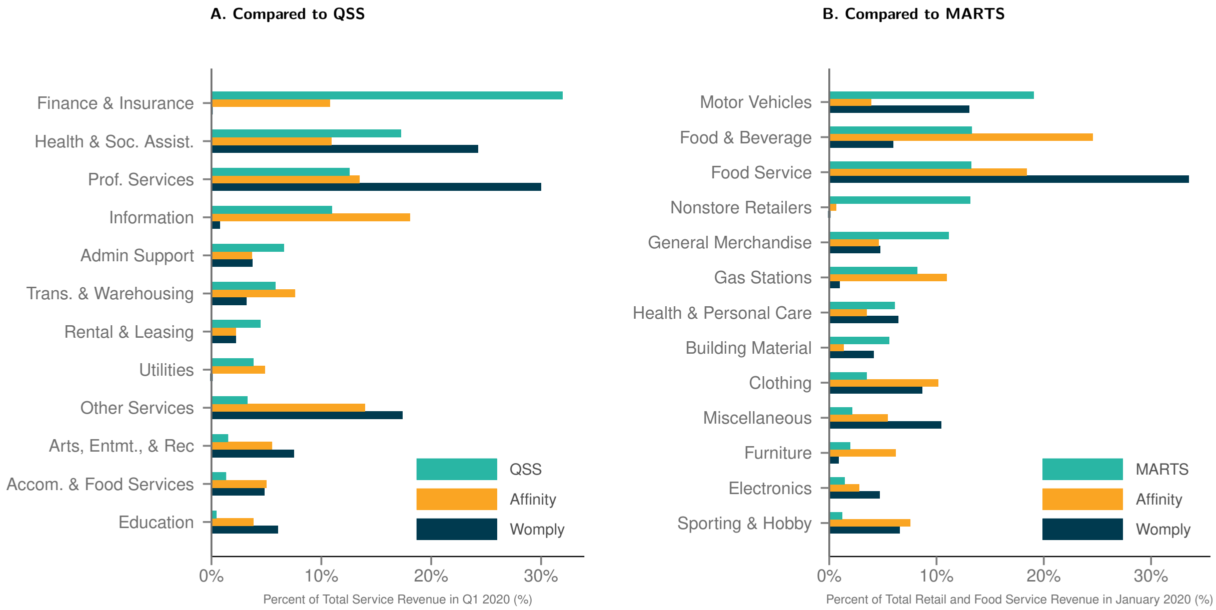

Notes: This figure compares the industry composition of spending in private sector datasets to the industry composition of spending in representative survey datasets. Panel A shows the NAICS 2-digit industry mix for transactions in the Affinity Solutions and Womply datasets compared with the Quarterly 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 Affinity Solutions and Womply datasets (respectively) by aggregating card transactions in Q1 2020, using the merchant to classify the purchase by sector. Panel B shows the NAICS 3-digit industry mix for the same two sector private datasets compared with the Advance Monthly Retail Trade Survey (MARTS), 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 MARTS, 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, QSS, MARTS.

Appendix Figure III
Industry Shares of Job Postings in Lightcast and JOLTS

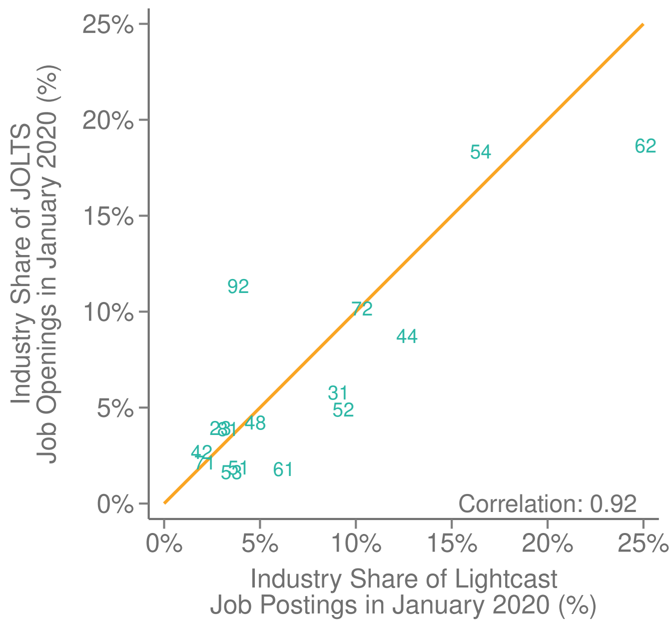

Notes: This figure presents a scatter plot showing the industry share of each 2-digit NAICS code of job postings in the Job Openings and Labor Turnover Survey (JOLTS) data in January 2020 vs. the corresponding industry share in job postings in Lightcast data in January 2020. The solid line is a 45 degree line. The annotation in the bottom right corner of the panel displays the correlation between 2-digit NAICS industry shares in the JOLTS vs. Lightcast data in January 2020, excluding NAICS 92 (Public Administration), and weighting according to total job openings in each 2-digit NAICS code in JOLTS in January 2020. Data sources: Lightcast, JOLTS.

Appendix Figure IV
Seasonal Fluctuations in Consumer Spending vs. Employment

A. Seasonal Fluctuations in Consumer Spending in MARTS Data

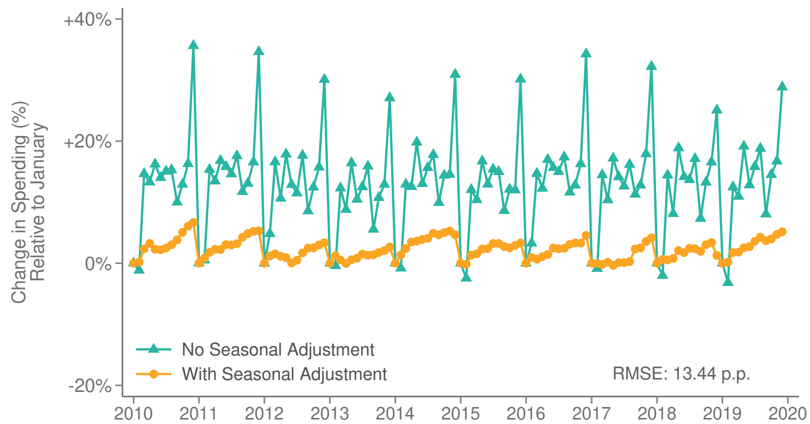

B. Seasonal Fluctuations in Employment in CES Data

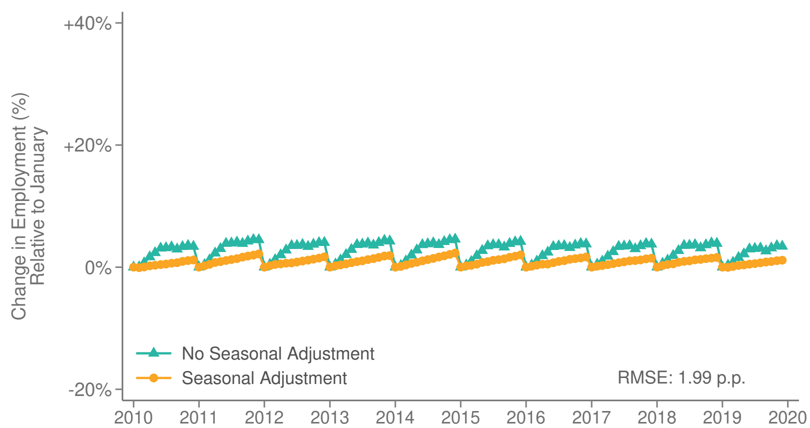

Notes: This figure compares seasonal fluctuations in Advance Monthly Retail Trade Survey (MARTS) data on consumer spending on retail sales and food services (excluding motor vehicle and gas spending) vs. Current Employment Statistics (CES) data on private sector non-farm employment. Panel A shows seasonal fluctuations in consumer spending in MARTS data. The series marked in triangles shows trends in consumer spending without seasonal adjustment, expressed as percentage changes in consumer spending in each month relative to January of the same year. The series marked in circles shows trends in consumer spending, as seasonally adjusted by the U.S. Census Bureau, expressed as percentage changes in consumer spending in each month relative to January of the same year. The annotation in the lower right hand corner displays the RMSE for the difference between the two series. Panel B replicates Panel A using CES data on private sector, non-farm employment. Data sources: MARTS, CES.

Appendix Figure V
Consumer Spending in National Accounts vs. Credit and Debit Card Data

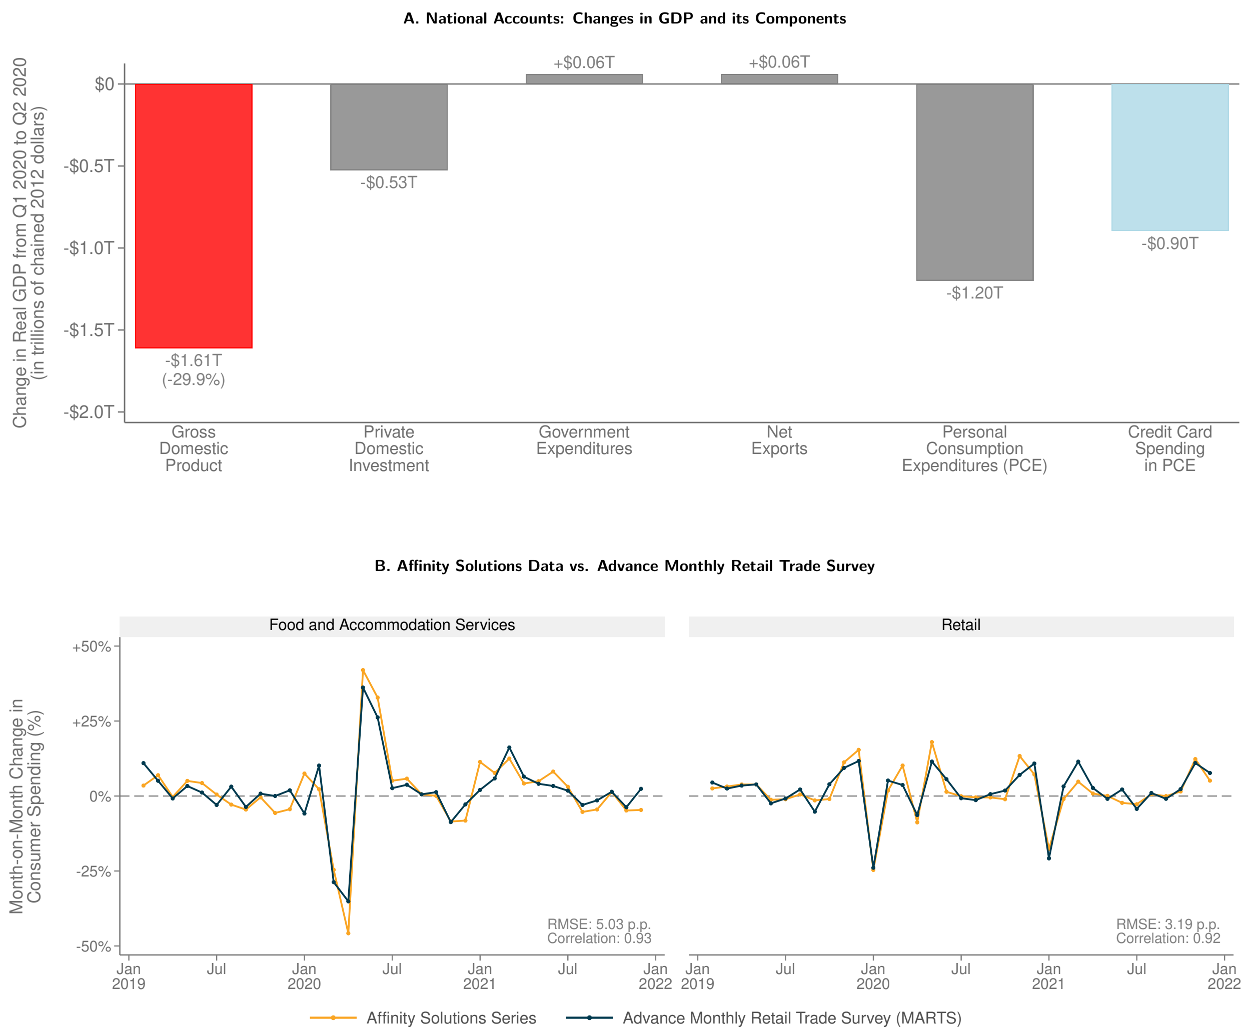

Notes: This figure examines changes in consumer spending measured in Affinity Solutions credit and debit card data, National Income and Product Accounts (NIPA) data, and Advance Monthly Retail Trade Survey (MARTS) data. Panel A shows the change in GDP from Q1 to Q2 2020 using NIPA data (Tables 1.1.1, 1.1.6 and 2.3.2). The first bar shows the seasonally-adjusted decline in real GDP ($1.61T). In parentheses under the first bar we report the compound annual growth rate corresponding to this change in real GDP (-29.9%). Bars two through five decompose the change in real GDP, estimated using NIPA Table 1.1.1. The final bar shows the contribution of components of Personal Consumption Expenditures (PCE) that are likely to be captured in credit card spending ($0.90T), estimated using NIPA Table 2.3.2. 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 indicates total spending (including spending in other modes of payment such as cash) in categories of goods and services which are likely to be well represented in card spending data, rather than total card spending itself. Panel B reports month-on-month changes in average daily spending for each month in the Affinity Solutions credit and debit card data and the Advance Monthly Retail Trade Survey (MARTS). The Food and Accommodation Services series in Panel B restricts to NAICS 72; the Retail series restricts to NAICS 44-45. The MARTS series in Panel B are constructed by dividing the total spending in each category by the number of days in that month, and then indexing to the average daily spending of January 2019. The Affinity series are constructed by taking the monthly average of the seven-day moving average series, indexed to January 2019. We also report the root mean squared error (RMSE) corresponding to the difference between these two series. Data sources: Affinity Solutions, NIPA.

Appendix Figure VI
Consumer Spending Benchmarks, Affinity vs. MARTS and CoinOut

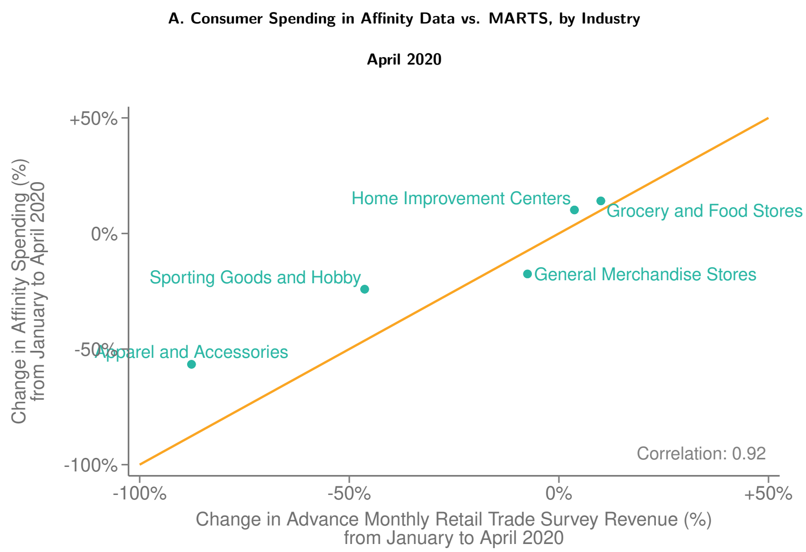

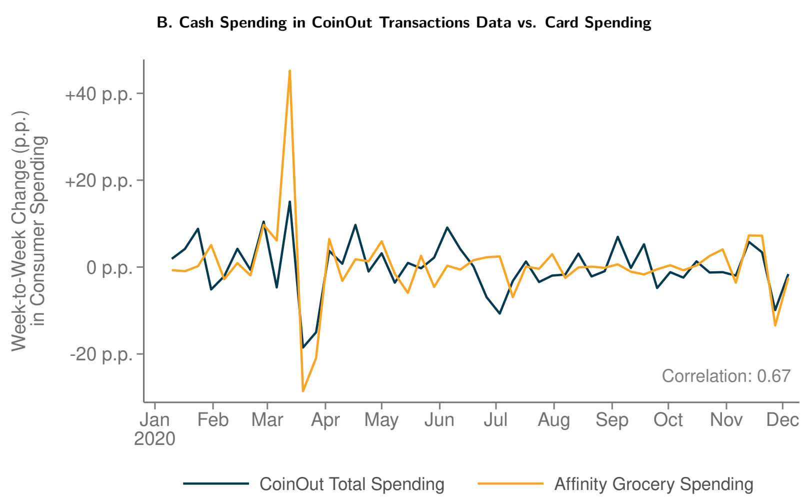

Notes: This figure benchmarks Affinity Solutions data against MARTS and CoinOut data. Panel A displays a scatter plot of changes in spending at the three-digit NAICS code level between January and April 2020 in the Affinity data vs. the MARTS data, restricting to industries where the industry definitions in the Affinity Solutions data align closely with a three-digit NAICS code surveyed in the MARTS. We report the correlation between changes in the Affinity and MARTS data, weighted by total MARTS spending in January 2020. Panel B compares week-to-week changes of national trends in cash transactions in CoinOut data vs. card spending on groceries in Affinity Solutions data in 2020. See Appendix B.3 for a description of the CoinOut data. Data sources: CoinOut, Affinity Solutions, MARTS.

Appendix Figure VII
Association Between COVID-19 Incidence and Mobility

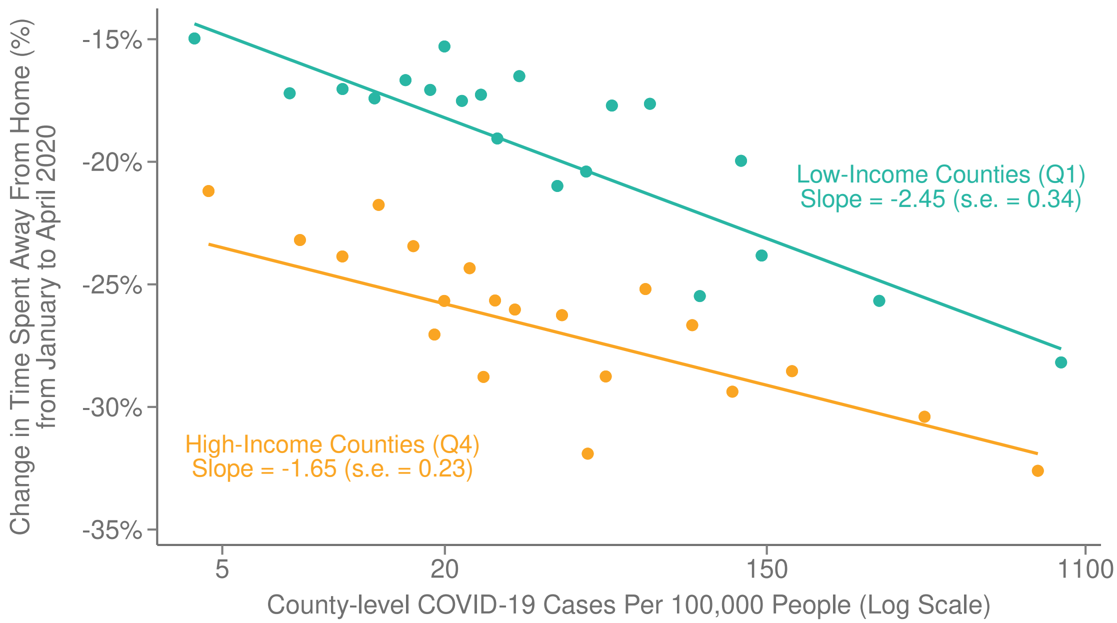

Notes: This figure presents a county-level binned scatter plot, constructed as described in Figure II. The y-axis presents the change in time spent away from home from the base period (January 3 to February 6, 2020) to the three-week period of March 25 to April 14, 2020 (see Appendix H for details on the time-away-from-home series from Google Community Mobility Reports). The x-axis variable is the logarithm of the county's cumulative COVID case rate per capita as of April 14, 2020; with axis labels showing the levels on a logarithmic scale. We plot values separately for counties in the top and bottom quartiles of median household income (measured using population-weighted 2014-2018 ACS data). Data sources: Google Community Mobility Reports, New York Times.

Appendix Figure VIII

Small Business Revenue Changes vs. Consumer Spending Changes

A. Food Services and Accommodation

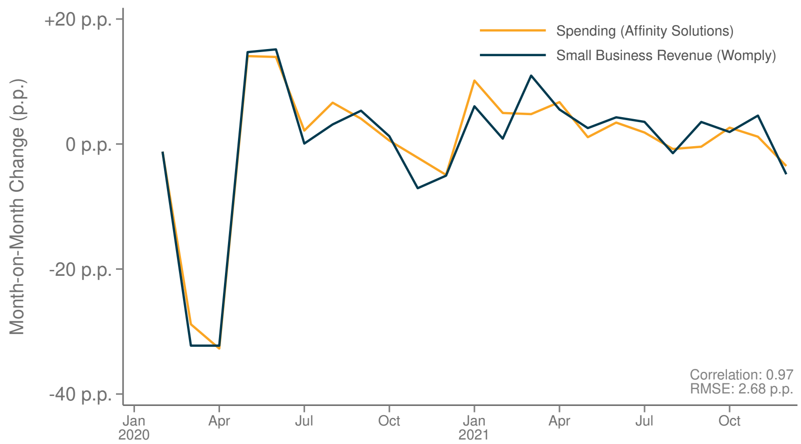

B. Retail Services

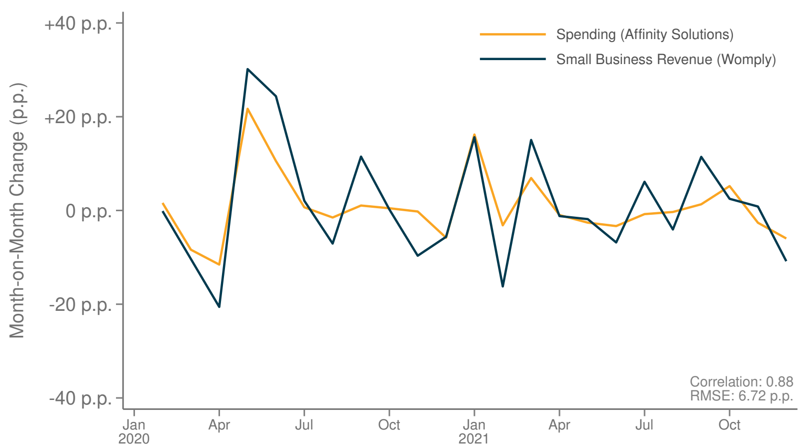

Notes: This figure compares month-on-month changes in total consumer spending (from Affinity Solutions data) and small business revenue (from Womply data) between January 2020 and December 2021. The spending series is expressed as a percentage change relative to January 6 to February 2, 2020 and the small business revenue series is expressed as a percentage change relative to January 4 to 31, 2020. We do not seasonally adjust spending or small business revenue in this figure because seasonal fluctuations provide useful variation to assess whether the consumer spending series tracks the small business revenue series. Panel A restricts to food services and accommodation (NAICS code 72), and Panel B restricts to retail trade sectors (NAICS code 44-45). The bottom right corner of each panel reports the root mean squared error (RMSE) corresponding to the difference between the two lines and the correlation between the month-on-month changes series. Data sources: Affinity Solutions, Womply.

Appendix Figure IX
National Trends in Small Business Revenue and Job Postings for Low-Education Workers

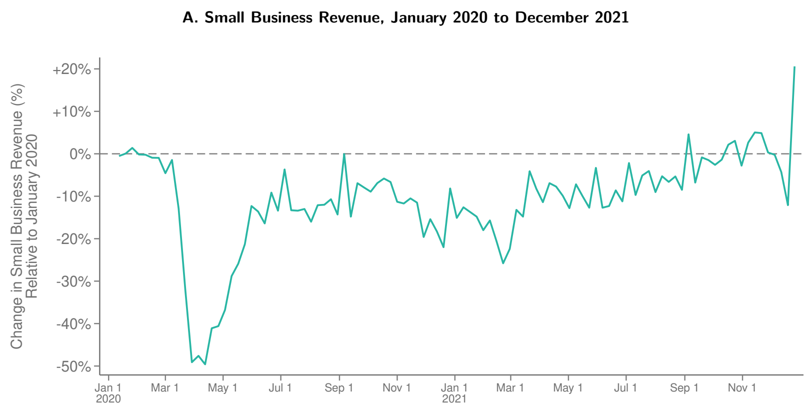

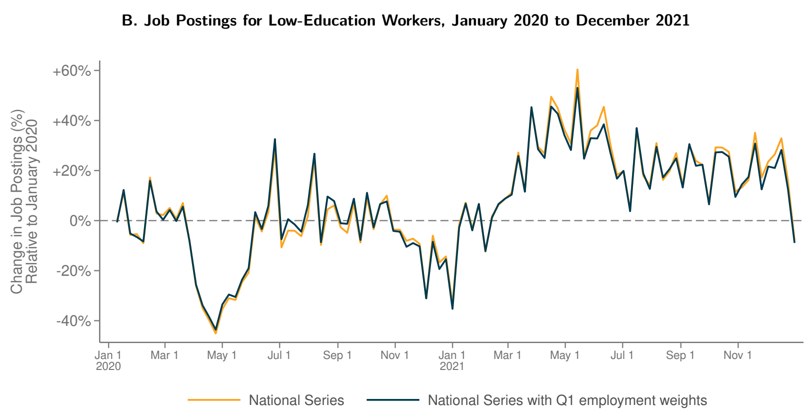

Notes: Panel A shows the changes in small business revenue between January 2020 and December 2021, relative to January 4 to 31, 2020. Panel B shows the changes in job postings for low-education workers between January 2020 and December 2021, relative to January 4 to 31, 2020. To address potential spatial mismatch in job postings, in the orange series of panel B, we reweight county-level job postings by the number of bottom-wage-quartile workers in January 2020 (as derived from the Paychex-Intuit data) and aggregate to the national level. Data sources: Womply, Lightcast, Paychex, Intuit.

Appendix Figure X
Changes in Small Business Revenues by ZIP Code

A. San Francisco

March 23 to April 12, 2020

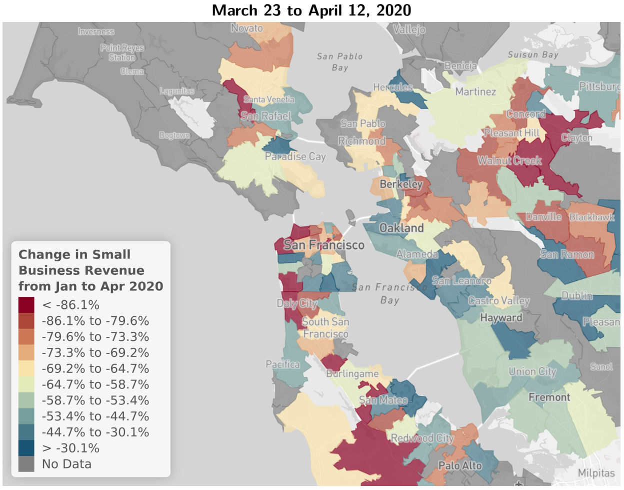

B. Chicago

March 23 to April 12, 2020

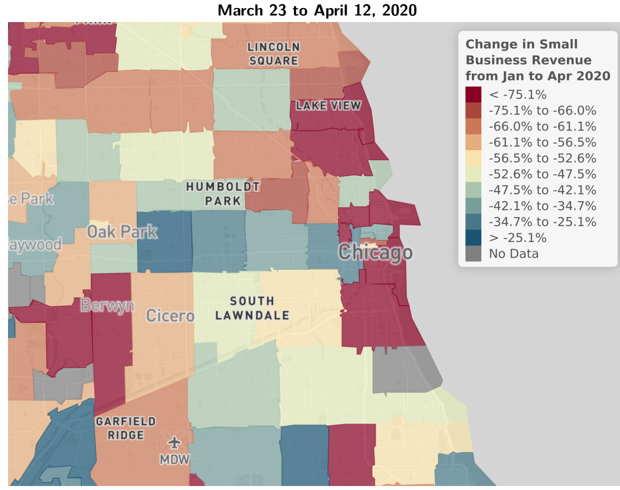

C. New York City

March 23 to April 12, 2020

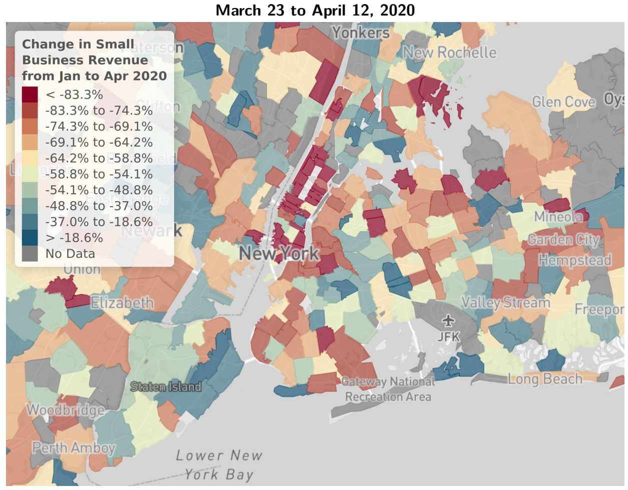

Notes: This figure plots seasonally-adjusted changes in small business revenue by ZIP code in the MSAs corresponding to San Francisco-Oakland-Hayward, CA (Panel A), Chicago-Naperville-Elgin, IL-IN-WI (Panel B), and New York-Newark-Jersey City, NY-NJ-PA (Panel C). The changes are measured during March 23 to April 12, 2020 relative to January 4 to 31, 2020. We seasonally-adjust revenue in each week by dividing the indexed value relative to January for that week in 2020 by the corresponding indexed value from 2019. These maps must be viewed in color to be interpretable; dark red colors represent areas with larger revenue declines, while dark blue colors represent areas with smaller declines. Data source: Womply.

Appendix Figure XI
Changes in Small Business Revenues vs. ZIP Code Characteristics

A. Median Income, by ZIP

April 2020

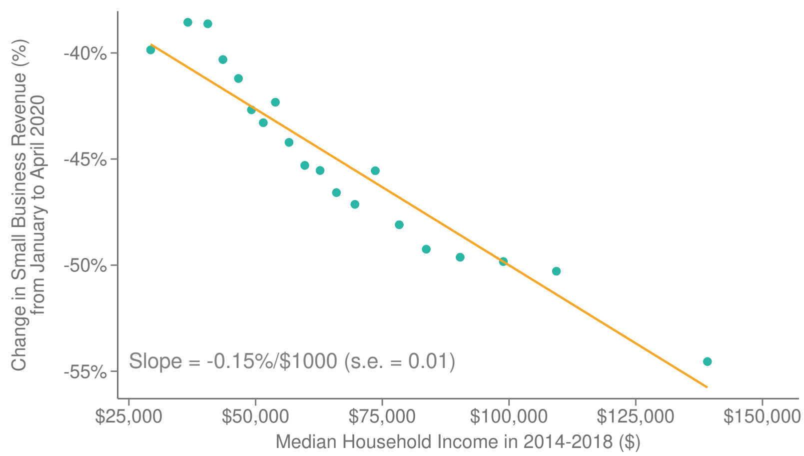

B. Population Density, by ZIP

April 2020

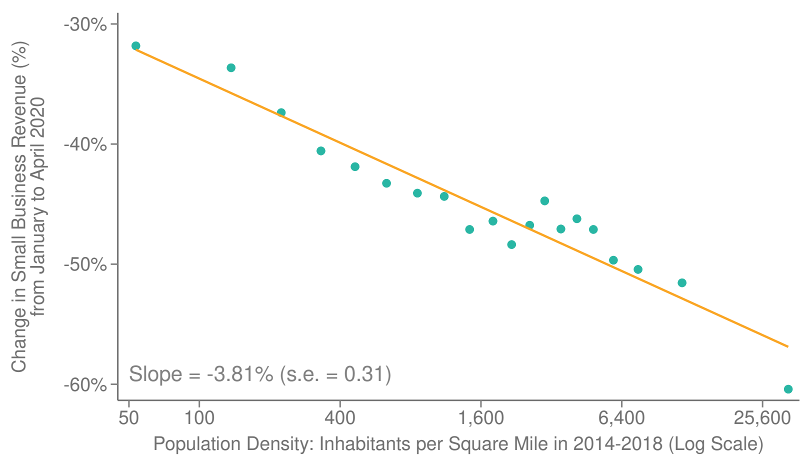

Notes: This figure presents binned scatter plots showing the relationship between changes in seasonally-adjusted small business revenue in Womply data vs. various local area characteristics at the ZIP code level. The binned scatter plots are constructed as described in Figure II. In each panel, we measure changes in small business revenue as the average value of our index at the ZIP code level between March 23 and April 12, 2020 (see Section II.B.2 and Appendix C for details on the construction of our small business revenue series). In Panel A, the x-axis variable is median household income at the ZIP code level from the 2014-2018 ACS. In Panel B, the x-axis variable is the logarithm of the number of ZIP code inhabitants per square mile in the 2014-18 ACS; with axis labels showing the levels on a logarithmic scale. Data sources: Womply, ACS.

Appendix Figure XII
Changes in Small Business Revenues vs. Local Characteristics

A. Median Two Bedroom Rent, by ZIP; Controlling for Sector FEs

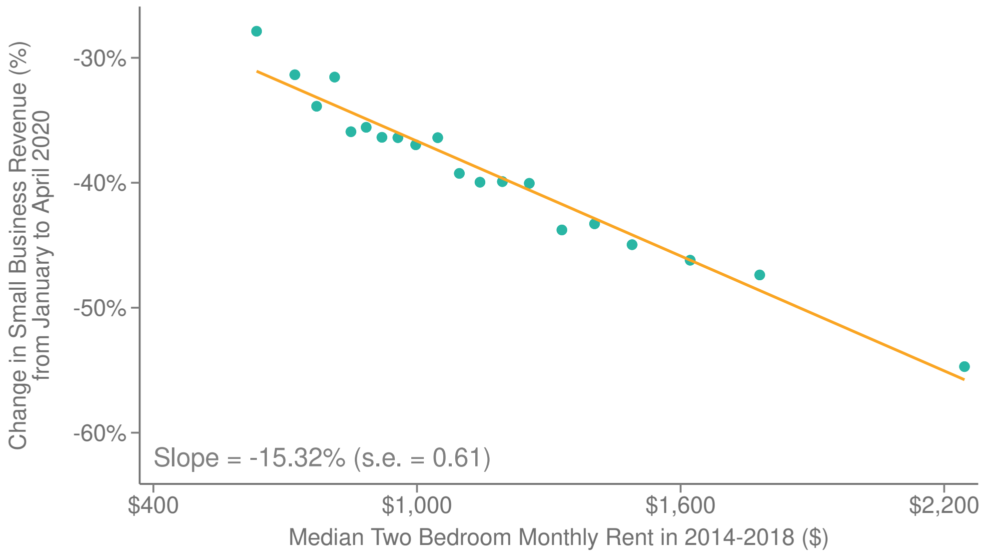

B. Income Share of the Top 1% of the Local Income Distribution, by County

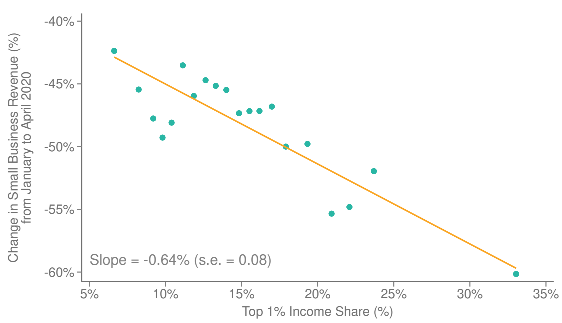

C. Share of the Population Below the Poverty Line, by County

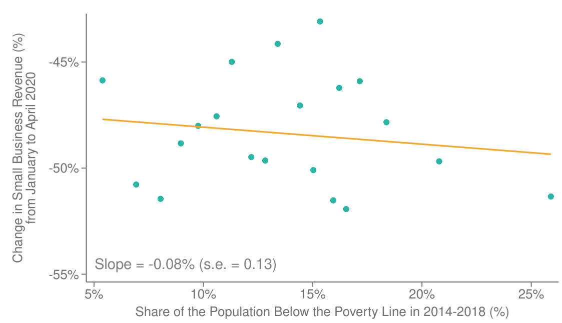

Notes: This figure shows the association between ZIP-level or county-level characteristics and changes in small business revenues between January 4 to 31, 2020 and March 23 to April 12, 2020, as measured in Womply data. Panel A presents a binned scatter plot of changes in small business revenue at the ZIP x sector (2-digit NAICS) level vs. median two-bedroom rent at the ZIP level, controlling for sector fixed effects. Panels B and C replicate Figure III but compare the declines in small business revenue with various measures of the distribution of income at the county level. Panel B presents a binned scatter plot of changes in small business revenue vs. the income share of the top 1% of the income distribution within each county, constructed using the distribution of parent incomes in Chetty et al. (2014). The top 1% of the income distribution is defined using the distribution of incomes within each county, rather than the national income distribution. Panel C presents a binned scatter plot of changes in small business revenue vs. the share of the county population with incomes below the poverty line in the 2014-2018 ACS. The binned scatter plots are constructed as described in Figure II. Data source: Womply.

Appendix Figure XIII
Changes in Employment Rates Over Time

A. Pooling All Industries

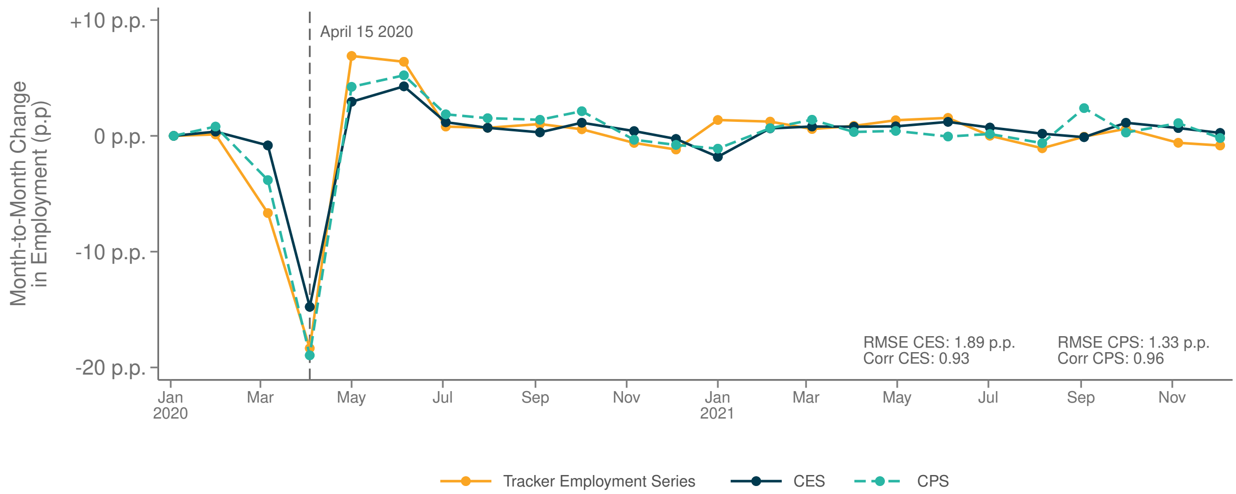

B. Accommodation and Food Services vs. Professional Services

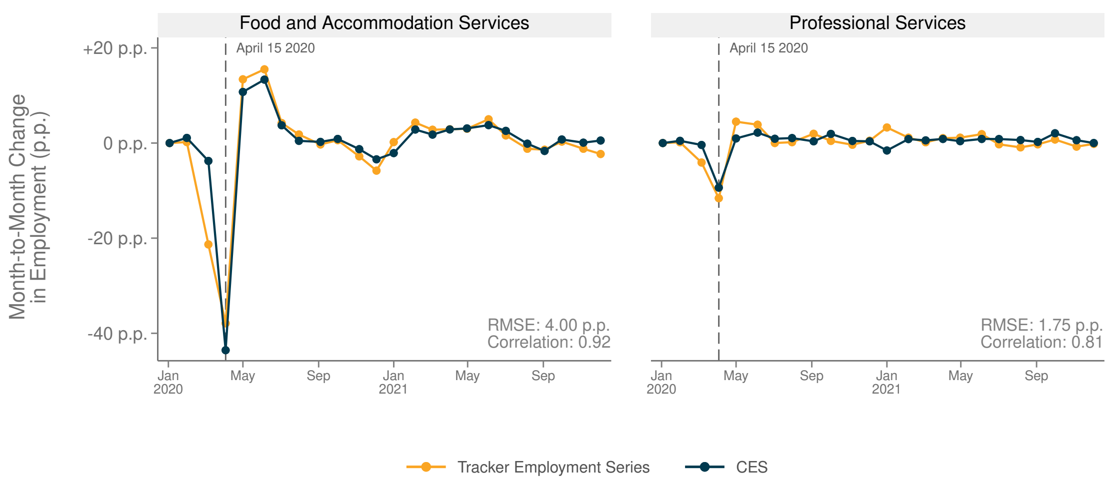

Notes: This figure compares month-to-month changes in employment, as measured in three different datasets. The series for the combined Paychex-Intuit dataset plots values of the changes measured as of the week of the 15th of a given month relative to the same value on the week of the 15th of the prior month. The series for the Current Employment Statistics (CES) and Current Population Survey (CPS) plot the change to the current monthly value relative to the prior monthly value. Panel A presents data from these three series pooling employment in all private non-farm sectors. Panel B repeats Panel A restricting to employment in the Food and Accommodation sector (NAICS 72) and the Professional and Business Services sector (NAICS 61-62). Data sources: Paychex, Intuit, CES, CPS.

Appendix Figure XIV

Changes in Employment as of July 2020 in Paychex-Intuit Data vs. CES and QCEW

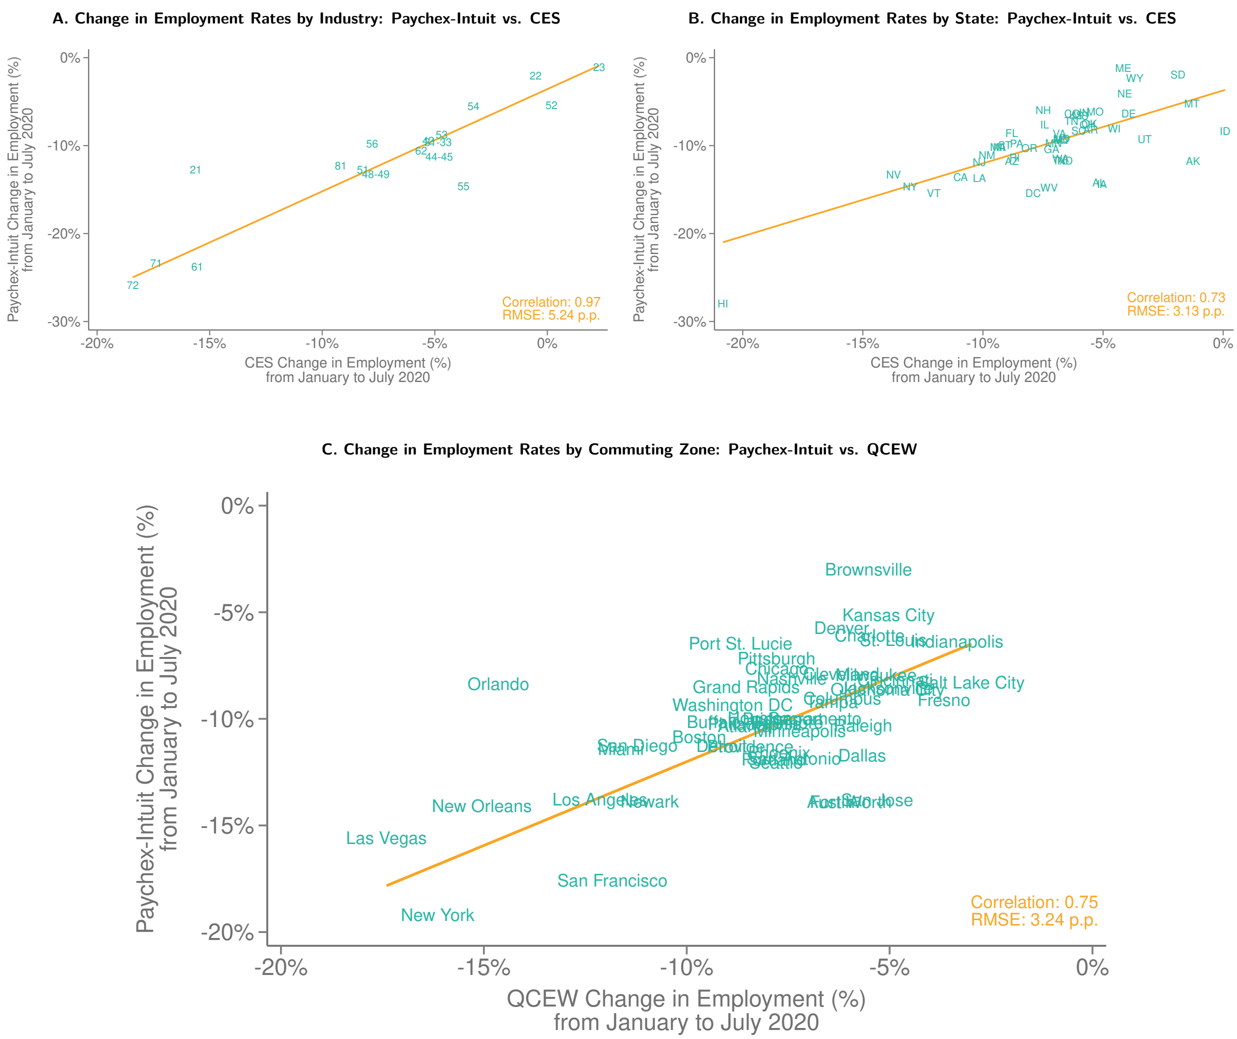

Notes: This figure benchmarks the Paychex-Intuit combined employment series to the Current Employment Statistics (CES) and the Quarterly Census of Employment and Wages (QCEW). Panel A shows a scatter plot of changes in employment in Paychex-Intuit combined data between January 4 to 31, 2020 and July 17, 2020 vs. changes in CES employment between January and July by industry (2-digit NAICS code), for private non-farm industries. Panel B shows a scatter plot of changes in employment in Paychex-Intuit combined data between January 4 to 31, 2020 and July 17, 2020 vs. changes in CES employment between January and July, by state. Panel C shows a scatter plot of changes in employment in Paychex-Intuit combined data between January 4 to 31, 2020 and July 17, 2020 vs. changes in QCEW employment between January and July, by commuting zone for the 50 largest commuting zones by population. In all panels, the bottom right corner displays the correlation between the data points in each graph, weighted respectively by CES employment in each NAICS code (Panel A), state population (Panel B), and commuting zone population (Panel C). Data sources: Paychex, Intuit, CES, QCEW.

Appendix Figure XV
Employment in Paychex-Intuit Data vs. ADP and CPS

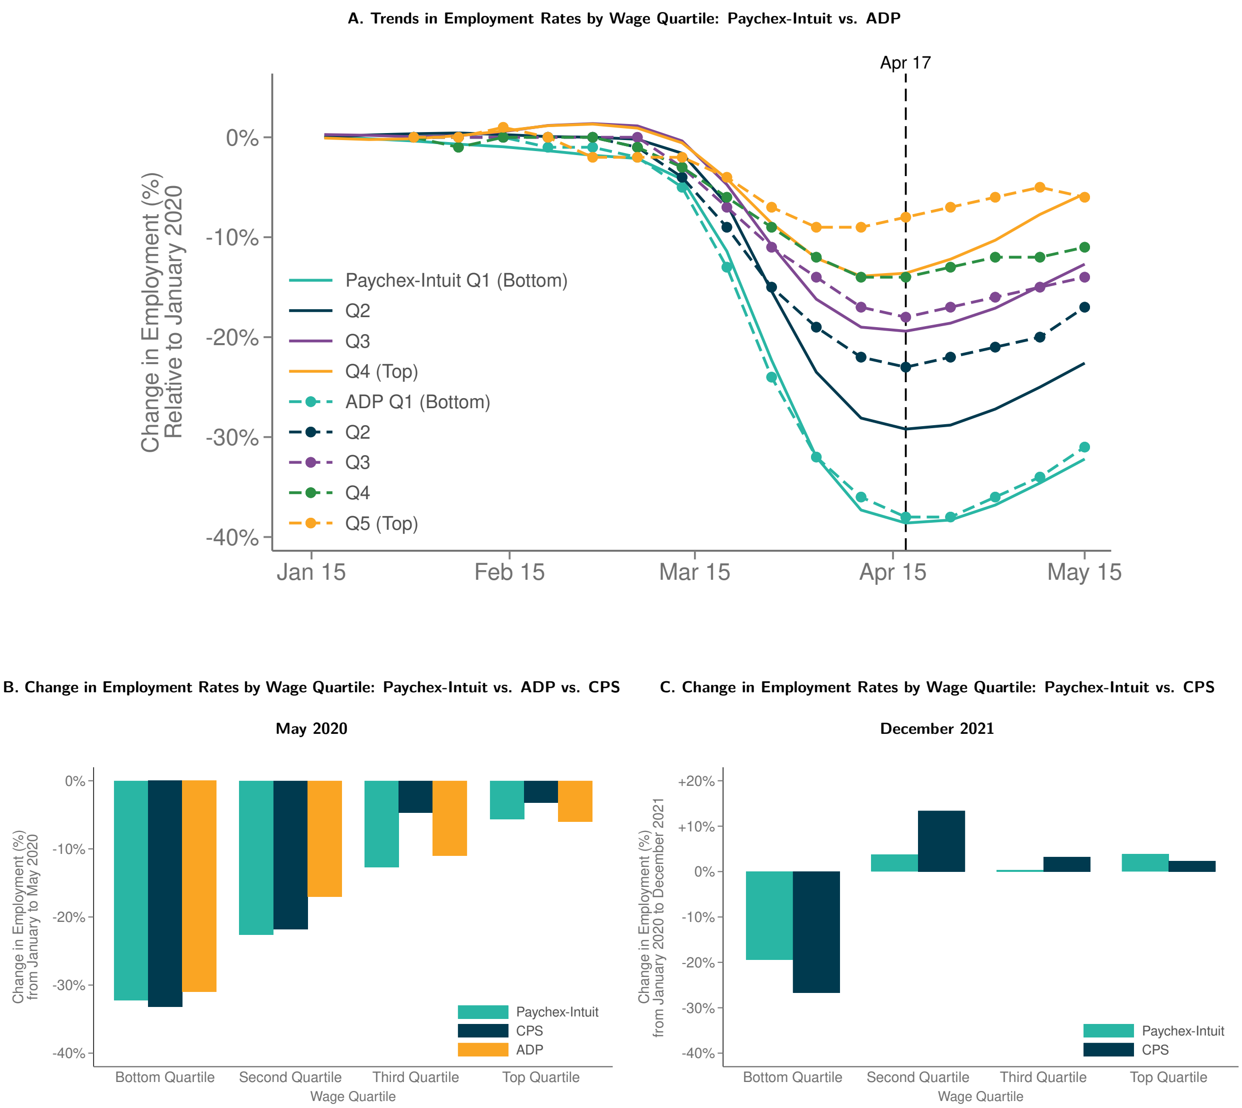

Notes: This figure benchmarks the Paychex-Intuit combined employment series to the Current Population Survey (CPS) and estimates based on ADP data in Cajner et al. (2020). Panel A shows employment trends in the Paychex-Intuit combined data (solid series) and ADP data (dashed series), split by wage quartile (combined Paychex-Intuit data) or wage quintile (ADP data). The Paychex-Intuit series is expressed as a percentage change relative to January 4 to 31, 2020. The ADP series (from Cajner et al. 2020) is expressed as a percentage change relative to February 1, 2020. Panel B shows changes in employment in the Paychex-Intuit, ADP, and CPS datasets from January to May 2020, split by wage quartile. The CPS series is expressed as a percentage change relative to January 2020. The ADP series is expressed as a percentage change relative to February 1, 2020. We omit the third wage quintile of the ADP series and compare the fourth and fifth quintiles of the ADP series to the third and fourth quartiles of the Paychex-Intuit series. Panel C replicates Panel B in December 2021 for the combined Paychex-Intuit series and CPS. For Panel C, we drop CA, MA, and NY since these three states raised the minimum wage above our upper wage threshold for bottom wage quartile employment after July 2020. Data sources: Paychex, Intuit, CPS, ADP.

Appendix Figure XVI
Changes in Low-Wage Employment by ZIP Code

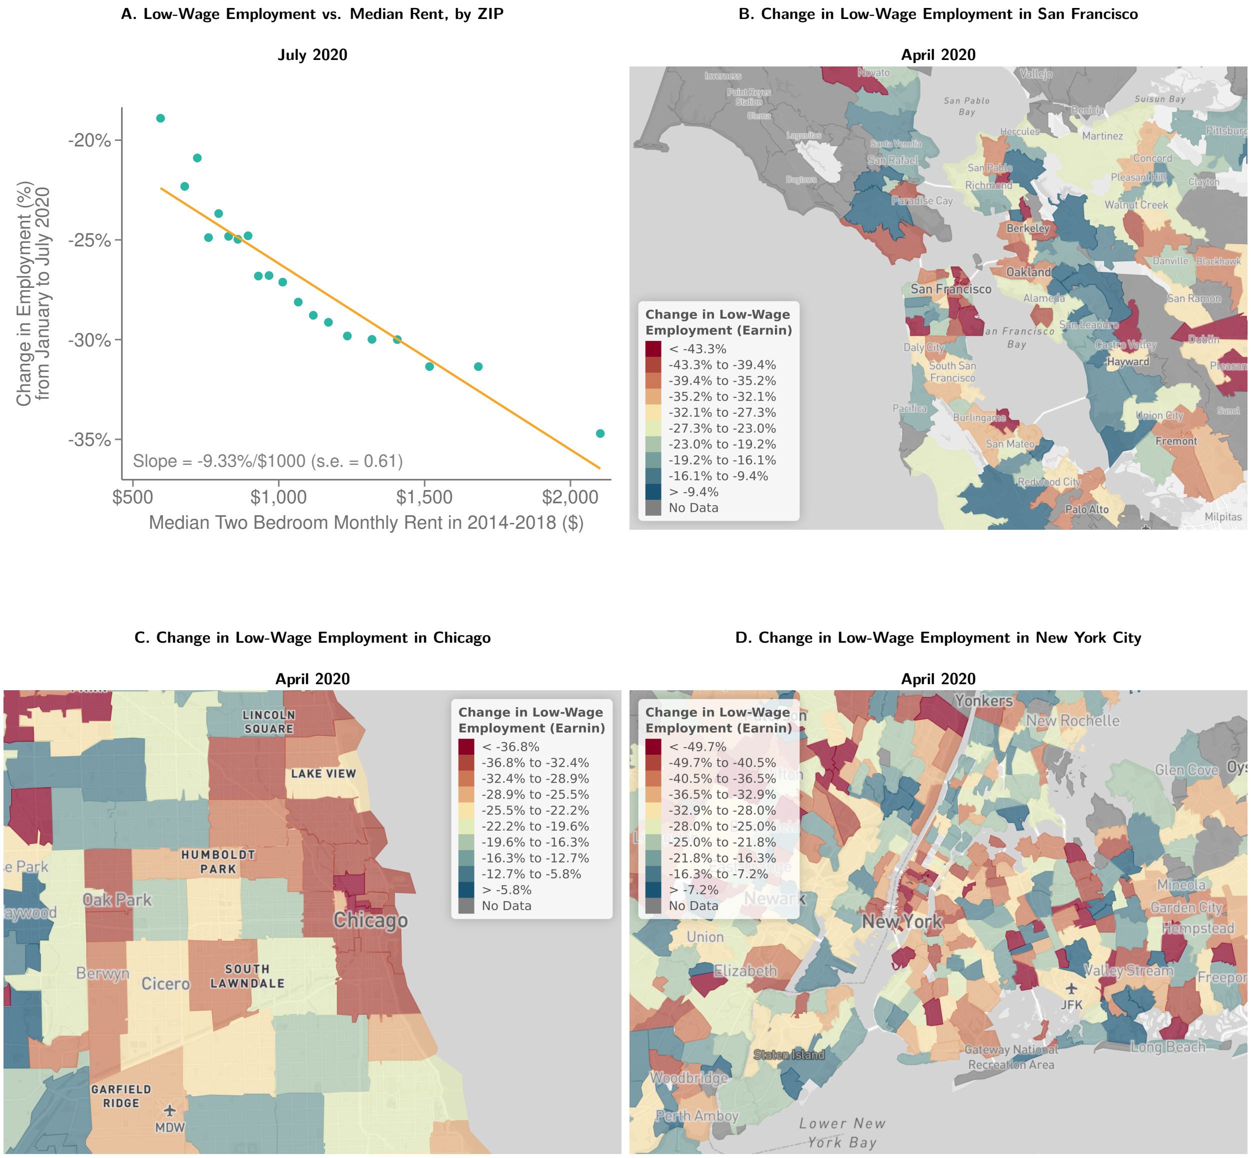

Notes: Panel A shows a binned scatter plot of the relationship between low-wage employment and median rent at the ZIP code level, constructed as described in Figure II. The x-axis variable is the median rent within a ZIP code for a two-bedroom apartment in the 2014-2018 ACS. The y-axis variable is the average value of low wage employment at the ZIP code level from Earnin during the month of July 2020 (see Section II.B.4 and Appendix E for more detail on the construction of the Earnin employment series). Panels B, C and D replicate Appendix Figure X using Earnin data on changes in employment among low-wage workers, plotted by employer ZIP code. We measure the change in employment as total average weekly employment between March 27 and April 24, 2020 divided by total average weekly employment between January 4 and 31, 2020. These maps must be viewed in color to be interpretable; dark red colors represent areas with larger employment declines, while dark blue colors represent areas with smaller declines. Data sources: Earnin, ACS.

Appendix Figure XVII
Geography of Employment Losses in the Great Recession vs. COVID Recession

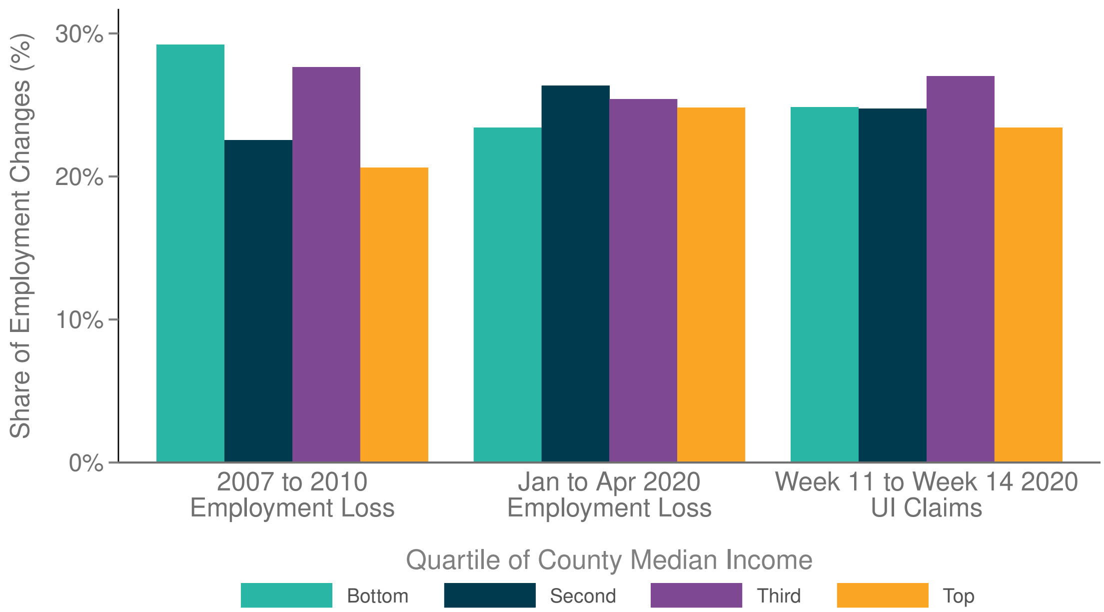

Notes: This figure displays the share of job losses occurring in low vs. high income counties during the Great Recession and the COVID Recession. We split counties into population-weighted quartiles by median household income in the 2006 ACS for the Great Recession (left bars) and the 2014-2018 ACS for the COVID Recession (middle and right bars). To construct the first set of four bars, we use BLS data to measure the share of the national employment losses from 2007 and 2010 occurring within counties in each quartile of median household income. The second set of bars replicates the first set of bars using the employment losses from January to April 2020. The third set of bars reports the share of total initial UI claims within each county income quartile between March 15, 2020 (the first week of COVID-related UI claims) and April 12, 2020. In this third set of bars, we only include counties within states that issue weekly reports of county-level UI claims data; these states include 53% of the U.S. population. The increase in unemployment rates between February and April 2020 (11%) was only two-thirds as large as the decrease in employment (16%). The difference was due to a 5% decline in the labor force: many people lost their jobs but were not actively searching for a new job in the midst of the pandemic (Coibion, Gorodnichenko, and Weber 2020). In the three prior recessions, the labor force continued to grow by 0.3% to 0.8% annually. We therefore focus on the decline in employment rates to obtain comparable statistics on job loss across recessions. Data source: BLS.

Appendix Figure XVIII
Changes in Employment by Wage Quartile and Consumer Spending, Retail Trade

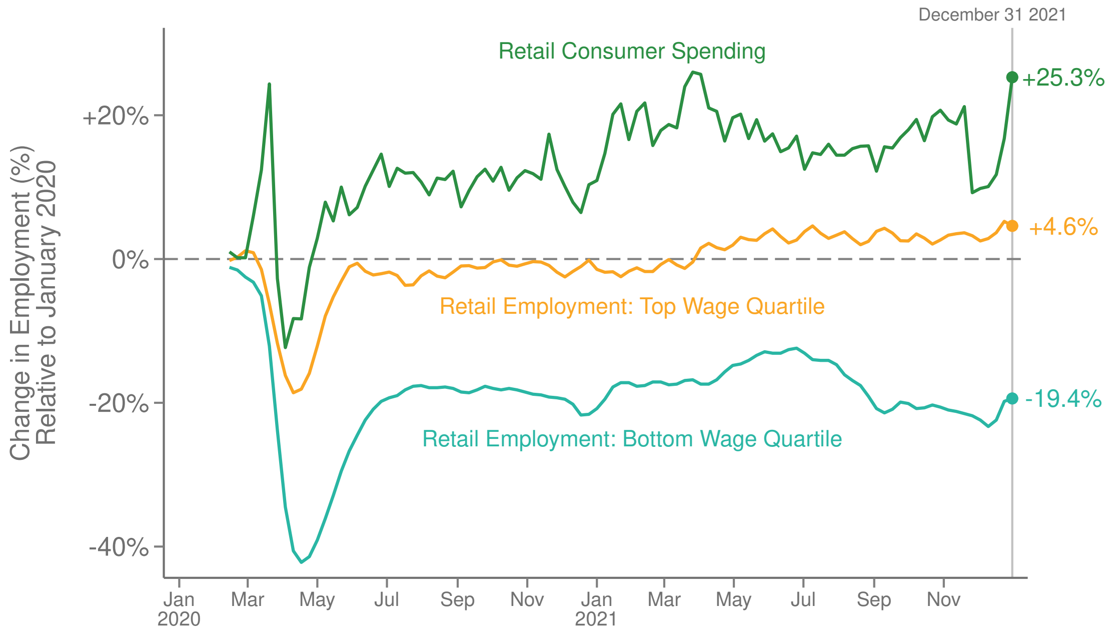

Notes: This figure plots our combined Paychex-Intuit employment series for top and bottom wage quartile jobs in the Retail Trade sector (NAICS 44-45), along with our consumer spending series for this sector (see Section II.B.1 and Appendix B for details on the construction of the consumer spending series). Data sources: Paychex, Intuit, Affinity Solutions.

Appendix Figure XIX
Relationship between Employment and Rent in CPS Panel

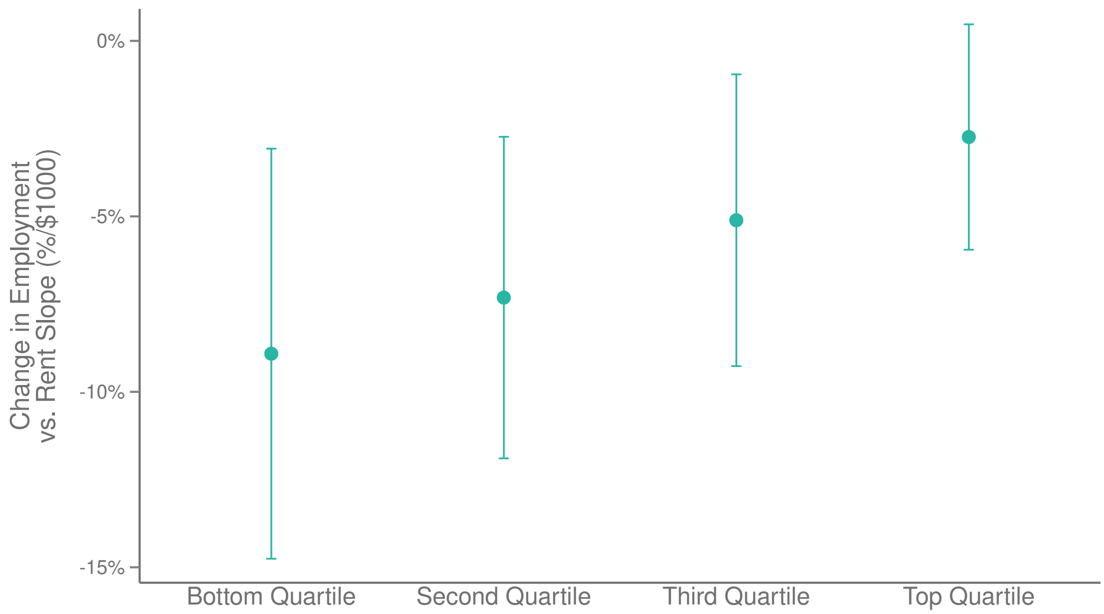

Notes: This figure plots regression estimates of the relationship between state-level changes in employment, measured separately for each wage quartile, and median two-bedroom rent in the CPS panel. We obtain the point estimates and 95% confidence intervals from population-weighted state-level OLS regressions. We use the change in employment between each respondent's initial survey (conducted 12 months before their follow-up survey) and their follow-up survey in July 2020 to February 2021. To construct the CPS panel, we first restrict the sample to the set of individuals who were employed during their initial survey and had their follow-up survey between July 2020 and February 2021. We then classify these individuals by their wage quartile during their initial survey (July 2019 to February 2020) based on inflation-adjusted values of the Federal Poverty Line, adding uniform noise between [-$0.50, $0.50] to whole number wages to smooth out spikes in the wage distribution at whole numbers. See Section II.B.4 for more details on the wage quartile thresholds and smoothing. The sample omits California, Massachusetts, and New York due to mismeasurement of bottom-quartile employment changes as a result of minimum wage increases; see Appendix E.2 for more details. Data sources: CPS.

Appendix Figure XX
Changes in Labor Supply due to Population Shifts

A. Trends in Low-Wage Employment for US-Born vs. Naturalized Citizens vs. Non-Citizens

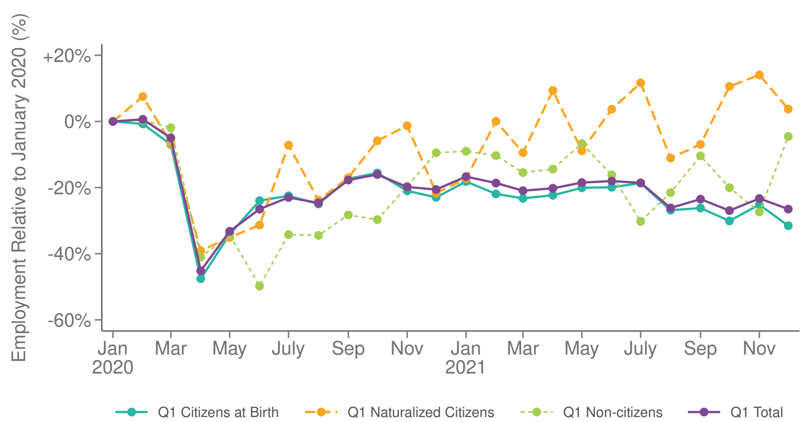

B. Net In-Migration vs. Rent in 2021

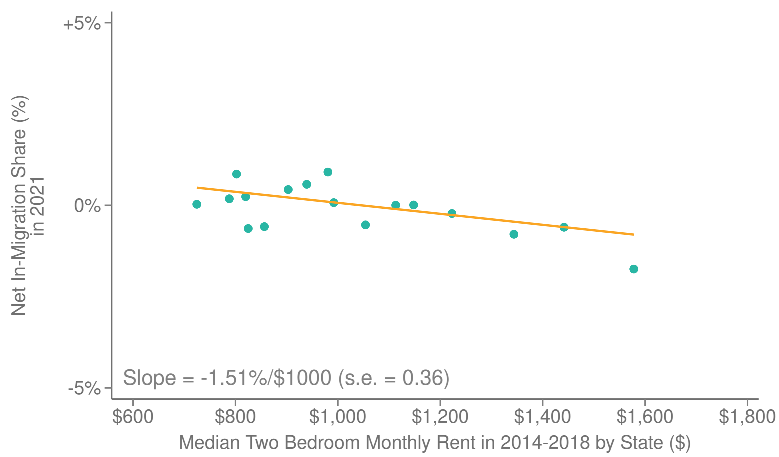

Notes: Panel A uses CPS data to show the employment change relative to January 2020 for workers in the bottom wage quartile, comparing those who are U.S. citizens at birth, naturalized U.S. citizens, or non-U.S. citizens. Panel B presents a binned scatter plot showing the association between net in-migration shares in 2021 and median rent at the state level in the ACS 2014-2018. The binned scatter plot is constructed as described in Figure II. For each state we calculate the net in-migration share as in-migration minus out-migration divided by the state population, using Current Population Survey Annual Social and Economic Supplement (CPS ASEC). Data sources: CPS, ACS.

Appendix Figure XXI
Effects of Stimulus Payments on Spending, without Adjusting for Pre-Trends

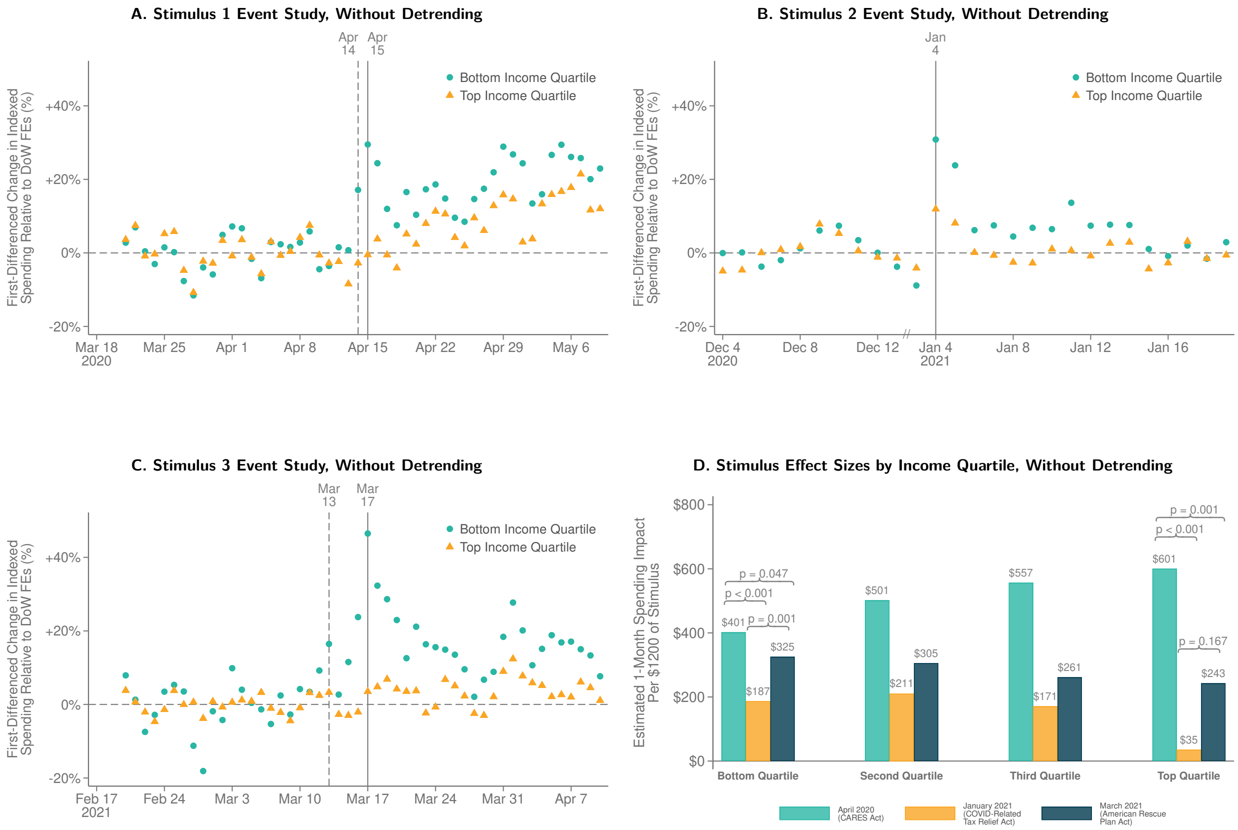

Notes: This figure shows event studies of spending around stimulus payments without adjusting for linear pre-trends. Panel A shows the equivalent of Figures VII.A and VII.B without adjusting for linear pre-trends. Panel B is identical to Figure VII.C, since both figures do not adjust for linear pre-trends. Panel C shows the equivalent of Figure VII.D without adjusting for linear pre-trends. Panel D shows the equivalent of Figure VIII without adjusting for linear pre-trends. Data source: Affinity Solutions.

Appendix Figure XXII
Robustness of Stimulus Effect Size to Post-Stimulus Event Window

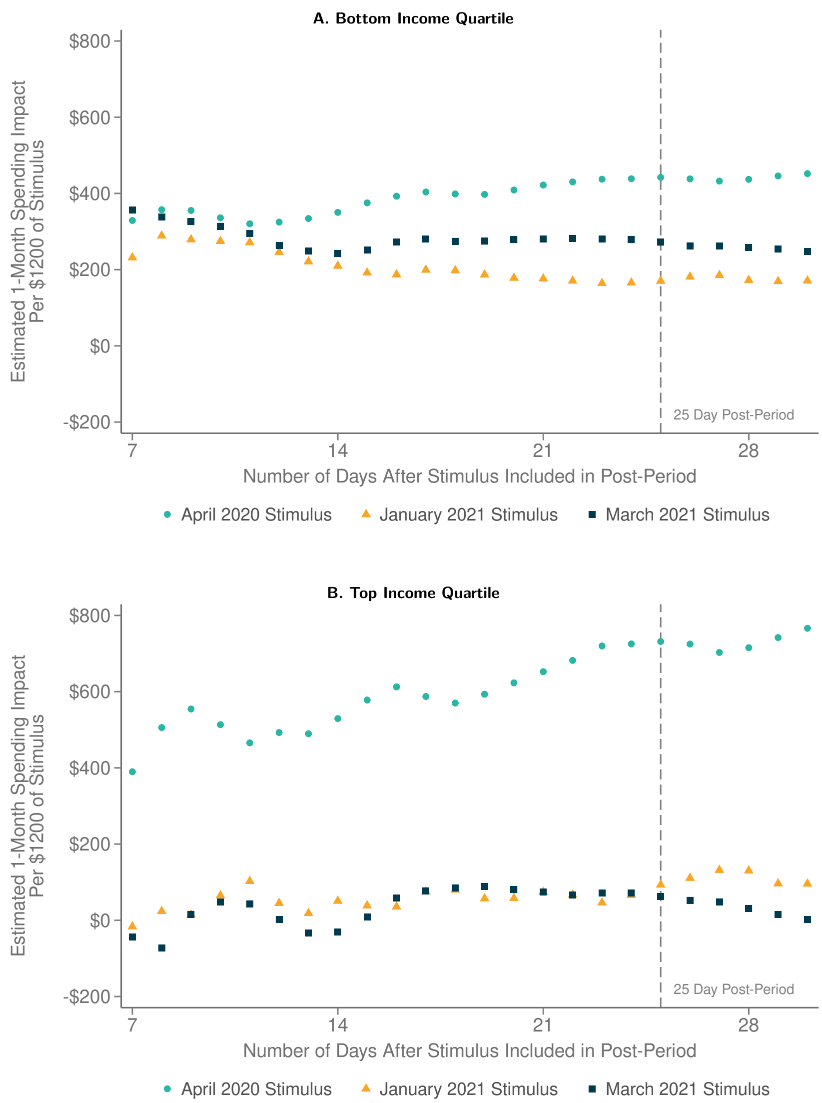

Notes: This figure displays coefficient estimates from varying post-period windows as a robustness check for the analysis displayed in Figure VII, Figure VIII, and Appendix Table VI. Each dot corresponds to a different estimate for the "Combined Dollar" effects from Appendix Table VI, Column 5. Panel A plots estimates for cardholders residing in the bottom income quartile of ZIP codes, varying the post-period from 7 to 30 days, for each of the three rounds of stimulus. Panel B repeats this analysis for cardholders residing in the top income quartile of ZIP codes. For the April 2020 and March 2021 stimulus rounds, the estimate at 25 days matches the estimate in Appendix Table VI, Column 5; for the January 2021 stimulus round, the estimate at 16 days matches the estimate in Appendix Table VI, Column 5. See Section IV.A and the notes to Figure VII and Appendix Table VI for details on how these estimates were calculated. Data source: Affinity Solutions.

Appendix Figure XXIII
Daily Variance of Consumer Spending, by Week

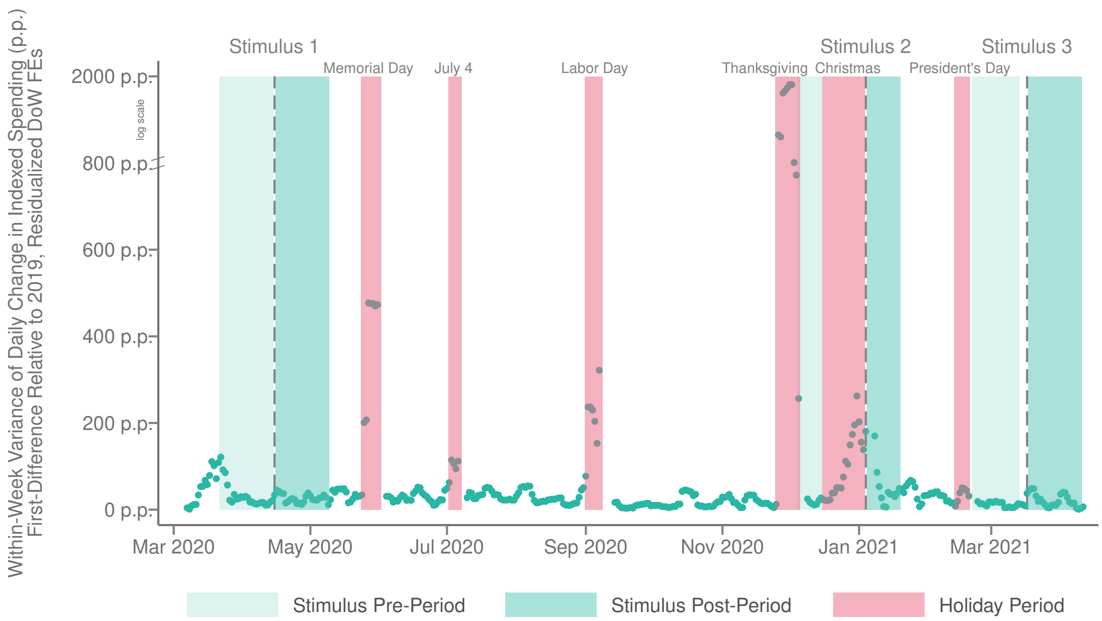

Notes: This figure depicts the within-week daily variance in the first-differenced change in indexed spending, which we construct as described in Appendix K.1. We then plot a 7-day trailing average of variance, omitting the points in the 6 days after holidays. The values in the Thanksgiving period range from 804p.p. to 1834p.p. – values between 800p.p. and 2000p.p. are plotted on a log scale rather than a linear scale to preserve visual clarity. Three vertical dashed lines mark the timing of deposits for each of the three rounds of stimulus payments, with light green shading to mark the pre-periods and dark green shading to mark the post-periods in the respective event studies in Figure VII and Appendix Table VI. The red shaded dates denote holiday periods that we drop from the event study analysis (in Figure VII.C) and from the permutation analysis (in Appendix Figures XXIV and XXV). Data source: Affinity Solutions.

Appendix Figure XXIV
Permutation Test for Differences in Stimulus Effects, Bottom Income Quartile

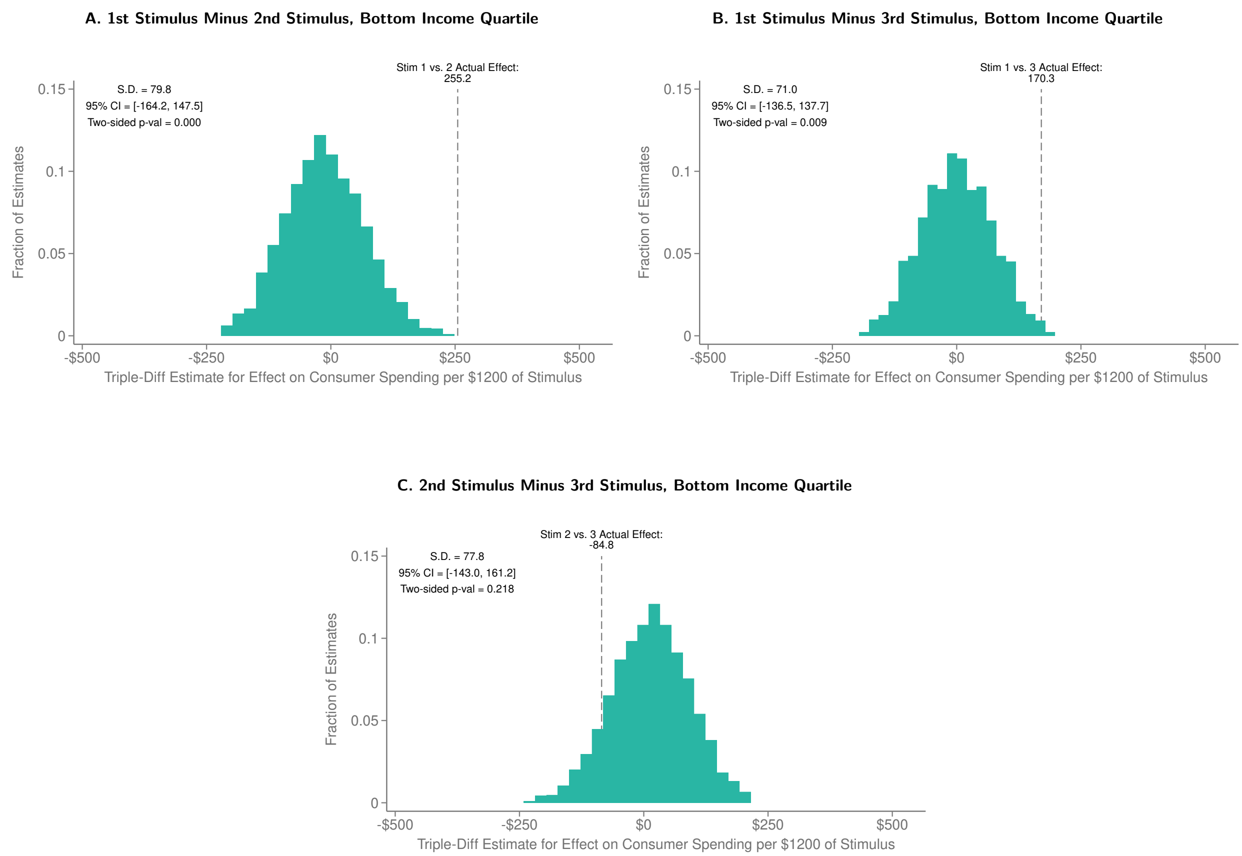

Notes: These figures present a permutation test of the null hypothesis that there is no difference in spending impacts across rounds of stimulus payments for households in the bottom income quartile of ZIP codes. We consider all placebo dates for stimulus payments from August 1, 2020 to May 1, 2021, dropping dates that fall within 25 days of the beginning or end of an actual pre- or post-stimulus window or within a high-variance holiday period (see Appendix Figure XXIII for more details). We then calculate the "combined dollar" estimate for each placebo date (as in Appendix Table VI, Column 5) for each income quartile and stimulus round, following the approach described in Appendix K.3. We include 25 days pre-event and then either 25 days post-event (1st and 3rd stimulus) or 16 days post-event (2nd stimulus). We drop days in the pre- or post-period that fall within a holiday period, and we omit the estimate entirely if a holiday falls in the post-period. We adjust for linear pre-trends in both the treatment and control series whenever there are more than 20 days in the pre-period. This leads to 89 placebo estimates for the first stimulus and third stimulus, and 59 placebo estimates for the second stimulus. We take the difference between each possible pair of placebo estimates (1st minus the 2nd stimulus; 1st minus the 3rd stimulus; and 2nd minus the 3rd stimulus) and plot the distribution of the calculated difference in "Combined Dollar" effects. Panel A plots the distribution of placebo differences between the 1st minus the 2nd stimulus. Panel B repeats this for the 1st minus the 3rd stimulus, and Panel C repeats this for the 2nd minus the 3rd stimulus. On each panel, we mark the actual difference in estimates (taken from the appropriate difference between estimates in Appendix Table VI, Column 5) with a dashed vertical line. We also report the standard deviation of the placebo draws, the 95% confidence interval (the smallest interval that covers 95% of placebo differences), and the two-sided p-value (twice the minimum of the fraction of placebo differences to the right of the actual estimate and the fraction to the left of the actual estimate). Data source: Affinity Solutions.

Appendix Figure XXV
Permutation Test for Differences in Stimulus Effects, Top Income Quartile

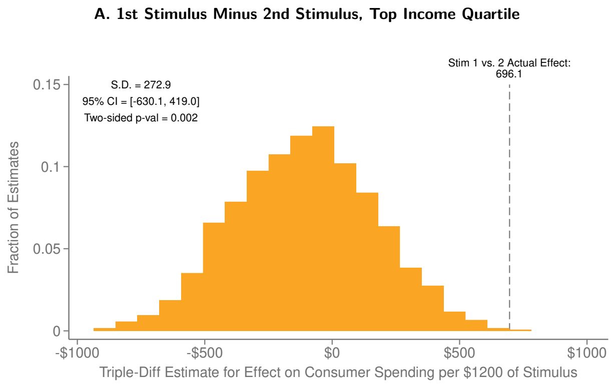

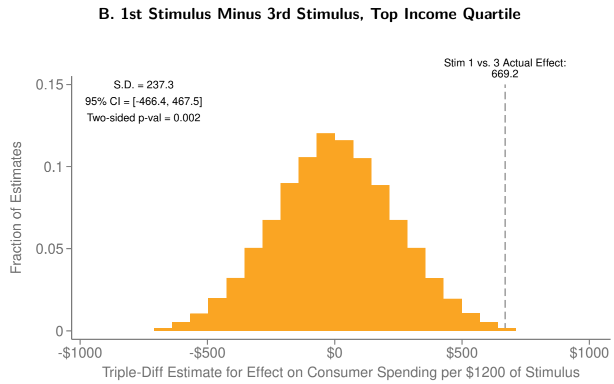

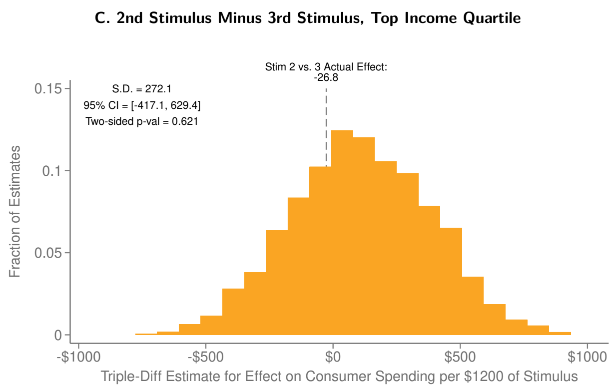

Notes: These figures present a permutation test of the null hypothesis that there is no difference in spending impacts across rounds of stimulus payments for households in the top income quartile of ZIP codes. Each panel is constructed analogously to the corresponding panel in Appendix Figure XXIV. Data source: Affinity Solutions.

Appendix Figure XXVI
Effects of Reopenings on Economic Activity

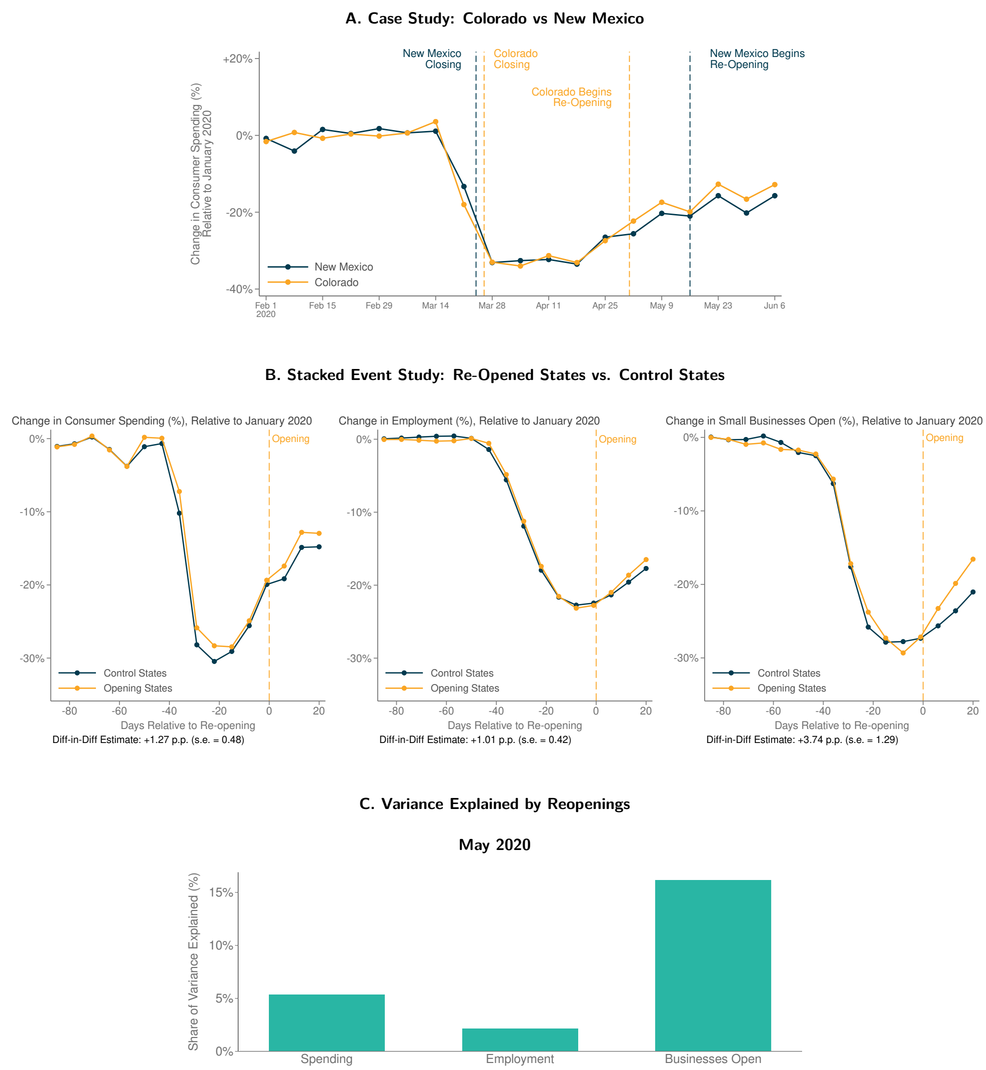

Notes: Panel A plots the change in seasonally-adjusted consumer spending relative to January 6 to February 2, 2020 for New Mexico and Colorado. Colorado partially reopened non-essential businesses on May 1, 2020, while New Mexico did not do so until May 16, 2020. Panel B1 plots an event study of the same outcome variable for five "treated" states (SC, AK, GA, MN and MS) that partially reopened non-essential businesses between April 20 and 27, 2020. For each reopening, the treated states are matched to multiple control states (listed in Appendix Table X) that did not reopen within the subsequent 3 weeks but had similar trends of the outcome variable in the preceding 3 weeks. We then stack the resulting event studies by time relative to the reopenings. Panels B2 and B3 replicate Panel B1 with, respectively, the change in employment and the seasonally-adjusted change in small businesses open. In Panels B1 to B3, we report the coefficient from a difference-in-differences regression comparing treated vs. untreated states in the two weeks after vs. the two weeks before the reopening (also reported in Appendix Table XI). Panel C reports the share of the variance in outcomes explained by reopenings as of May 18, 2020. To estimate these variance shares, we first calculate the variance of each outcome across states on May 18, 2020. Then, we add the difference-in-differences estimate for the effect of reopening on a given outcome to all states not open on May 18 (adding only half of the effect if the state opened between May 11 and 18, 2020). We then recalculate the variance in this counterfactual in which all states had reopened. The share of variance explained by reopenings for each outcome is 1 − (counterfactual variance/actual variance). Data sources: Affinity Solutions, Paychex, Intuit, Womply.

Appendix Figure XXVII
Effects of State-Ordered Business Closures on Consumer Spending

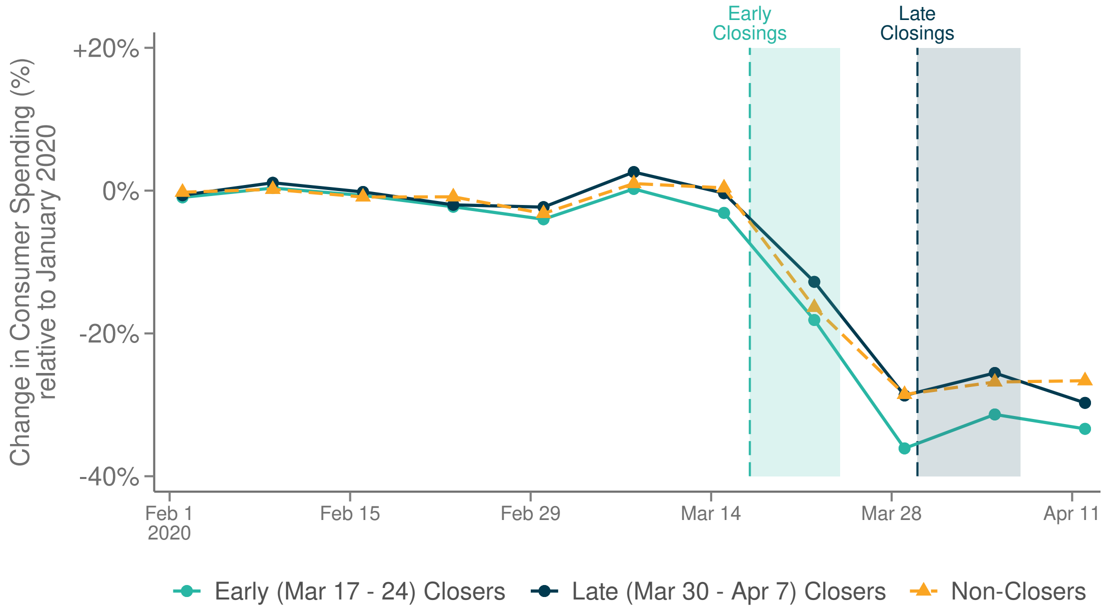

Notes: This figure displays trends in seasonally-adjusted consumer spending in the Affinity Solutions data, pooling states by the date on which a state-wide order closed non-essential businesses and weighting by state population. States are aggregated into three groups: “Early” (state-wide closure order issued between March 17 and 24, 2020), “Late” (state-wide closure order issued between March 30 and April 7, 2020), and “Non-Closers” (no state-wide closure order issued by April 7, 2020). Dashed lines denote the first date on which state-wide orders closing non-essential businesses were issued by “Early Closers” (March 17) and “Late Closers” (March 30). The blue and grey shaded areas denote the range of closure dates for “Early Closers” and “Late Closers” respectively. Data source: Affinity Solutions.

Appendix Figure XXVIII
Effects of the Paycheck Protection Program on Employment

A. Change in Employment by PPP Eligibility, All Industries Excluding Food Services

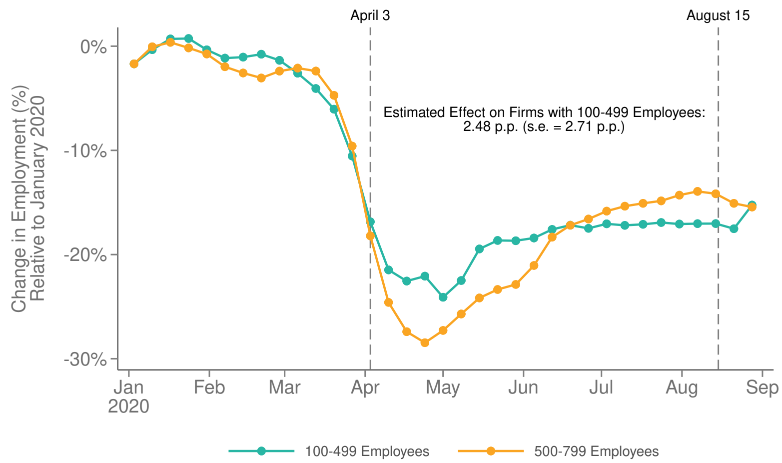

B. Change in Employment by Firm Size, All Industries Excluding Food Services

June 2020

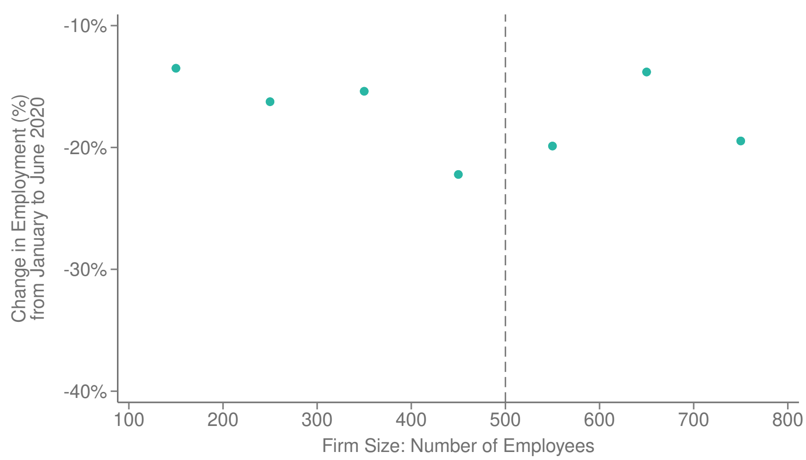

Notes: This figure analyzes the effects of the Paycheck Protection Program on employment using the threshold in eligibility at 500 employees. We pool all industries except Accommodation and Food Services (NAICS 72), which was subject to different eligibility rules (discussed in Appendix L.2). Panel A compares employment trends measured in Paychex and Earnin data among firms with 100-499 employees (generally eligible for PPP loans) to firms with 500-799 employees (generally ineligible for PPP loans). To construct these employment trends, we begin by calculating weekly employment changes relative to January 4 to 31, 2020 disaggregated by county, industry (2-digit NAICS), wage quartile and firm size bin. We reweight these cells so that employment shares by industry within each eligibility group match the overall employment shares by industry over January 4 to 31, 2020. We plot the "control" series (firms with 500-799 employees) directly as the mean weekly value of the reweighted employment series. We plot the "treated" series (firms with 100-499 employees) as the sum of the control series and the coefficients from an event study specification (controlling for county x wage quartile x week fixed effects and leaving out January 3, 2020 as the reference dummy). We use reweighted employment over January 4 to 31, 2020 as regression weights. The regression estimate in Column 1 of Appendix Table XII, which uses an estimation window of March 11 to August 15, 2020, is also reported in the figure. Panel B presents a binned scatter plot of changes in reweighted employment from January to June 2020 vs. firm size. Data sources: Paychex, Earnin.