Building an Infrastructure for Real-Time Policy Analysis using Administrative Data: A Case Study of the Impacts of Stimulus Checks in the COVID-19 Crisis

Raj Chetty, John Friedman and Michael Stepner

A recurring challenge for evidence-based policymaking is an absence of timely evidence—especially during crises. In this talk, we plan to (1) discuss a real-time economic tracker that we built using private sector data to address this problem during the COVID-19 crisis, (2) present an application that shows how these data have been useful in informing the design of the current stimulus bill, and (3) discuss an effort to institutionalize this approach by constructing "real time" national accounts in collaboration with the main statistical agencies of the government. Many of those researchers will be attending the Thursday morning session on research using government administrative data, and our talk will fit into the theme with a forward-looking emphasis on the future of official government statistics. We believe a talk in the morning administrative data session would be useful in letting people know about this initiative, both in the NBER PE group and among government staff, and potentially provide further opportunities for collaborators to contribute to this public good. What follows below is an abstract describing what we would plan to present and how we would plan to structure the talk, drawing on a new analysis of the recent stimulus payments that is based on a completed paper (currently at the R&R stage at the QJE).

During the course of the COVID-19 pandemic, we (jointly with Nathaniel Hendren and the Opportunity Insights Team) built a publicly available database that tracks economic activity at a granular level in real time. Using anonymized data from a dozen private companies, we report daily statistics on consumer spending, business revenues, employment rates and other key indicators. As one example of the utility of these new forms of data, we analyze the heterogeneous consumption responses to recent Economic Impact Payments (stimulus checks) within three weeks of those checks being deposited into banks accounts. We find that the $600 stimulus checks deposited at the beginning of January 2021 increased spending among lower-income households significantly, but had little impact on spending among higher-income households. By contrast, the initial round of stimulus checks deposited in April 2020 raised spending significantly and similarly among households of all income levels. The difference in consumption responses is consistent with other evidence on the differential trajectory of the recession across the income distribution: while high income workers returned to pre-pandemic employment levels and amassed substantial savings, low-income workers have experienced prolonged unemployment and remain substantially below baseline. This analysis illustrates how changing economic conditions alter the impacts of economic policies, and demonstrates how real-time data can provide continually-updated evidence to guide policymaking (in this case, targeting the $1,400 stimulus checks currently under consideration).

We detail the construction of this database and analyze the effects of an initial set of policies enacted in the first half of 2020 in our paper "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." During this 20 minute talk, we will briefly discuss the data from that paper but focus on the more recent analysis of Economic Impact Payments that we released on January 26th 2021, which is also attached.

Attachments

1. Nontechnical one-page summary of Economic Impact Payments analysis
2. Technical appendix to Economic Impact Payments analysis
3. Paper: "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data"

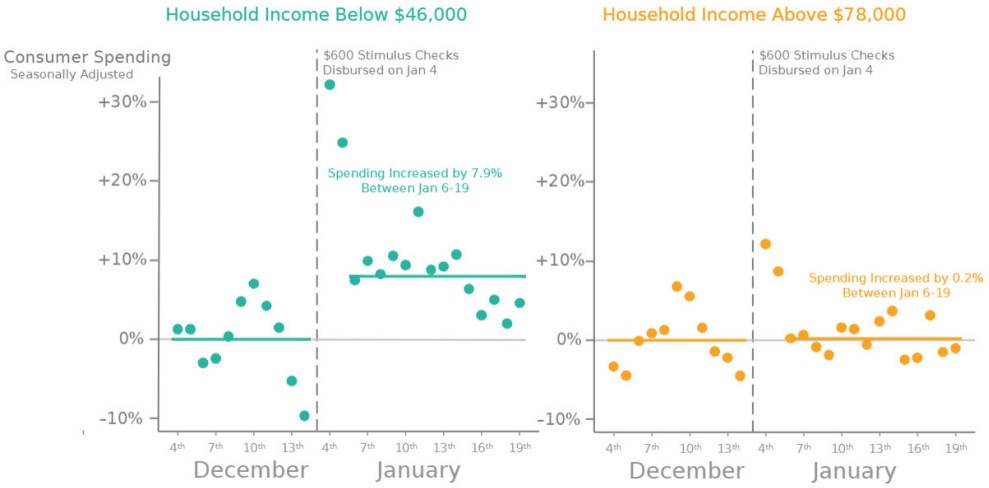
OPPORTUNITY INSIGHTS
ECONOMIC TRACKER

Effects of January 2021 Stimulus Payments on Consumer Spending

Raj Chetty, John Friedman & Michael Stepner

We use real-time spending data from the Opportunity Insights Economic Tracker to assess the effects of the COVID-Related Tax Relief Act of 2020 stimulus payments on consumer spending across different income groups. We find that the $600 stimulus checks increased spending among lower-income households significantly, but had little impact on spending among higher-income households.[1]

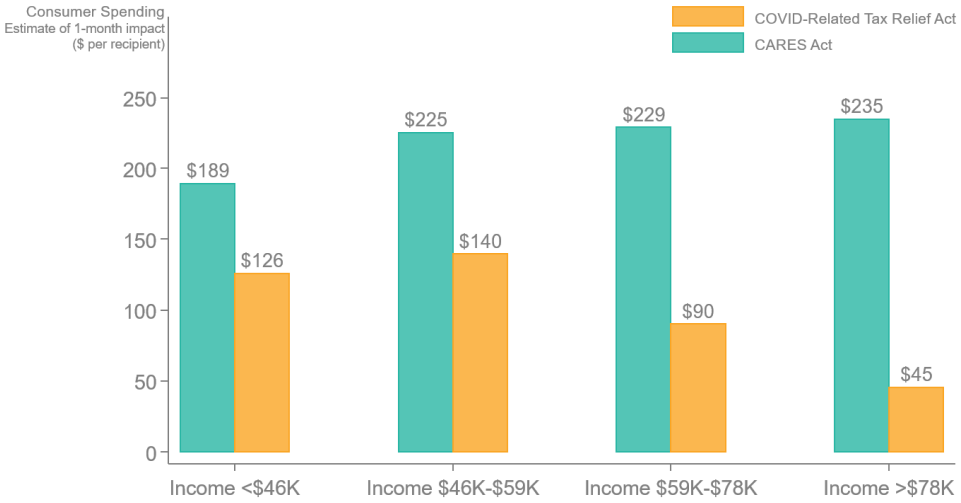
DAILY CONSUMER SPENDING ON DEBIT AND CREDIT CARDS, PRE VS. POST $600 STIMULUS



Data Source: Affinity Solutions

Over the first month after stimulus payments were received, our estimates imply that households with incomes above $78,000 will spend only $45 of the $600 payments they received. This is much less than the effect of the April stimulus payments on the spending of these higher-income households, which arrived at a time when unemployment was surging. Since then these households have largely returned to work, and have even accrued additional savings.

TOTAL MONTHLY SPENDING EFFECT OF FIRST AND SECOND STIMULUS PAYMENTS, BY HOUSEHOLD INCOME GROUP



Data Source: Affinity Solutions and ACS. Since the April stimulus payments were $1,200 for most families, we halve the values in this chart in order to make them comparable to the per-dollar effects of the recent $600 stimulus payments.

Based on these results, we estimate that **households earning more than $78,000 will spend only $105 of the $1400 stimulus check they receive** - implying that $200 billion of additional government expenditure will lead to only $15 billion of additional spending.[2] Targeting the next round of stimulus payments toward lower-income households would save substantial resources that could be used to support other programs, with minimal impact on economic activity.

A technical appendix for this analysis can be found on the Opportunity Insights website as well as more analysis on the economic effects of the COVID-19 pandemic.

1. Low- and high-income households are defined as those residing in zip codes with average household incomes below $46,000 and above $78,000 respectively.
2. This figure is calculated based on reduced spending among multi-person households with incomes greater than $78,000 and single-headed households with incomes greater than $50,000 respectively.

ABOUT OPPORTUNITY INSIGHTS

Opportunity Insights is a non-partisan, not-for-profit organization located at Harvard University that seeks to translate insights from rigorous, scientific research to policy change by harnessing the power of "big data" using an interdisciplinary approach.

Requests for additional information on the data or technical questions can be directed to info@opportunityinsights.org.

Updated as of February 4th, 2021

Technical Appendix to "Effects of January 2021 Stimulus Payments on Consumer Spending"

Opportunity Insights
Updated February 5, 2021

This appendix provides further details on our approach to estimating the effects of direct stimulus on consumer spending for payments made under the CARES Act 2020 and the COVID-Related Tax Relief Act 2020.

We estimate effects on consumer spending by constructing daily national consumer spending series using data from Affinity Solutions Inc, an aggregator of consumer credit and debit card information capturing nearly 10% of debit and credit card spending in the U.S; for further details on Affinity Solutions, see Section II.A and Appendix B of Chetty et al. (2020). The spending series we construct from the Affinity data closely track patterns in the nationally representative Advance Monthly Retail Trade Survey (MARTS) (Figures 1b-1c, Chetty et al. 2020). We then use these series to estimate the effects of each round of stimulus payments on national consumer spending. Finally, we use data on eligibility for stimulus payments to assess the impact per recipient.

Our main finding is that January 2021 stimulus payments substantially increased spending among lower-income households, but had little impact on spending among higher-income households. This contrasts with April 2020 stimulus payments, which increased consumer spending for all household income quartiles.

These differences stem from the different contexts in which the stimulus payments were distributed. Appendix Figure 1a shows that the national savings rate increased substantially since the onset of the COVID-19 pandemic. Using Affinity data to disaggregate national trends, we find that between April 2020 and the end of December 2020, high-income (top ZIP Code income quartile) households reduced their spending relative to 2019 by a cumulative total of slightly over $11,000 (Appendix Figure 1b).[1] As the employment impacts of COVID-19 were relatively small for high-income workers (Figure 7, Chetty et al. 2020), high-income households built up substantial savings balances over this period. By contrast, low-income (bottom ZIP Code income quartile) households, which reduced their spending by a smaller proportion (Appendix Figure 1b) and were more likely to lose employment (Figures 2 and 7, Chetty et al. 2020), built up smaller savings balances, especially after the expiry of supplemental unemployment insurance (Farrell et al. 2020). As households with more liquid balances tend to have lower marginal propensities to consume from income shocks (Johnson et al. 2006; Broda and Parker 2014; Jappelli and Pistaferri 2014), these trends between April 2020 and January 2021 dampened high-income households' response to stimulus payments, while leaving low-income households largely unaffected.

This appendix is organized in two sections. The first section details our methodology in constructing estimates of the effects of stimulus payments on consumer spending. As we do not observe incomes in the Affinity data, these estimates are constructed using ZIP Code income quartile as a proxy for income. Sections 2 examines the implications of this imputation approach by comparing our spending series using

[1] We use 2010 Census ZIP Code Tabulation Areas (ZCTAs) to perform all geographic analyses of ZIP Code-level data. Throughout the text, we refer to these areas simply as "ZIP Codes".

ZIP Code income quartiles to spending series in which income is directly observed. We also conduct a bounding exercise examining the implications of our ZIP Code-level results for underlying variation in MPCs by household income.

1. Methodology to construct main estimates

Estimating the effects of stimulus payments on indexed consumer spending by neighborhood income level

Construction of national spending series. We first construct a national daily consumer spending series by ZIP Code income quartile using data from Affinity Solutions. We receive raw data from Affinity Solutions at the county x ZIP Code income quartile x industry x date level, where cells with fewer than five unique cards with transactions are masked. We allocate ZIP Codes to quartiles using 2014-2018 ACS estimates of ZIP Code median household income, and then calculating the national population-weighted distribution of ZIP Code median incomes. For further details, see Appendix A of Chetty et al. (2020).

We first aggregate the raw data to calculate total card spending on each date, in each ZIP Code income quartile. We then index spending in each ZIP Code income quartile relative to spending in January 2019. Finally, we residualize indexed spending with respect to day-of-week fixed effects calculated using data from 2019.

Calculation of effects of stimulus payments on consumer spending. We use a difference-in-difference approach to estimate the effects of stimulus payments on indexed consumer spending within each income quartile. The most recent date for which Affinity Solutions data are available is January 19, 2021. Using Daily Treasury Statements on Economic Impact Payments distributed either by check or by electronic funds transfer (EFT), we compute that approximately two-thirds of the total $166 billion allocated to stimulus payments under the COVID-related Tax Relief Act of 2020 was distributed on January 4 2021, with another 13% distributed on January 6 2021. We therefore use the period January 4-January 19, 2021 as the treatment period for stimulus payments. We seasonally-adjust consumer spending on each date by subtracting indexed spending in each date x income quartile in December 2019 (January 2020) from the corresponding date x ZIP Code income quartile in December 2020 (January 2021). This creates first-differenced data on changes in consumer spending in December 2020 and January 2021 relative to the previous year. We then take the second difference by subtracting mean indexed consumer spending within each ZIP Code income quartile during December 4-14th, 2020 from indexed consumer spending on each date in January 2021. We choose December 4-14th as the pre-stimulus control period because the daily volatility of spending increases markedly during the holiday period, beginning (in the Affinity data) on December 15th.

Figure 1 of "Effects of January 2021 Stimulus Payments on Consumer Spending". The first figure in our analysis, titled "Daily Consumer Spending on Debit and Credit Cards, Pre vs. Post $600 Stimulus", displays these daily values of seasonally-adjusted consumer spending relative to December 2020 over the periods December 4-December 14, 2020 and January 4-January 19, 2021, for ZIP Codes in the bottom income quartile (left panel) and the top income quartile (right panel). Each point represents seasonally-adjusted consumer spending at each date, constructed as described above.

The annotation in each panel displays the estimated effect of stimulus payments on consumer spending over the period January 6-19, 2020. We calculate this effect by averaging seasonally-adjusted consumer

spending over the period January 6-January 19, 2021 relative to the period December 4-14, 2020, within each income quartile. We exclude the short-term effect on consumer spending over the first two days after the receipt of stimulus payments.

Estimated portion of stimulus check spent

We take several steps to convert these estimates on indexed consumer spending to estimates of the portion of stimulus checks that was spent. The procedure differs in subtle ways between the April 2020 and January 2021 stimulus rounds. We therefore describe the procedure for January 2021 in full, noting in places where the procedure for April 2020 differs.

Calculation of indexed monthly spending effects of the COVID-related Tax Relief Act of 2020. We first estimate effects on indexed consumer spending over the first month after stimulus payments were received. To do so, we begin by calculating the effect on consumer spending separately over the first two days after the payments (Jan 4-5), following the procedure described above (replacing Jan 6-19 with Jan 4-5). We do so due to the clear "spike" in spending on these days, relative to the days that follow. For the bottom three ZIP Code income quartiles, we then calculate the monthly effect in each ZIP Code income quartile as the mean daily effect over the period January 4-5 multiplied by 2, plus the mean daily effect over the period January 6-19 multiplied by 14, plus the daily effect on January 19 multiplied by 15, making the assumption that stimulus effects observed on January 19 remain constant through February 3. For the top ZIP Code income quartile, we assume that spending remains constant at the level observed over the period January 6-19 through February 4. This assumption differs from the one we make for the other income quartiles because the second-differenced spending level observed for the top income quartile on January 19th is negative.

Calculation of indexed monthly spending effect of the CARES Act of 2020. We estimate the effect of the economic impact payments disbursed in April 2020 as part of the CARES Act on consumer spending using a similar methodology to the one outlined above with the following exceptions: (1) we use April 3rd-April 11th as the pre-period and April 15th-April 30th as the post period; and (2) we obtain our monthly impacts for all ZIP Code income quartiles by taking the sum of consumer spending over April 15th-16th and adding the average spending over the period April 17th-April 30th, multiplied by 29. We impute spending on Easter in both 2019 and 2020 by taking the average of the indexed, seasonally adjusted value of the adjacent days in each year.

Calculating Spending Changes per Check Recipient. We calculate total Affinity Solutions spending in each ZIP Code income quartile in January 2019. We then inflate the spending observed in Affinity data to estimate total spending for the full U.S. population by multiplying by the ratio of January 2019 total spending, as measured in the National Income and Product Accounts, on components of PCE that are likely captured in credit/debit card spending to January 2019 total spending in the Affinity data. For further details on the components of PCE that are likely captured in credit/debit card spending, see notes to Figure 1 of Chetty et al. (2020). We then multiply our estimated effects on indexed spending by January 2019 total spending to calculate effects on spending in dollar terms. We express the effect on consumer spending per stimulus check by dividing our estimated effects on consumer spending (in dollar terms) by the estimated number of recipients of stimulus checks in each ZIP Code income quartile.

Calculation of eligibility rates. To calculate the number of recipients in each ZIP Code income quartile, we begin by calculating stimulus payment eligibility rates by ZIP Code income quartile. We treat households

receiving stimulus checks of amounts less than $600 as eligible in proportion to the checks they received, *e.g.* we assign an eligibility value of 0.5 to households receiving $300 stimulus checks. We retrieve ACS 2014-2018 data on the count of individuals in single vs. married households and the count of households in each income bin in each ZIP Code.

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. We use total reported income in the ACS as a proxy for adjusted gross income (AGI), which determines eligibility for stimulus payments.

We then calculate the eligibility rate at the ZIP Code level in three steps: (1) computing the share of people in single-headed households who are eligible for stimulus payments; (2) computing the share of people in joint-filing households who are eligible for stimulus payments; and (3) combining these estimates using data on marriage rates to create overall ZIP Code-level eligibility rates.

(1) *Eligibility among single-headed 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. That is, we assume that stimulus payments were phased out for single-headed households earning over $75,000 at a rate of $5 per person per $100 increment over $75,000, with payments fully phased out at incomes over $87,000 ($99,000) for the January 2021 (April 2020) round of stimulus payments. 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. 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. Under the assumption that incomes are uniformly distributed within each ZIP Code x income bin, we assign an eligibility rate of 25% for these households.

(2) *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; instead, we assume that stimulus payments were phased out for households earning over $150,000 at a rate of $5 per person per $100 increment over $150,000, with payments fully phased out at household incomes over $174,000 ($198,000) for the January 2021 (April 2020) round of stimulus payments. 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 together. Under the assumption that incomes are uniformly distributed within each ZIP Code x income bin, we assign an eligibility rate of 25% for these households.

(3) *ZIP Code-level eligibility rates.* Steps (1) and (2) 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; intuitively, 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 $600 and by the population within each ZIP Code income quartile to calculate total expected stimulus spending in each ZIP Code income quartile.

Share of stimulus spending received. Not all households received payments at the beginning of our treatment period. To calculate stimulus amounts received, we multiply the total amount appropriated for stimulus payments under the COVID-related Tax Relief Act of 2020 by the share of stimulus spending that had been distributed by January 10, 2021. The total amount appropriate for stimulus payments under the COVID-related Tax Relief Act of 2020 is $166 billion (Rows 1-2 of Table 2, Sherlock, Gravelle, Marples and Keightley 2021). We use Daily Treasury Statements to calculate that total spending on stimulus payments prior to January 10 was roughly $130 billion. We therefore find that roughly 78% of payments had been distributed by January 10. A similar calculation for payments made under CARES Act in April 2020 yields estimates that roughly 58% of payments were distributed by April 21.[2] Under the assumption that this rate is constant across ZIP Code income quartiles, we multiply total expected stimulus spending in each ZIP Code income quartile by 78%. We then divide our effects on aggregate consumer spending by the estimated number of recipients to calculate effects per recipient. 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; this would cause us to overestimate the total number of stimulus recipients in the lowest ZIP Code Income Quartile, leading to an underestimate of the amount spent per recipient for these households.

Figure 2 of "Effects of January 2021 Stimulus Payments on Consumer Spending". The second figure in our analysis, titled "Portion of $600 Stimulus Check Spent, by Household Income Group" presents these estimates by income quartile. Since the April stimulus payments were $1,200 for most families, we halve the values for the CARES Act in this chart to make them comparable to the per-dollar effects of the recent $600 stimulus payments. For example, the "$189" spending statistic for the lowest ZIP Code income quartile households corresponds to first-month spending of $189 x 2 = $378 out of each $1200 check paid in April 2020.

Implications for proposed $1,400 stimulus checks

Estimated proposed spending on stimulus payments to high-income households. We calculate proposed spending on multi-headed households earning over $78,000 and single-headed households earning over $50,000 by applying the income thresholds in the Caring for Americans with Supplemental Help (CASH) Act 2020 to 2014-2018 ACS microdata. As before, for simplicity, we ignore the effects of the number of children on phase-out rates.

Estimated effect of $1,400 stimulus payments on spending among high-income households. We assume that the response of consumer spending to $1,400 stimulus payments would be the same in proportional terms as the response of consumer spending to $600 stimulus payments, within each ZIP Code income

[2] In an update to this analysis on February 4th, 2021, we adjusted our calculation of the fraction of stimulus checks received before April 21st to its current value of 58%. Our previous analysis used a higher value of this number, leading to lower estimated spending effects for each ZIP Code income quartile in Figure 2.

quartile. We also assume that the response of consumer spending that we estimate within the top ZIP Code income quartile, with median household incomes above $78,000, approximates the response of consumer spending among high-income households. Sections 2 presents the case for the validity of this assumption.

Making these assumptions, we calculate the per-recipient effect of the proposed $1400 stimulus checks as the effect of $600 stimulus checks, multiplied by ($1400/$600); following this procedure, we estimate a per-recipient effect of stimulus checks of around $105 among high-income households. To calculate the national aggregate effect on consumer spending, we note that a spending response of $105 represents increased spending of 7.5% of the initial stimulus amount; we therefore multiply total spending on high-income households ($200 billion) by 7.5% to calculate a national effect on consumer spending of around $15 billion.

2. Implications for Stimulus Effects at the Household Level

We do not directly observe cardholders' incomes in the Affinity data. For this reason, we use the median household income of the cardholders' residence, as measured in 2014-2018 ACS data, to split households into groups for our analysis. The section examines the validity of inferring something about the spending patterns of high- vs. low-income households based on this area-level grouping. We discuss these issues in two steps: First, we assess whether the differences in households between areas can be understood solely as differences in the distribution of underlying income; second, we use the specific household income distributions within each ZIP Code income quartile bucket to translate our results into bounds on the underlying spending responses of high-income households.

Validity of Area-Level Estimates. The key assumption to using median ZIP Code income as a proxy for individual income is that there are no sources of individual-level variation in consumer responses to the COVID-19 pandemic that are correlated with median ZIP Code income, other than income itself. For instance, if low-income households in high-income ZIP Codes experienced larger declines in spending than low-income households in low-income ZIP Codes, then our ZIP Code imputation approach would tend to overstate the difference in spending between high-income and low-income households.

Bachas et al. (2020), who observe incomes directly for JPMorgan Chase clients, address this assumption in Table 3, which shows that very little of the relationship between income and the initial decline in consumer spending in response to COVID-19 is driven by geographic fixed effects. This suggests that remaining differences in consumer spending between high-income and low-income areas reflect the differing income composition of households in those areas, rather than reflecting the effects of other characteristics that are correlated with median income.[3] Consistent with this evidence, the series we construct using our ZIP Code imputation approach are generally well aligned with those in Bachas et al. (2020) from January through to May 2020, the last date for which spending series are available in Bachas et al. (2020).

We also directly assess the accuracy of our ZIP Code imputation approach when estimating the effects of April 2020 stimulus payments. (Unfortunately, similar results are not yet available for the January 2021 stimulus payments.) To do so, we compare our estimates of the effects of stimulus payments by ZIP Code

[3] Of course, the fact that the key assumption for our imputation approach holds in the initial period of response to the pandemic does not guarantee that it will continue to hold throughout the full sample period.

income quartile to estimates constructed using the data in Figure 5 of Bachas et al. (2020). The data in Bachas et al. (2020) are available on a weekly basis; for this reason, we cannot directly implement our main approach to estimating the effects of stimulus payments, which treats the first two days after stimulus payments were received differently to subsequent days. Instead, in the Bachas et al. (2020) data, we compare mean indexed spending over the two weeks prior to the receipt of stimulus payments to spending over the subsequent five weeks.[4] In the Affinity data, we use the difference-in-difference approach described in Section 1, except restricting the sample to the same period as in the Bachas et al. (2020) data. Appendix Figure 2 demonstrates that estimated effects on indexed consumer spending are similar in the two samples.[5]

Based on these comparisons, we conclude that the key assumption stated above is likely to hold: households differ in their spending responses between areas with higher and lower ZIP Code median incomes due to the differences in the underlying household income distributions, rather than other correlated factors.

Implications of Area-Level Estimates for Household-Level Responses. We now use the underlying household income distribution, as it varies between ZIP Code income quartiles, to translate our ZIP Code income quartile-level results into results on how the underlying stimulus effects vary by individual household income. We focus in this section specifically on estimating the spending responses of high-income households.

We begin by characterizing the distribution of stimulus-eligible household within the fourth quartile of ZIP Code median household incomes. We do so using ACS 2014-2018 data on the count of individuals in single vs. married households and the count of households in each income bin in each ZIP Code. As above when computing eligibility rates, 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 within a ZIP Code). We assign eligibility for stimulus payments as described above, restricting our attention to the January 2021 round of stimulus payments. Appendix Figures 3a and 3b display the resultant distribution of household incomes among married and unmarried households respectively, restricting to stimulus-eligible households in top income quartile ZIP Codes. Then, in order to combine the income distributions for singles and married households, we follow the Congressional Budget Office and assume that single households have equivalent spending responses to married households with 1.41 times higher household incomes (the "square-root equivalence scale") (Congressional Budget Office 2020).

[4] Stimulus payments on April 15 fall in the middle of a week in the Bachas et al. (2020) data. We treat the week of April 12-18 as a post-stimulus week.

[5] Appendix Figure 2 displays the effects of stimulus payments on indexed consumer spending, whereas Figure 2 of "Effects of January 2021 Stimulus Payments on Consumer Spending" displays spending in dollar terms per recipient of stimulus payments. Spending per recipient of stimulus payments is relatively constant across income quartiles. By contrast, the effect of stimulus payments on indexed consumer spending is decreasing with income. This is because (1) lower-income households were more likely to be eligible for stimulus payments, and (2) spending among lower-income households was lower at baseline, and so a constant increase in spending in dollar terms leads to a larger effect on indexed spending among lower-income households.

Appendix Figure 3c plots this combined distribution. Even in the top quartile of ZIP Codes, by median income, 50% of stimulus-eligible households have incomes lower than $74,000, with a substantial fraction (25%) even below $42,000.

Bound on household-level stimulus response. We now rely on the distribution from Appendix Figure 3c to bound the underlying stimulus response for households with incomes above $75,000. Because roughly half of stimulus payments in these ZIP Codes went to households with incomes above and below $75,000, the overall spending response for these ZIP Codes (in Figure 2) is the average of the spending response for these two groups. Under the assumption that the response for those households earning more than $75,000 is not greater than that for households earning less than $75,000 – an assumption that is consistent with the results in Figure 2 as well as previous literature – then the response for households earning above $75,000 cannot be less than the overall estimated response of $45 per $600 in Figure 2.

More generally, the underlying distributions of household incomes suggest that the differences in Figure 2 of our analysis understate the underlying response differences, as they vary with household income levels. Intuitively, households in high-median-income ZIP Codes are on average higher income than those in low-median-income ZIP Codes, but there is substantial overlap in the distributions. As a result, the responses for high-income households themselves are likely even lower than suggested by Figure 2, once one corrects for the larger responses from those low-income households who live in high-income ZIP Codes. Conversely, the response for low-income households is likely greater than the $126 per $600 suggested by Figure 2, since there are a number of higher-income households whose lack of response pulls down the average. We thus consider our ZIP Code-level estimates as bounds for the underlying household-level responses.

References

Bachas, Natalie, Peter Ganong, Pascal J. Noel, Joseph S. Vavra, Arlene Wong, Diana Farrell, and Fiona E. Greig. 2020. "Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data." NBER Working Paper No. 27617.

Broda, Christian, and Jonathan A. Parker. 2014. "The Economic Stimulus Payments of 2008 and the aggregate demand for consumption." *Journal of Monetary Economics* 68: S20-S36.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and the Opportunity Insights Team. 2020. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." NBER Working Paper No. 27431.

Congressional Budget Office. 2020. "The Distribution of Household Income, 2017."

Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, Daniel Sullivan, and Joseph Vavra. 2020. "The unemployment benefit boost: Trends in spending and saving when the $600 supplement ended." JPMorgan Chase Institute Policy Brief.

Jappelli, Tullio and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6(4): 107-136.

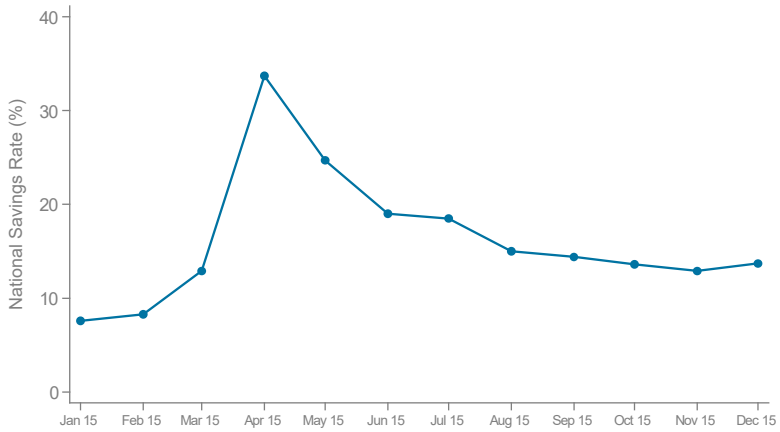
Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96(5): 1589-1610.

Sherlock, Molly F., Jane G. Gravelle, Donald J. Marples, and Mark P. Keightley. 2021. “The COVID-Related Tax Relief Act of 2020 and Other COVID-Related Tax Provisions in P.L. 116-260.” Congressional Research Service Report R46649.

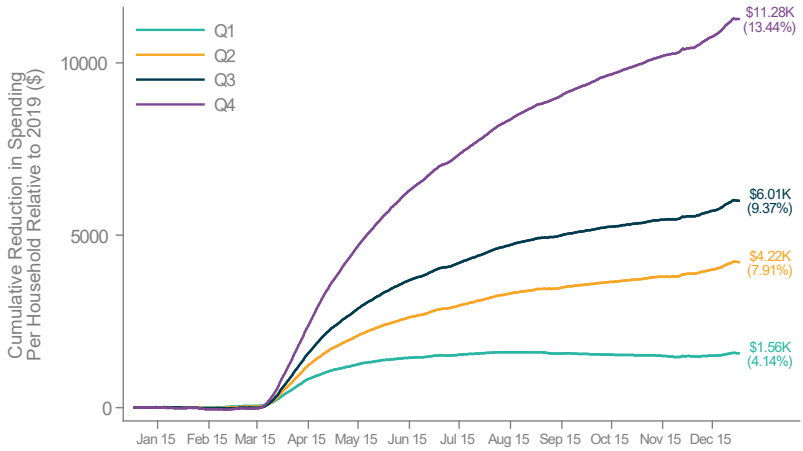
APPENDIX FIGURE 1

Trends in Saving Around COVID-19 Pandemic

A. National Savings Rate in Bureau of Economic Analysis Data



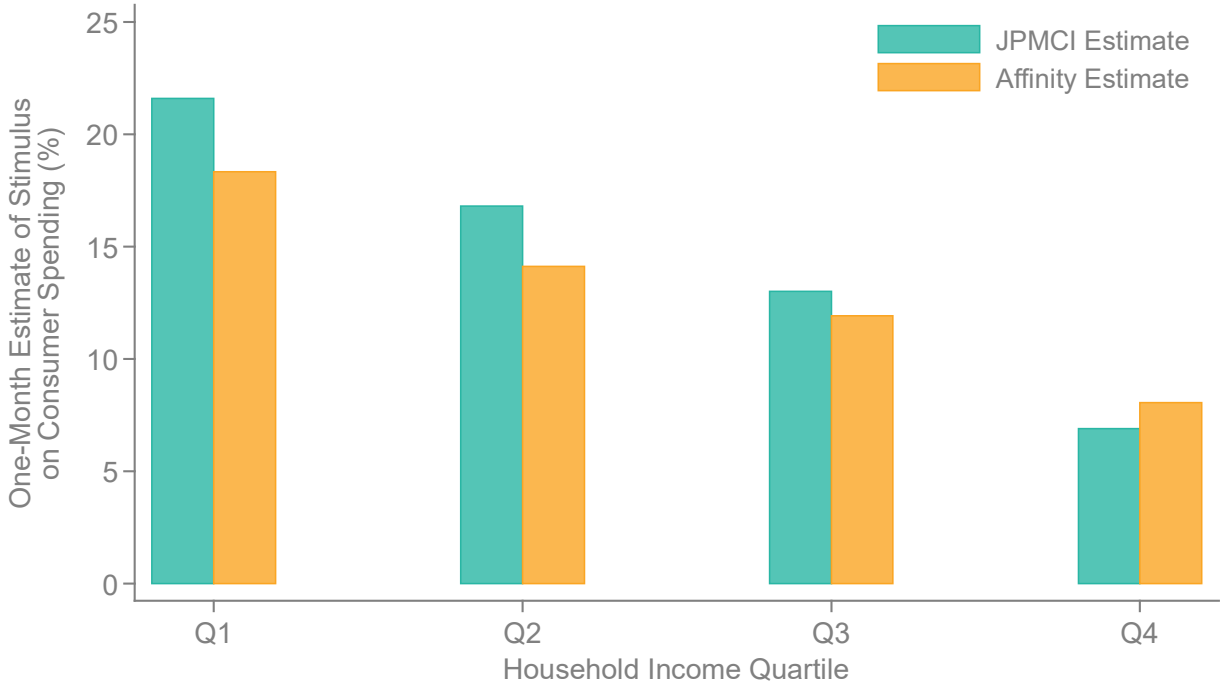
B. Reduction in Spending per Household by Income Quartile in Affinity Data



Notes: This figure shows trends in savings throughout 2020. Panel A shows the seasonally-adjusted national personal savings rate in Bureau of Economic Analysis data. Panel B shows the cumulative reduction in spending per household relative to 2019 in Affinity Solutions data, by income quartile. To construct Panel B, we first compute total spending by income quartile on each day in 2019 and 2020. We then rescale from Affinity spending to national spending by multiplying by the ratio of January 2019 total spending on components of PCE that are likely captured in credit/debit card spending to January 2019 total spending in the Affinity data. We measure spending on components of PCE that are likely captured in credit/debit card spending using the National Income and Product Accounts; for details, see notes to Figure 1 of Chetty et al. (2020). We then calculate the difference between 2020 spending and 2019 spending on each date, and plot the cumulative reduction in spending over time, divided by the number of households in each ZIP Code income quartile, as measured in 2014-2018 ACS data. The annotation to the right of the panel displays the cumulative reduction in spending over the year 2020 for each income quartile.

APPENDIX FIGURE 2

Estimated One-Month Impacts of Stimulus on Consumer Spending in JP Morgan Chase Institute (JPMCI) vs. Affinity Data

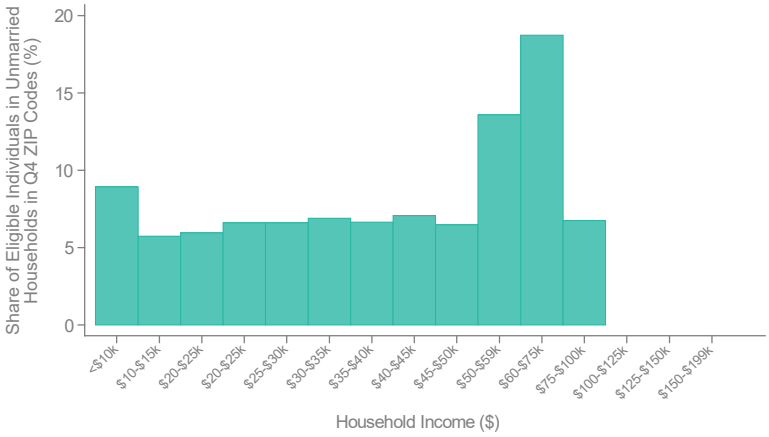


Notes: This figure compares estimates of the impact of stimulus spending in JP Morgan Chase Institute data contained in Bachas et al. (2020) to corresponding estimates in Affinity data. In the JP Morgan Chase Institute data, household income quartiles are directly observed; in the Affinity data, household income quartiles are imputed using ZIP Code median income. To estimate the impact of stimulus payments in JP Morgan Chase Institute data, we begin by retrieving weekly data on indexed consumer spending by income quartile from Figure 5 of Bachas et al. (2020). We estimate the effect of stimulus payments in each income quartile as mean indexed spending over the five weeks after the receipt of stimulus payments minus mean indexed spending over the two weeks prior to the receipt of stimulus payments. Stimulus payments on April 15 fall in the middle of the week April 12-18 in Bachas et al. (2020); we treat this entire week as a post-stimulus week. In the Affinity data, we use a difference-in-difference approach as described in Section 1, restricting the sample to the same period as in the Bachas et al. (2020) data.

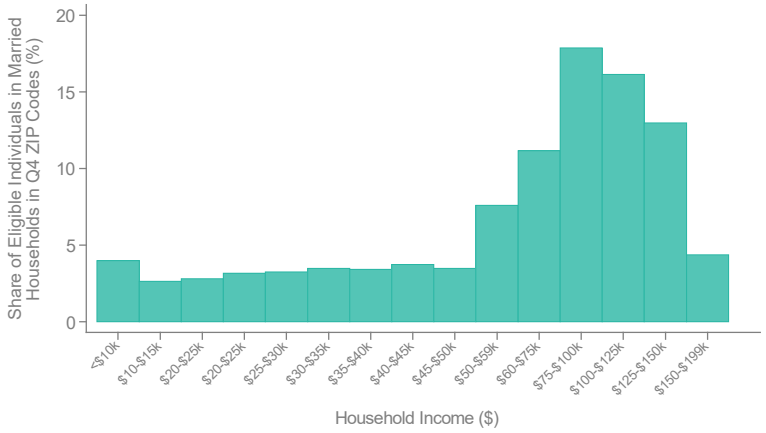
APPENDIX FIGURE 3

Household Income Distributions in Top Income Quartile ZIP Codes

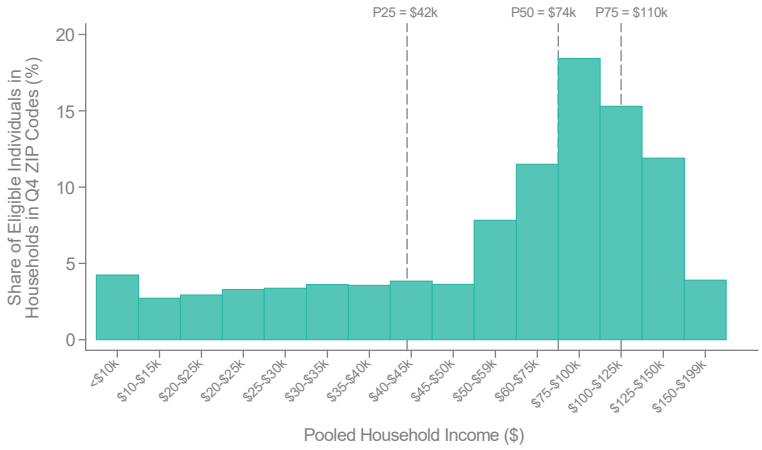
A. Distribution of Household Incomes Among Eligible Unmarried Households in Top Income Quartile ZIP Codes



B. Distribution of Household Incomes Among Eligible Married Households in Top Income Quartile ZIP Codes



C. Distribution of Household Income Among Married Households Eligible for January 2021 Stimulus Payments Residing in Top Income Quartile ZIP Codes



Notes: This figure presents the household income distribution of individuals residing in top income quartile ZIP Codes who are eligible for January 2021 stimulus checks. To construct each panel, we use 2014-2018 ACS data on the distribution of incomes and household types by ZIP Code. We assign eligibility for stimulus payments as described in section 1. Panel A shows the income distribution among unmarried households. Panel B shows the income distribution among married households. To construct Panel C, we multiply household income among unmarried households by 1.41, under the assumption that incomes are uniformly distributed within each income bin x ZIP Code. The annotations in Panel C display the 25th, 50th and 75th percentiles of the income distribution.

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

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

First Version: May 2020
This Version: November 2020

Abstract

We build a publicly available database that tracks economic activity at a granular level in real time using anonymized data from private companies. We report daily statistics on consumer spending, business revenues, employment rates, and other key indicators disaggregated by ZIP code, industry, income group, and business size. Using these data, we study how COVID-19 affected the economy by analyzing heterogeneity in its impacts. We first show that high-income individuals reduced spending sharply in mid-March 2020, particularly in areas with high rates of COVID-19 infection and in sectors that require in-person interaction. This reduction in spending greatly reduced the revenues of small businesses in affluent ZIP codes. These businesses laid off many of their employees, leading to widespread job losses especially among low-wage workers in affluent areas. High-wage workers experienced a "V-shaped" recession that lasted a few weeks, whereas low-wage workers experienced much larger job losses that persisted for several months. Building on this diagnostic analysis, we estimate the causal effects of policies aimed at mitigating the adverse impacts of COVID-19. State-ordered reopenings of economies had small impacts on spending and employment. Stimulus payments to low-income households increased consumer spending sharply, but little of this increased spending flowed to businesses most affected by the COVID-19 shock, dampening its impacts on employment. Paycheck Protection Program loans increased employment at small businesses by only 2%, implying a cost of $377,000 per job saved. These results suggest that traditional macroeconomic tools – stimulating aggregate demand or providing liquidity to businesses – have diminished capacity to restore employment when consumer spending is constrained by health concerns. During a pandemic, it may be more fruitful to mitigate economic hardship through social insurance. More broadly, this analysis shows how public statistics constructed from private sector data can support many research and policy analyses without compromising privacy, providing a new tool for empirical macroeconomics.

---

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

†The Opportunity Insights Economic Tracker Team consists of 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, Abby Hiller, Matthew Jacob, Tyler Jacobson, Margaret Kallus, Laura Kincaide, Cailtin Kupsc, Sarah LaBauve, Maddie Marino, Kai Matheson, Christian Mott, Kate Musen, Danny Onorato, Sarah Oppenheimer, Trina Ott, Lynn Overmann, Max Pienkny, Jeremiah Prince, Sebastian Puerta, Daniel Reuter, Peter Ruhm, Tom Rutter, Emanuel Schertz, Shannon Felton Spence, Krista Stapleford, Kamelia Stavreva, James Stratton, Clare Suter, Elizabeth Thach, Nicolaj Thor, Amanda Wahlers, Kristen Watkins, Alanna Williams, David Williams, Chase Williamson, Shady Yassin, and Ruby Zhang.

I Introduction

Since Kuznets (1941), macroeconomic policy decisions have been made on the basis of publicly available statistics constructed from recurring surveys of households and businesses conducted by the federal government. Although such statistics have great value for understanding total economic activity, they have two limitations. First, survey-based data typically cannot be used to assess variation across geographies or subgroups; due to relatively small sample sizes, most statistics are typically reported only at the national or state level and breakdowns for demographic subgroups or sectors are unavailable. Second, such statistics are typically available only at low frequencies, often with a significant time lag. For example, data on consumer spending disaggregated by geography are only available for selected large metro areas at a bi-annual level in the Consumer Expenditure Survey (CEX). Because of these limitations, existing publicly available macroeconomic statistics are insufficient to study the sources of economic fluctuations and the causal impacts of macroeconomic policies.

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

We organize the paper in three parts. First, we describe how we construct statistics on consumer spending, business revenues, employment rates, job postings, and other key indicators – disaggregated by area (ZIP code or county), industry, income level, and business size – by combining data from credit card processors, payroll firms, and financial services firms. The main challenge in using private sector data sources to measure economic activity is a tension between research value and privacy protection. For research, it is beneficial to use raw, disaggregated data – ideally down to the individual consumer or business level – to maximize precision and flexibility of research designs. But from a privacy perspective, it is preferable to aggregate and mask data to reduce the risk of disclosure of private information. To balance these conflicting interests, one must construct statistics that are sufficiently aggregated and masked to mitigate privacy concerns yet sufficiently granular to support research.[1]

1. An alternative approach pursued by several recent studies (summarized at the end of this section) is to use confidential private data sources for research under non-disclosure agreements with companies. Although a valuable complement to the public data approach we pursue here, the need to acquire data from each company separately typically leads most studies to use one or two datasets and limits the number of researchers who can analyze the data.

We navigate this tradeoff by combining several statistical methods: reporting only changes since January 2020 (rather than raw levels), masking small cells, and pooling data from multiple companies to comply with regulations governing the disclosure of material non-public information. We then clean the raw data by removing data artifacts (e.g., breaks that arise from changes in platforms) and smoothing seasonal fluctuations. Finally, we address the challenge that statistics from specific private firms may not be representative of the population in general. To minimize potential selection biases, we start by obtaining data from companies that have large samples (e.g., at least one million individuals) and span well-defined sectors or subgroups (e.g., small businesses). We then compare each time series to publicly available benchmarks based on representative surveys and use only the series that track publicly available data closely.[2] After establishing these protocols, we report the final statistics using an automated pipeline that ingests data from businesses and reports statistics within a week after the relevant transactions occur.

In the second part of the paper, we use these new public data to analyze the economic impacts of COVID-19. National accounts reveal that GDP fell in the second quarter of 2020 following the COVID-19 shock primarily because of a reduction in consumer spending. We therefore begin by examining the drivers of changes in consumer spending using credit and debit card data. We find that spending fell primarily because *high-income* households started spending much less. As of mid-July, 49% of the reduction in total spending since January came from households in the top income quartile, while 7% came from households in the bottom quartile.[3] This is both because the rich account for a larger share of spending to begin with and because they cut spending more in percentage terms: top-quartile households spent 13% less as of mid-July than in January 2020, whereas bottom-quartile households spent only 4% less. Spending reductions were concentrated in services that require in-person physical interaction, such as hotels and restaurants, consistent with contemporaneous work by Alexander and Karger (2020) and Cox et al. (2020). These findings suggest that high-income households reduced spending primarily because of health concerns rather

Our goal is to assess whether one can produce public statistics that can deliver insights analogous to those obtained from confidential microdata, thereby expanding the scale and timeliness of empirical macroeconomic research.

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

3. We impute income as the median household income (based on Census data) in the cardholder's ZIP code. We verify the quality of this imputation procedure by showing that our estimates of the gap in spending reductions by income group are aligned with those of Cox et al. (2020), who observe income directly for JPMorgan Chase clients.

than a reduction in income or wealth, perhaps because they were able to self-isolate more easily than lower-income individuals (e.g., by substituting to remote work).

Next, we turn to the impacts of the consumer spending shock on businesses. To do so, we exploit geographic variation in the demand shocks businesses face arising from the fact that in-person services are typically produced by small businesses (e.g., restaurants) that serve customers in their local area. Small business revenues in the highest-income and highest-rent ZIP codes (e.g., the Upper East Side of Manhattan) fell by more than 65% between March and mid-April, compared with 30% in the least affluent ZIP codes. These reductions in revenue resulted in a much higher rate of small business closure in affluent areas within a given county than in less affluent areas. This was particularly the case for non-tradable goods that require physical interaction – e.g., restaurants and accommodation services. Small businesses that provide fewer in-person services – such as financial or professional services firms – experienced much smaller losses in revenue even in affluent areas.

As businesses lost revenue, they passed the shock on to their employees, particularly low-wage workers. Building on results first established using other data sources by Cajner et al. (2020), we show that employment rates fell by 37% around the trough of the COVID recession (April 15, 2020) for workers with wages rates in the bottom quartile of the pre-COVID wage distribution. By contrast, employment rates fell by 14% for those in the top wage quartile. Employment for high-wage workers also rebounded much more quickly: employment levels for workers in the top wage quartile were almost back to pre-COVID levels by the end of May, but remained 20% below baseline for low-wage workers even as of October 2020. The greater persistence of job losses for low-wage workers is not explained purely by sectoral differences. Even in sectors where spending rebounded to baseline levels, such as retail trade, employment of low-wage workers remained far below baseline levels, suggesting that economic activity may have shifted toward modes of production using less low-wage labor (Jaimovich and Siu 2020).

Low-wage individuals working at small businesses in affluent areas were especially likely to lose their jobs. At small businesses located in the highest-rent ZIP codes, more than 45% of workers were laid off within two weeks after the COVID crisis began; in the lowest-rent ZIP codes, fewer than 25% lost their jobs. Job postings also fell much more sharply in more affluent areas, particularly for positions requiring less education. As a result, unemployment rates surged even in affluent areas that have typically had relatively low unemployment rates in previous recessions. The impacts on displaced workers persist over time: low-income people working in more affluent ZIP codes pre-COVID (e.g., Manhattan) remain less likely to be employed months after losing their jobs and

reduce their own spending levels more sharply than their peers working in nearby less affluent areas (e.g., the Bronx).

In summary, the initial impacts of COVID-19 on economic activity were largely driven by a reduction in spending by higher-income individuals due to health concerns, which in turn affected businesses that cater to the rich and ultimately reduced the incomes and expenditure levels of low-wage employees of those businesses. In the third and final part of the paper, we analyze the impacts of three sets of policies that were enacted shortly after the crisis began in an effort to break this chain of events and mitigate economic losses: state-ordered shutdowns and reopenings, stimulus payments to households, and loans to small businesses.[4]

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

Stimulus payments made to households in mid-April 2020 increased spending among low-income households sharply, nearly restoring their spending to pre-COVID levels by late April, consistent with empirical evidence from Baker et al. (2020) and models of consumption that generate excess sensitivity (e.g., Kaplan and Violante 2014). However, most of this increase in spending was in sectors that require limited physical interaction: purchases of durable goods surged, while consumption of in-person services increased much less. As a result, little of the increased spending flowed to businesses most affected by the COVID-19 shock, potentially limiting the capacity of the stimulus to increase economic activity and employment because of diminished multiplier effects (a broken Keynesian cross), as discussed by Guerrieri et al. (2020).

Loans to small businesses as part of the Paycheck Protection Program (PPP) also had small impacts on employment rates. Employment rates at firms with fewer than 500 employees (which were eligible for PPP assistance) increased by only 2 percentage points after the PPP was enacted

4. This set of policies is not exhaustive: a vast set of other policy efforts ranging from changes in monetary policy to various state-level programs were also undertaken in response to the crisis. We focus on these three policies because they illustrate the ways in which the data we have assembled here can be used for real-time policy analysis, and we hope that future work will use these data to analyze other policies.

relative to larger firms that were ineligible for PPP. Our point estimates imply that the cost per job saved by the PPP was $377,000 ($119,000 at the lower bound of the 95% confidence interval). Contemporaneous studies by Granja et al. (2020), Autor et al. (2020), and Hubbard and Strain (2020) reach similar conclusions using other data sources and research designs. The PPP had modest marginal impacts on employment likely because the vast majority of PPP loans went to inframarginal firms that were not planning to lay off many workers. As a result, providing liquidity to firms is an expensive way to maintain employment rates in the short run, although it remains possible that the PPP may have long-term benefits by reducing permanent business closures.

Our findings suggest that economic recovery from a pandemic ultimately requires restoring consumer confidence by addressing health concerns themselves (e.g., Allen et al. 2020, Romer 2020). Traditional macroeconomic tools – stimulating aggregate demand or providing liquidity to businesses – have diminished capacity to fully restore employment when demand is constrained by health concerns.[5] In such a setting, it may be especially valuable to provide social insurance to reduce hardship for those who have lost their jobs, e.g., via unemployment benefit extensions (Guerrieri et al. 2020). In addition, it may be useful to target employment assistance to low-income people who were working in places that suffered the largest job losses (such as affluent, urban areas), since geographic disparities in unemployment persist for many years due to limited migration (Blanchard and Katz 1992, Austin, Glaeser, and Summers 2018, Yagan 2019).

Our work builds on two literatures: a longstanding literature on macroeconomic measurement and a nascent literature on the economics of pandemics. In the macroeconomic measurement literature, our work is most closely related to recent studies showing that private sector data sources can be used to forecast government statistics (e.g., Abraham et al. 2019, Aladangady et al. 2019, Ehrlich et al. 2019, Cajner et al. 2019, Gindelsky, Moulton, and Wentland 2019, Dunn, Hood, and Driessen 2020). In the COVID-19 pandemic literature, several recent papers – whose results we compare to ours in the course of our analysis below – have used confidential private sector data to analyze consumer spending (e.g., Baker et al. 2020, Chen, Qian, and Wen 2020, Cox et al. 2020), business revenues (e.g., Alexander and Karger 2020), and labor market trends (e.g., Bartik et al. 2020, Cajner et al. 2020, Kurmann, Lalé, and Ta 2020, Forsythe et al. 2020).

Our paper makes two contributions to these literatures. First, it sheds light on the mechanisms

5. During the period we study, the federal government injected substantial income into the economy (e.g., increased UI benefits and stimulus checks) and the Federal Reserve eased monetary policy significantly. These policies may have averted a Keynesian demand-driven recession on top of the COVID supply shock recession. Our point is not that these policies were ineffective, but rather that the remaining shortfall in spending and employment was driven primarily by health concerns and hence may not have been fixable through conventional macroeconomic tools.

through which pandemics affect economic activity. Other studies of the COVID-19 pandemic have focused on one outcome (e.g., spending or employment or job postings) at broad geographies. By combining data sources on multiple outcomes at the ZIP code level, we exploit local geographic variation to provide an integrated picture of how COVID-19 affected the macroeconomy – from changes in consumer spending to in-person business revenue losses to employment changes and impacts on displaced workers. The heterogeneity in impacts we document across ZIP codes – which, to our knowledge, has not been analyzed in prior studies – provides a novel source of local variation to understand macroeconomic dynamics, similar to the geographic variation widely exploited to understand the Great Recession (e.g., Mian and Sufi 2009). In addition, analyzing a suite of outcomes allows us to characterize the impacts of major stabilization policies more fully, from changes in consumer behavior to impacts on businesses' employment and hiring decisions.

Second, and perhaps more importantly, this study opens new approaches to empirical macroeconomics by constructing a new public database of granular macroeconomic statistics. Unlike the aforementioned studies of COVID-19, which use confidential microdata, the results reported here are all produced from what are now *publicly available* data. Where those studies overlap with ours, results from our public data are very similar to those obtained from their microdata. The capacity to use public data substantially expands the number of researchers who can conduct such studies and opens many new avenues for policy and research. For instance, policymakers could adjust policies as they observe their impacts on the economy (much as the Paycheck Protection Program was repeatedly adjusted, but without the benefit of evidence on its ongoing impacts). Moreover, such impacts can be analyzed heterogeneously across areas and subgroups, permitting tailored responses by local governments. In this sense, the data assembled here provide a prototype for a new system of real-time, granular national accounts that we hope will be refined in future work, much as Kuznets (1941) and Summers and Heston (1984, 1991) developed prototypes for national accounts that were refined in subsequent work (e.g., Feenstra, Inklaar, and Timmer 2015).

The paper is organized as follows. The next section describes how we construct the public data series. In Section III, we analyze the effects of COVID-19 on spending, revenue, and employment. Section IV analyzes the impacts of policies enacted to mitigate COVID's impacts. Section V concludes. Technical details are available in an online appendix, and the data used to produce the results can be downloaded from this GitHub repository.

II Construction of Public Database

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

First, we remove artifacts in raw data that arise from changes in data providers' coverage or systems. For instance, firms' clients often change discretely, sometimes leading to discontinuous jumps in series, particularly in small cells. We systematically search for large jumps in series (e.g., >80%), study their root causes by consulting with the data provider, and address such discontinuities by imposing continuity using series-specific methods described below. We also winsorize some outcomes at the 99th percentile to reduce the influence of outliers when analyzing small cells.

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

Third, we take a series of steps to protect the confidentiality of businesses and their clients. Instead of reporting levels of each series, we report indexed values that show percentage changes relative to mean values in January 2020.[6] We suppress small cells and exclude outliers to meet privacy and data protection requirements, with thresholds that vary across datasets as described below. For data obtained from publicly traded firms – whose ability to disclose data is restricted by Securities and Exchange Commission regulations governing the disclosure of material non-public information – we combine data from multiple firms so that the statistics we report do not reveal information about any single company's activities.[7]

Finally, we address the challenge that our data sources capture information about the cus-

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

7. For publicly traded firms, our platform serves as a coordination device that allows multiple firms to pool and release their data even though each firm faces restrictions that limit its capacity to share its own data publicly.

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

We release each data series at the highest available frequency using an automated pipeline that ingests data from data providers, constructs the relevant statistics and conducts quality control tests, and outputs the series publicly (see Appendix A for details on the engineering of this pipeline). To limit revisions, we allow for a lag to account for reporting delays (typically one week). Wherever feasible, we disaggregate each series by industrial sector, geography (ZIP, county, or state), and income quartile.

In the rest of this section we describe each of the series in turn, discussing the raw data sources, construction of key variables, and cross-sectional comparisons to publicly available benchmarks. We benchmark trends in each series over time to publicly available data in the context of our analysis in Section III. All of the data series described below can be freely downloaded from the Economic Tracker website.

II.A Consumer Spending: Affinity Solutions and CoinOut

Affinity Solutions. We measure consumer spending primarily using aggregated and anonymized data on credit and debit card spending collected by Affinity Solutions Inc, a company that aggregates consumer credit and debit card spending information to support a variety of financial service products, such as loyalty programs for banks. Affinity Solutions captures nearly 10% of debit and credit card spending in the U.S. We obtain raw data from Affinity Solutions disaggregated by county, ZIP code income quartile, industry and day starting from January 1, 2019. We first remove discontinuous breaks caused by entry or exit of card providers from the sample (see Appendix B for details). We then construct daily values of the consumer spending series using a seven-day moving average of the current day and the previous six days of spending.

Because spending exhibits very large seasonal fluctuations (Appendix Figure 1a), we seasonally adjust our spending series by dividing each week's 2020 value by its corresponding value from 2019. We then index the seasonally-adjusted series relative to pre-COVID-19 spending by dividing each

value by the mean of the seasonally-adjusted average spending level in the first four complete weeks of 2020 (January 4-31).

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

Comparison to QSS and 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 2 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 potentially 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 below.

II.B Small Business Revenue: Womply

We obtain data on small business transactions and revenues from Womply, a company that aggregates data from several credit card processors to provide analytical insights to small businesses and other clients. In contrast to the Affinity series on consumer spending, which is a cardholder-based panel covering total spending, Womply is a firm-based panel covering total revenues of small busi-

nesses. The key distinction is that location in Womply data measures the location of the business as opposed to the location where the cardholder lives.

We obtain raw data on small business transactions and revenues from Womply at the ZIP-industry-day level starting from January 1, 2019. After excluding outliers, we aggregate these raw data to form two publicly available series at the county by industry level: one measuring total small business revenue and another measuring the number of small businesses open (see Appendix C for details). For each series, we construct seasonally-adjusted daily values using similar methods to those used to construct the consumer spending series.

Comparison to QSS and MARTS. Appendix Figure 2 shows the distribution of revenues observed in Womply across industries in comparison to national benchmarks. Womply revenues are again broadly distributed across sectors, particularly those where card use is common. A larger share of the Womply revenue data come from industries that have a larger share of small businesses, such as food services, professional services, and other services, as one would expect given that the Womply data only cover small businesses.

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

We combine several data sources to obtain information on employment and earnings: payroll data from Paychex and Intuit, worker-level data from Earnin, and time sheet data from Kronos. We describe each of these data sources in turn and then discuss how we combine them to construct a weekly series that measures private non-farm employment rates in the U.S.[8] Further details are provided in Appendix D.

Paychex. Paychex provides payroll services to approximately 670,000 small- and medium-sized businesses across the United States and pays nearly 10% of U.S. private-sector workers (Paychex 2020). We obtain aggregated weekly data on total employment and payroll earnings for each county by industry (two-digit NAICS) by 2019 hourly wage quartile by 2019 firm size bin by pay frequency. Hourly wage quartiles are based on fixed thresholds of the hourly wage distribution in 2019 (<\$13.00, \$13.00-\$18.18, \$18.18-29.17, >\$29.17). Firm size is measured in Dun & Bradstreet data on employment, broken into a set of broad groups (e.g., 1-49 employees, 50-99 employees, 100-199 employees, ..., 900-999 employees, >999 employees). We obtain data from Paychex on checks processed by week in each of these groups. We convert these data into an employment series using methods analogous to those developed by Cajner et al. (2020); see Appendix D for details.

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

Intuit. Intuit offers payroll services to businesses as part of its Quickbooks program, covering approximately one million businesses as of January 2020. Businesses that use Quickbooks tend to be very small (fewer than 20 employees). 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. Earnin is a financial management application that provides its members with access to their income as they earn it, in advance of their paychecks. Workers sign up for Earnin individually using a cell phone app, which connects to the bank account in which paychecks are deposited. Earnin complements the firm-based payroll datasets discussed above by providing a *worker-level* sample. This yields insight into employment rates at a much wider spectrum of firms – ranging from the largest firms in the U.S. to small businesses – at the expense of having fewer workers per firm. Since employment decisions are highly correlated across workers within firms at business cycle frequencies, Earnin's coverage of a large set of firms proves to be a valuable complement to the firm-based payroll datasets for our analysis. Because its users tend to have lower income levels, Earnin primarily provides information on employment for low-wage workers.[9]

We obtain anonymized data from Earnin from January 2020 to present at the paycheck level with information on home ZIP code, workplace ZIP code, unemployment status, wages, and earnings. We assign firm sizes to Earnin employers by matching to firm size data obtained from ReferenceUSA, and we assign NAICS codes using a custom-built crosswalk constructed by Digital Divide Data. We convert these data into an employment series using an approach similar to that used to construct the Paychex employment series.

Kronos. Kronos is a workforce management service used by many firms across the U.S. The data we obtain from Kronos cover approximately 30,000 mid-sized firms which together employed about 3.2 million workers pre-COVID. We obtain anonymized and aggregated weekly data on the total number of "punches," representing an employee clocking into work on an electronic time sheet. We obtain these data by county, industry, and firm size at the point that the firm entered the Kronos database. The employees in the database are primarily hourly workers who must record their time, and are concentrated in the bottom quartile of the wage distribution: assuming full-time

9. The median worker in the Earnin sample in January 2020 has an hourly wage rate of $10.80, which falls at the 16th 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 $8.03-$14.33 (corresponding to the 3rd and 40th percentiles of the national distribution).

employment, their wage rates translate to average earnings of $24,000 per year (with a 10th-90th percentile range of $7,200 to $45,600).

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

Homebase. Homebase provides scheduling tools for small businesses (on average, 8.4 employees) such as restaurants (which comprise 64% of employees for whom sectoral data are available). We obtain de-identified individual-level data on days worked and construct an employment series at the industry level. We include Homebase as a point of comparison because it has been widely used in other studies of small business employment in the COVID pandemic, but we do not include it in our primary employment indices because it does not track benchmarks of overall employment at small businesses as closely as our other data sources (see Section III.C below).

Combined Employment Series. To protect business privacy and maximize precision, we combine the data sources above to construct our primary employment series (see Appendix D for details). Because Paychex covers all sectors and wage levels, we use Paychex data disaggregated by industry, wage quartile and geographic area (county, state and national) as the base for the combined employment series. We then use Earnin and Intuit to refine the series in cells represented by those datasets. As noted above, Earnin best represents lower-wage workers. We therefore combine Earnin data with Paychex data to construct employment estimates for the bottom wage quartile. Next, we combine Intuit with the Paychex-Earnin data, accounting for the fact that Intuit data are available at a lower frequency and are not broken down by wage level or industry. We report seven-day moving averages of these series, expressed as a percentage change relative to the first four complete weeks of 2020 (January 4-31). We do not seasonally adjust our employment series because we have incomplete data in 2019; fortunately, seasonal fluctuations in employment are an order-of-magnitude smaller than those in spending (Appendix Figure 1b) and hence are unlikely to materially affect our results.

The employment series constructed based on payroll data is generally available only with a one month lag because people are paid after completing work over multiple prior weeks. To obtain more current estimates, we use Kronos time sheet data along with Paychex data from firms with

weekly paycycles to forecast employment rates (see Appendix D for the forecasting model).

Comparisons to QCEW and OES. Appendix Table 1 compares industry shares in each of the data sources above to nationally representative statistics from the Quarterly Census of Employment and Wages (QCEW). The Paychex-Earnin combined sample is broadly representative of the industry mix in the U.S. Intuit is concentrated primarily in professional services, construction, other services, and health care and social assistance. Kronos has fairly broad coverage, but over-represents the manufacturing and transportation and warehousing sectors. Homebase covers primarily food services.

In Appendix Table 2, we assess how these datasets compare to national benchmarks in terms of wage rates by comparing median wage rates to nationally representative statistics from the BLS's Occupational Employment Statistics. Median wage rates in Paychex-Earnin combined data closely match the OES estimates. Average wages in Intuit closely mirror OES estimates in the industries that Intuit covers.

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

II.D Job Postings: Burning Glass

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

We receive raw data from Burning Glass on job postings disaggregated by industry, week, job qualifications and county. Industry is defined using select NAICS supersectors, aggregated from 2-digit NAICS classification codes. Job qualifications are defined using ONET job zones, which classify jobs into five groups based on the amount of preparation they require. We also obtain analogous data broken down by educational requirements. We report job postings at the weekly level, expressed as changes in percentage terms relative to the first four complete weeks of 2020.

Comparison to JOLTS. Burning Glass data have been used extensively in prior research in economics; for instance, see Hershbein and Kahn (2018) and Deming and Kahn (2018). Carnevale, Jayasundera, and Repnikov (2014) show that the Burning Glass data are reasonably well aligned

with government survey-based statistics on job openings and characterize the sample in detail. In Appendix Figure 3, we compare the distribution of industries in the Burning Glass data to nationally representative statistics from the Bureau of Labor Statistics' Job Openings and Labor Market Turnover Survey (JOLTS) in January 2020. In general, Burning Glass is well aligned across industries with JOLTS, with the one exception that it under-covers government jobs. We therefore view Burning Glass as a sample representative of private sector jobs in the U.S.

II.E Education: Zearn

Zearn is a non-profit math curriculum publisher that combines in-person instruction with digital lessons. Zearn was used by approximately 925,000 students in the U.S. in Spring 2020. Many schools continued to use Zearn as part of their math curriculum after COVID-19 induced schools to shift to remote learning. We obtain data on the number of students using Zearn Math and the number of lessons they completed at the school-grade-week level (see Appendix E for details).

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

Comparison to American Community Survey. In Appendix Table 3, we assess the representativeness of the Zearn data by comparing the demographic characteristics of the schools for which we 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.

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

In addition to the new private sector data sources described above, we also collect and use three sets of data from public sources to supplement our analysis: data on unemployment benefit claims obtained from the Department of Labor and state government agencies; data on COVID-19 cases and deaths obtained from the New York Times; and data on the amount of time people spend at

home vs. other locations obtained from Google's COVID-19 Community Mobility Reports. Further details on these data sources are provided in Appendix F.

III Economic Impacts of COVID-19

To structure our analysis of the economic impacts of COVID-19, we begin from national accounts data. According to the Bureau of Economic Analysis (2020), GDP fell by $1.73 trillion (an annualized rate of 31.7%) from the first quarter of 2020 to the second quarter of 2020, shown by the first bar in Figure 1a. GDP fell primarily because of a reduction in personal consumption expenditures (consumer spending), which fell by $1.35 trillion. Government purchases and net exports did not change significantly, while private investment fell by $0.47 trillion.[10] We therefore begin our analysis by studying the determinants of this sharp reduction in consumer spending. We then turn to examine downstream impacts of the reduction in consumer spending on business activity and the labor market. We characterize the dynamics of these outcomes at a daily level from March to October 2020, the first eight months of the COVID pandemic in the U.S, focusing primarily on two critical periods: the trough of the recession (March 25-April 14) and the initial economic recovery (July 4-31), after which most outcomes remained relatively stable until October.

III.A Consumer Spending

We analyze consumer spending using data on aggregate credit and debit card spending. National accounts data show that spending that is well captured on credit and debit cards – essentially all spending excluding housing, healthcare, and motor vehicles – fell by approximately $1.03 trillion between the first quarter of 2020 and the second quarter of 2020, comprising roughly 75% of the total reduction in personal consumption expenditures.[11]

Benchmarking. We begin by assessing whether our Affinity data track patterns in corresponding spending categories in the national accounts. Figure 1b plots spending on retail services (excluding auto-related expenses) and food services in the Affinity data vs. corresponding series from the

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

11. The rest of the reduction is largely accounted for by healthcare expenditures; housing and motor vehicle expenditures did not change significantly. We view the incorporation of additional private sector data sources to study these other components of spending as an important direction for future work.

Advance Monthly Retail Trade Survey (MARTS), one of the main inputs used to construct the national accounts.[12] All series are expressed as a percentage change relative to January of each calendar year; each point shows the average level of daily spending in a given month divided by spending in January of that year. We do not seasonally adjust spending for this benchmarking analysis because seasonal fluctuations provide useful variation to assess whether the card spending series tracks the MARTS. The Affinity spending series tracks the MARTS closely both before and after the COVID crisis. In particular, both series show a rapid drop in food services spending in March and April 2020, while total retail spending fluctuates much less. The root-mean-squared-error of the Affinity series relative to the MARTS is small relative to the fluctuations induced by COVID.

Figure 1c plots the change in spending from January to April 2020 in the Affinity spending series against the decline in consumer spending as measured in the MARTS. Despite the fact that the MARTS category definitions are not perfectly aligned with those in the card spending data, the relative declines are generally well aligned across sectors, with a correlation of 0.88. Given that Affinity data track the MARTS at the national level quite well, we proceed to use it to disaggregate the national series in several ways to understand why consumer spending fell so sharply.[13]

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

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

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

application.[14]

Figure 2a plots a seven-day moving average of total daily card spending for households in the bottom vs. top quartile of ZIP codes based on median household income.[15] The solid line shows data from January to October 2020, while the dashed line shows data for the same days in 2019 as a reference. Spending fell sharply on March 15, when the National Emergency was declared and the threat of COVID became widely discussed in the United States. Spending fell from $7.9 billion per day in February to $5.6 billion per day by the end of March (a 30% reduction) for high-income households; the corresponding change for low-income households was $3.5 billion to $2.8 billion (a 20% reduction).

Because high-income households cut spending more in percentage terms and accounted for a larger share of aggregate spending to begin with, they accounted for a much larger share of the decline in total spending in the U.S. than low-income households. In Column 1 of Table 1a, we estimate that as of mid-April, top-quartile households accounted for roughly 41% of the aggregate spending decline after the COVID shock, while bottom-quartile households accounted for only 12% of the decline.

This gap in spending patterns by income grew even larger over time. By mid-July, spending had essentially returned to 2019 levels among households in the bottom quartile, whereas spending among high-income households remained 13% below baseline levels. These differences persisted through the first stage of economic recovery to the end of September. This heterogeneity in spending changes by income is larger than that observed in previous recessions (Petev, Pistaferri, and Eksten 2011, Figure 6) and plays a central role in the downstream impacts of COVID on businesses and the labor market, as we show below.

A potential concern with our card-based estimates of spending changes is bias from substitution out of cash purchases; for instance, if individuals sought to use more contactless methods to pay or began placing more orders online, trends in card spending might exhibit excess volatility relative to overall spending. To assess the importance of such substitution, we examine cash purchases using receipts data from CoinOut. Appendix Figure 4a plots aggregate cash purchases in the CoinOut

14. Cox et al. (2020) report an eight percentage point larger decline in spending for the highest income quartile relative to the lowest income quartile in the second week of April. Our estimate of the gap is also eight percentage points at that time, although the levels of the declines in our data are slightly smaller in magnitude for both groups.

15. We estimate total card spending for the U.S. population by inflating the daily spending in the Affinity Solutions data, multiplying by the ratio of January 2020 (or 2019) total spending for components of PCE that are likely captured in credit/debit card spending (shown in the last bar of Figure 1a) to the January 2020 (or 2019) total spending in the Affinity data.

data vs. aggregate card spending at grocery stores over time.[16] The time trends are very similar between the two series (with a signal correlation of 0.9 at the weekly level), showing a sharp spike in spending in late March (as households stocked up on groceries), followed by a more sustained increase in spending from the latter half of April. These results – together with the fact that our card spending series closely track estimates from the MARTS (Figures 1b and 1c) – indicate that aggregate fluctuations in card spending do not appear to have been offset by opposite-signed changes in cash spending. Rather, households shifted spending similarly across both modes of payment. We therefore proceed to focus on card spending in the rest of our analysis given the larger sample sizes and greater granularity of the card spending data.

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

The left bar in Figure 2b plots the share of the total decline in spending from the pre-COVID period to mid-April accounted for by various categories. In this and all subsequent figures analyzing consumer spending, we seasonally adjust spending levels by comparing them to 2019 levels in the same week (see Appendix B for details). Consistent with the findings of Cox et al. (2020), roughly two-thirds of the reduction in spending came from reduced spending on goods or services that require in-person contact (and thereby carry a risk of COVID infection), such as hotels, transportation, and food services. This is particularly striking given that these goods accounted for only one-third of total spending in January, as shown by the right bar in Figure 2b. Panel B of Table 1 shows that these gaps only grew larger as the pandemic progressed, as consumer spending increased above pre-pandemic levels for durable and non-durable goods by mid-July, but remained sharply depressed for in-person services.

Next, we zoom in to specific subcategories of spending that differ sharply in the degree to which they require physical interaction in Figure 2c. Seasonally-adjusted spending on luxury goods such as installation of home pools and landscaping services – which do not require in-person contact – *increased* slightly after the COVID shock. In contrast, spending on restaurants, beauty shops, and airlines all plummeted sharply. Consistent with these substitution patterns, online spending increased sharply: online purchases increased by 37% from the first to the second quarter of 2020

16. 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.

(U.S. Department of Commerce 2020). A conventional reduction in income or wealth would typically reduce spending on all goods as predicted by their Engel curves (income elasticities); the fact that the spending reductions vary so sharply across goods in line with their health risks lends further support to the hypothesis that health concerns (either one's own health or altruistic concerns about others' health) rather than a lack of purchasing power drove spending reductions.

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

Heterogeneity by COVID Incidence. To further evaluate the role of health concerns, we next turn to directly examine the association between incidence of COVID across areas and changes in spending from a baseline period prior to the COVID shock (January 4 to January 31, 2020) to the trough in consumer spending immediately after the COVID shock (March 25 to April 14, 2020). Figure 3a presents a binned scatter plot of changes in spending vs. the rate of detected COVID cases by county, separately for low- and high-income counties (median household income in the bottom vs. top income quartile). Spending fell more in counties with higher rates of COVID infection in both low- and high-income areas.[17]

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

Figures 3a-b show that at all levels of COVID infection, higher-income households reduced spending and time spent outside more than lower-income households. Figure 3c establishes this point more directly by showing that change in time spent outside home falls monotonically with

17. The relationship shown in Figure 3a also holds in a county-level regression of changes in consumer spending on the logarithm of cases per 100,000 people with controls for household income and state fixed effects (coefficient = -2.21, s.e. = 0.30). Note that there was a substantial reduction in spending even in areas without high rates of realized COVID infection, which is consistent with widespread concern about the disease even in areas where outbreaks did not actually occur at high rates.

household income across the distribution. These results help explain why the rich reduced spending more, especially on goods that require in-person interaction: high-income people apparently self-isolated more, perhaps by working remotely or because they had larger living spaces.

In sum, disaggregated data on consumer spending reveal that spending in the initial stages of the pandemic fell primarily because of health concerns rather than a loss of current or expected income. Income losses were relatively modest because few high-income individuals lost their jobs – as we show in Section III.C below – and lower-income households who experienced job loss had their incomes more than replaced by supplemental unemployment benefits (Ganong, Noel, and Vavra 2020), which led unemployed households to increase their spending relative to pre-COVID levels (Farrell et al., 2020a).[18] Indeed, national accounts data actually show an *increase* in total income of 13% from March to April 2020. This finding implies that the central channel emphasized in Keynesian models that have guided policy responses to prior recessions – a fall in aggregate demand due to a lack of purchasing power – was less important in the early stages of the pandemic, partly thanks to government policies such as increases in unemployment benefits and easing of monetary policy. Rather, the key driver of residual changes in aggregate spending was a contraction in firms' ability to supply certain goods, namely services that carry no health risks – consistent with the mechanisms emphasized by Eichenbaum, Rebelo, and Trabandt (2020).[19] We now show that this source of spending reductions leads to a novel set of downstream impacts on businesses and the labor market, potentially calling for different policy responses than in prior recessions.

III.B Business Revenues

How did reductions in consumer spending affect business activity – whether to remain open, how many employees to retain, what wage rates to pay them, how many new people to hire? To analyze

18. It may be surprising that we do not see a decline in aggregate spending in Figure 2a when supplemental unemployment benefits expired at the end of July 2020. Farrell et al. (2020b) estimate that the expiration of supplemental unemployment benefits led to a $96 reduction in weekly spending among unemployed households (16% of the reduction in unemployment benefits). As of the last week of July, federal spending on supplemental unemployment benefits was $15.6 billion (Morath 2020), implying that aggregate spending would decline by roughly $2.5 billion in the week following the expiration of supplemental benefits, about 1.6% of mean weekly card spending in January 2020. From July to September 2020, the root-mean-squared-error from a regression of total weekly consumer spending on a linear trend is 1.6%. Hence, the aggregate effect of the expiry of supplemental unemployment benefits is small relative to typical fluctuations around trend, explaining why the impacts of the expiration of UI benefits are not visible in the aggregate data shortly after benefits expire. However, as emphasized by Farrell et al. (2020b), households maintained consumption by depleting savings they had built up when receiving supplemental benefits; as they exhaust their assets, expenditure could fall more sharply and have a larger impact on aggregate spending.

19. This explanation may appear to be inconsistent with the fact that the Consumer Price Index (CPI) showed little increase in inflation during the pandemic, given that one would expect a supply shock to increase prices. However, the CPI likely understated inflation because it did not capture the extreme shifts in the consumption bundle (Cavallo 2020) or changes in product variety and promotions offered to consumers (Jaravel and O'Connell 2020) during the COVID crisis.

these impacts, we exploit variation in shocks to firms' revenues across ZIP codes. The motivation for this geographic approach is that spending fell primarily among high-income households in sectors that require in-person interaction, such as restaurants. Most of these goods are non-tradable products produced by small local businesses who serve customers in their local area.[20] We therefore use differences in average incomes and rents across ZIP codes as a source of variation in the magnitude of the spending shock that small businesses face.[21]

Benchmarking. We measure small business revenues using data from Womply, which records revenues from credit card transactions for small businesses (as defined by the Small Business Administration). Business revenues in Womply closely track patterns in the Affinity total spending data, especially in sectors with a large share of small businesses, such as food and accommodation services (Appendix Figure 5).[22]

Heterogeneity Across Areas. We begin our analysis of the Womply data by examining how small business revenues changed in low- vs. high-income ZIP codes from a baseline period prior to the COVID shock (January 4 to 31, 2020) to the period immediately after the COVID shock prior to the stimulus program (March 25 to April 14, 2020).[23] Throughout our analysis, we seasonally adjust small business revenues using data from 2019 (see Appendix C for details).

Figure 4 maps the change in small business revenue by ZIP code in three large metro areas: New York City, Chicago and San Francisco (analogous ZIP-level maps for other cities are posted online). There is substantial heterogeneity in revenue declines across areas. For example, within New York City average small business revenue declined by more than 73% in the hardest-hit decile of ZIP codes, compared to declines of less than 13% in the least affected decile of ZIP codes.[24]

In all three cities, revenue losses were largest in the most affluent parts of the city. For example, small businesses lost 67% of their revenue in the Upper East Side in New York, compared with 45%

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

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

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

23. We use 2010 Census ZIP Code Tabulation Areas (ZCTAs) to perform all geographic analyses of ZIP-level data. Throughout the text, we refer to these areas simply as "ZIP codes".

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

in the East Bronx; 66% in Lincoln Park vs. 41% in Bronzeville on the South Side of Chicago; and 84% in Nob Hill vs. 27% in Bayview in San Francisco. Revenue losses were also large in the central business districts in each city (lower Manhattan, the Loop in Chicago, the Financial District in San Francisco), likely a direct consequence of the fact that many high-income workers who used to work in these areas transitioned to working remotely. But even within predominantly residential areas, businesses located in more affluent neighborhoods suffered much larger revenue losses, consistent with the heterogeneity in spending reductions observed in the Affinity data.[25] More broadly, cities that have experienced the largest declines in small business revenue on average tend to be affluent cities – such as New York, San Francisco, and Boston (Appendix Table 4, Appendix Figure 7a).

Figure 5a generalizes these examples by presenting a binned scatter plot of percent changes in small business revenue vs. median household incomes, by ZIP code across the entire country. We observe larger reductions in revenue at small businesses located in more affluent ZIP codes across the distribution.[26] In the richest 5% of ZIP codes, small business revenues fell by 50%, as compared with around 35% in the poorest 5% of ZIP codes, consistent with the differences observed in the Affinity consumer spending data across areas.[27]

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

Figure 5c combines the income and population density mechanisms by plotting revenue changes vs. median rents (for a two bedroom apartment) by ZIP code. Rents are a simple measure of the affluence of an area that combine income and population density: the highest rent ZIP codes tend to be high-income, dense areas such as Manhattan. Figure 5c shows a steep gradient of revenue

25. We find a similar pattern when controlling for differences in industry mix across areas; for instance, the maps look very similar when we focus solely on small businesses in food and accommodation services (Appendix Figure 6).

26. One may be concerned that some of this correlation is driven by high-income households relocating to second homes, leading to shifting of small business revenue to other areas rather than a reduction in total revenue. While a significant number of households own second homes in certain areas (e.g., the most affluent parts of New York), the fraction of households moving to second homes is negligible in most areas, particularly those with rents below the 95th percentile of the distribution, and hence plays a small role in explaining these results.

27. Of course, households do not restrict their spending solely to businesses in their own ZIP code. An alternative way to establish this result at a broader geography is to relate small business revenue changes to differences in income distributions across counties. Counties with larger top 1% income shares experienced larger losses of small business revenue (Appendix Figure 8a). Poverty rates are not strongly associated with revenue losses at the county level (Appendix Figure 8b), showing that it is the presence of the rich in particular (as opposed to the middle class) that is most predictive of economic impacts on local businesses.

changes with respect to rents: revenues fell by less than 30% in the lowest-rent ZIP codes, compared with more than 50% in the highest-rent ZIP codes. This relationship is essentially unchanged when we compare ZIP codes within the same county by regressing revenue changes on rent with county fixed effects (Table 2, Column 2).[28] It also remains similar when controlling for the (pre-COVID) density of high-wage workers in a ZIP code to account for differences that may arise from shifts to remote work in business districts (Table 2, Column 3). Furthermore, these spatial differences persisted even as the economy began to recover. In July, small business revenue was still around 25% below its January level in the highest-rent ZIP codes, but only around 5% below the January level in the lowest-rent ZIP codes (Appendix Figure 9a).

In Figure 5d, we examine heterogeneity in this relationship across sectors that require different levels of physical interaction: food and accommodation services and retail trade (which largely require in-person interaction) vs. finance and professional services (which largely can be conducted remotely). Revenues fell much more sharply for food and retail in higher-rent areas; in contrast, there was essentially no relationship between rents and revenue changes for finance and professional services. These findings show that businesses that cater *in person* to the rich were those that lost the most revenue.

As a result of this sharp loss in revenues, small businesses in high-rent areas were much more likely to close entirely. We measure closure in the Womply data as reporting zero credit card revenue for three days in a row, and seasonally adjust rates of business closure using 2019 data (see Appendix B for details). Appendix Figure 8c shows that 40% of small businesses in the highest-rent ZIP codes closed, compared with 25% in the lowest-rent ZIP codes. The extensive margin of business closure accounts for most of the decline in total revenues.

Because businesses located in high-rent areas lost more revenue in percentage terms and accounted for a greater share of total revenue to begin with, they accounted for most of the aggregate loss in small business revenue. Almost half of the total loss in small business revenues came from businesses located in the top-quartile of ZIP codes by rent; less than 15% of the revenue loss came from businesses located in the bottom quartile. Next, we examine how the incidence of this shock is passed on to their employees.

28. The relationship between higher local rents and larger revenue losses at small businesses is not driven by variation between rural and urban areas, since this pattern is unchanged within counties.

III.C Employment Rates

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

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

Figure 6b examines how our series performs in matching national statistics on trends across sectors. For illustration, we focus on two sectors that experienced very different trajectories: food services, where employment fell heavily, and professional services, where it did not. In both cases, our Paychex-Intuit-Earnin series closely tracks data from the CES. Appendix Figure 10d shows more generally that changes in employment rates across sectors (two-digit NAICS) are very closely aligned in our series and the CES, with a correlation of 0.95.

For comparison, we also examine trends in employment based on data from Homebase, a dataset that has been used to examine employment trends in the COVID recession in many studies. Homebase exhibits a much larger decline in employment than the other series (56% at the trough vs. 15% in the CES). This is primarily because half of the individuals in the Homebase data work in the food services sector, which suffered particularly large employment losses as noted above; however, even within food services, Homebase exhibits a larger decline in employment at the trough (64%) relative to the CES (47%), as shown in Figure 6b, and thus has a higher RMSE. Because Homebase does not track overall national benchmarks on employment very closely, we do not use it for the analysis that follows, although we note that it exhibits qualitative patterns similar to the other series within food services.

In Appendix Figures 10a-b, we compare trends by wage quartile in our data with estimates based on the Current Population Survey and estimates in Cajner et al. (2020), who report employment changes by wage quintile using data from ADP in the initial weeks after the COVID shock. We find broadly similar trends in all three datasets. We also examine employment changes by state and find that in almost all states (excluding North Dakota and Hawaii), employment changes from

January-April in our combined series align very closely with changes in the CES, with an overall correlation of 0.99 (Appendix Figure 10c).

Based on these benchmarking exercises, we conclude that our combined employment series provides a good representation of employment rates across sectors, wage groups, and geographic areas.

Heterogeneity by Wage Rates. Figure 7a plots the combined employment series by wage quartile. To construct this figure, we first construct hourly wage quartiles based on fixed thresholds of the hourly wage distribution in 2019 (<\$13.00, \$13.00-\$18.18, \$18.18-29.17, >\$29.17). The solid lines plot total employment (based on repeated daily cross-sections) in each of these bins relative to the January baseline, based on the combined Paychex-Intuit-Earnin data. Consistent with the findings of Cajner et al. (2020) in prior work using ADP data, we find sharp heterogeneity in job losses by wage rate. Employment rates fell by 37% around the trough of the recession (April 15) for workers in the bottom wage quartile (i.e., the total number of jobs paying <\$13/hour was 37% lower as of April 15 than in January). By contrast, employment rates fell by only 14% for those in the top wage quartile as of April 15.

High-wage workers not only were less likely to lose their jobs to begin with, but also experienced a much more rapid recovery. By late May, employment for high-wage workers had returned nearly to the pre-COVID baseline. But employment rates for low-wage workers remained roughly 20% below baseline levels even as of mid-September. Using time sheet data from Kronos, and payroll data from firms with weekly paycycles in Paychex – both of which are available with a shorter lag than payroll-based employment data containing all paycycles – we construct a prediction of employment rates up to October 14 as described in Section II.C (shown by the dashed lines in Figure 7a). These predictions suggest that the rate of recovery remained slow for low-wage workers in October.

In sum, COVID induced a short-term "V-shaped" recession for high-wage workers in terms of employment opportunities, but led to a much deeper and more prolonged recession for lower-wage workers. Why did employment trajectories for low-wage workers differ so sharply from those for high-wage workers? One potential explanation is that low-wage workers work in different sectors or areas that may have experienced larger reductions in consumer demand. We evaluate this hypothesis in Figure 7b by plotting employment for workers in the bottom wage quartile, reweighting the series to match baseline employment shares by county and industry (2 digit NAICS) in the top wage quartile. This reweighting closes very little of the gap between the two series, showing that

differences in industry and location do not explain the differences in employment trajectories. Figure 7c provides a specific illustration of this result by showing trends in employment and spending in the retail trade sector. Total retail spending was nearly 10% higher as of July 31 relative to the pre-COVID baseline. Employment of high-wage workers was comparable to baseline levels, yet employment of low-wage workers was still down by over 15%.

One explanation for these patterns is that firms have shifted their production processes to use more technology or find other efficiencies (Lazear, Shaw, and Stanton 2016) and economic activity may have shifted toward these more efficient firms. For example, retail spending may have shifted toward online retailers and larger firms that use more capital (or imports) than low-wage labor in the United States to produce goods. Such shifts could reduce demand for routine occupations more permanently – a phenomenon documented in previous recessions by Jaimovich and Siu (2020). These results raise the possibility of a "jobless recovery" absent efforts to help workers who have been displaced from their prior jobs (Berger 2012).

Heterogeneity Across Areas. To shed further light on why employment rates for low-wage workers fell so much, we next turn to examine geographic heterogeneity in employment losses, in connection to the heterogeneity in spending changes and business revenue losses examined above. We begin by using the Earnin data – which is available at the ZIP code level – to analyze heterogeneity across the ZIP codes where people *work* (not necessarily where they live). Figure 8 maps changes in employment rates from January to the trough in mid-April for low-wage workers at small- and mid-size businesses (fewer than 500 employees) by ZIP code in New York, Chicago and San Francisco (analogous ZIP-level maps for other cities are posted online).[29] The patterns closely mirror those observed for business revenues above. Employment rates for low-wage workers fell by more than 70% in the most affluent areas of these cities, as compared with 30% in the least affluent areas. We observe very similar spatial patterns when examining variation across commuting zones (aggregates of counties) at the national level using the combined Paychex-Intuit-Earnin data (Appendix Figure 7b).

Figure 9a presents a binned scatter plot of changes in low-wage employment rates (from January to the mid-April trough) vs. median rents in the employer's ZIP code, by firm size. Employment rates fell much more at businesses located in high-rent areas than low-rent areas in all groups,

29. We focus on small and mid-size businesses here because, as noted above, larger firms are more likely to serve markets extending well beyond the ZIP code in which they are located. Additionally, very large firms exhibit significantly smaller declines in employment (Appendix Figure 11). Because employment fell more slowly after the initial shock to consumer spending and business revenue, we use a slightly later period – from April 8-28 – to analyze the trough in employment than the March 25-April 14 period we used to study changes in spending and revenue.

supporting the view that the sharp reductions in business revenue in affluent areas induced firms to lay off low-wage workers. As employment recovered in the months after April 2020, the spatial differences observed in Figure 9 persisted. In July, low-wage employment was approximately 15% below baseline levels in low-rent ZIP codes, but remained 30% below baseline in the highest-rent ZIP codes (Appendix Figure 9b).

We observe a similar gradient with respect to local rents for workers at very large firms: from 25% in the lowest-rent ZIPs to over 35% in the highest-rent ZIPs. This presumably reflects that fact that multi-establishment firms such as Starbucks faced larger revenue losses at stores located in more affluent neighborhoods for the reasons documented above, which in turns induced them to reduce employment in those areas more heavily. While there is a similar gradient with respect to rent levels, the overall level of employment losses for workers at large firms was lower than at smaller firms. This may be because large firms lost less revenue as a result of the COVID shock given their line of business (e.g., fast food vs. sit-down restaurants) or had a greater ability to substitute to other modes of business (delivery, online retail).[30]

Table 2b presents a set of regression estimates quantifying these impacts. Low-wage workers consistently faced larger employment losses in higher-rent ZIP codes, even when comparing ZIP codes within the same county (Column 2) and controlling for the density of high-wage workers in the ZIP code (Column 3).[31]

Job Postings. Prior work suggests that the labor market impacts of the recession may depend as much upon job postings as they do on the rate of initial layoffs (e.g., Diamond and Blanchard 1989, Elsby, Michaels, and Ratner 2015). We therefore now turn to examine how the spending shocks and revenue losses have affected job postings using data from Burning Glass, building on prior work by Forsythe et al. (2020). We conduct this analysis at the county level, pooling firms of all sizes and sectors because workers can substitute across firms and areas when searching for a new job, making it less relevant which exact firm or ZIP code they work in.

Figure 9b presents a binned scatter plot of the change in job postings between January and the April trough vs. median rents by county for jobs that require minimal education. We find a pattern similar to what we find with current employment: job postings with minimal educational

30. We cannot measure changes in revenue by establishment for large firms because the Womply data on revenues only cover small businesses.

31. Employment in the combined Paychex-Intuit-Earnin data is only available at the county level. We see a similar pattern when regressing Paychex-Intuit-Earnin employment changes for low-wage workers on rent at the county level (coefficient = -8.69% per thousand dollars, s.e. = 1.09% per thousand dollars), although the magnitude of the gradient is attenuated as expected given the coarser geographic measures.

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

Comparison to Great Recession. The geographic pattern of job losses in the COVID recession differed sharply from that in the Great Recession. Using employment data from the Bureau of Labor Statistics, Figure 10 shows that counties in the bottom quartile of the household median income distribution accounted for 30% of job losses in the Great Recession (from 2007-2010), while those in the top quartile accounted for less than 20% of job losses. By contrast, in the COVID recession (from January to April 2020), counties in the top quartile accounted for a larger share of job losses than counties in the bottom quartile.[32]

Because job losses were concentrated in more affluent areas, UI claims were almost equally likely to come from high- and low-income counties in the COVID recession. For example, 16% of the labor force in Santa Clara county in California – the highest income county on the West Coast – claimed UI between March 15 and May 2. This claim rate is identical to that in Fresno CA, a low-income county in California's Central Valley. Unemployment rates reached 10% regularly in Fresno in prior recessions, but such rates are unprecedented in Santa Clara.

In the Great Recession, the areas of the country that experienced the largest increases in unemployment took many years to recover because workers did not move to find new jobs, and job vacancies remained depressed in hard-hit areas well after the national recession ended (Yagan 2019). Appendix Figure 9c shows early signs of a similar pattern in this recession: job postings remained significantly lower in high-rent counties than in low-rent counties even at the end of July. These findings suggest that the recovery for low-wage workers may take longer in richer areas in the COVID recession.

32. 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.

III.D Impacts on Displaced Workers

The analysis in the last section focused on changes in employment rates for workers by location in repeated cross-sections, showing for example that low-wage employment rates fell sharply in Manhattan relative to the Bronx. We close our analysis by examining how the COVID shock affected workers' employment and consumption trajectories longitudinally, tracking workers who lost their jobs and examining whether they were able to find employment elsewhere as the economy recovered. Did people who worked in Manhattan pre-COVID remain out of work and have lower spending levels than comparable workers in the Bronx (who were less likely to lose their jobs)? Or did they find jobs elsewhere, so that the initial incidence of the shock was effectively shared across workers over time, as one would expect in a frictionless labor market?

The traditional approach to studying displacement effects is to compare the employment trajectories of displaced vs. comparable non-displaced workers in individual-level panel data (e.g., Jacobson, LaLonde, and Sullivan 1993, Sullivan and Wachter 2009). Here, we implement such an analysis using the aggregated employment data from Earnin that we have made publicly available. To do so, we track employment rates by workers' ZIP code of residence; since relatively few people moved during the pandemic (Pew Research Center 2020), this effectively provides a panel of employment rates for a given set of workers. We then use data from the Census LEHD Origin-Destination Employment Statistics (LODES) database, which provides information on the matrix of residential ZIP by work ZIP for low-income workers in the U.S. in 2017, to compute the average workplace median rent level for each residential ZIP.

Figure 11a presents a binned scatter plot of changes in low-income employment by *home* (residential) ZIP code vs. average *workplace* rent. This figure shows that low-income individuals who were working in high-rent areas pre-COVID were much less likely to be employed after the shock hit in April – consistent with our findings above. Strikingly, these employment losses then persisted over time with no convergence over time in employment rates for workers who lived in different areas. Figure 11b plots employment trends for low-wage workers living in low-income ZIP codes, by the average workplace rent. As of April 15, employment fell for workers in high-workplace-rent (top quartile) ZIP codes by 45%, whereas employment fell by 32% among workers in low workplace rent (bottom quartile) ZIP codes. This gap in employment rates persisted for several months; in mid-July, employment in high workplace rent ZIP codes was around 33% below baseline levels, whereas employment in low workplace rent ZIP codes was 17% below baseline levels. More broadly, across

the entire distribution of workplace rents, the slope of the relationship between employment rates and workplace rents was almost the same in July as in April (Appendix Figure 9d). These gaps in employment persist even when comparing workers who live in the same county: adding county fixed effects to a regression of changes in employment (as of July) on ZIP workplace increases the slope coefficient from -15%/$1000 (s.e. = 0.7%) to -33.3%/$1000 (s.e. = 3.6%).[33]

These results show that the spatial patterns in the maps in Figure 8 are driven not by people switching from working in high-rent areas to low-rent areas, but rather by persistently lower employment rates for those who happened to be working in high-rent areas pre-COVID. A worker working at a restaurant in Manhattan remains less likely to be employed months later than a worker in the Bronx. One explanation for this lack of re-equilibration even within labor markets is that, unlike in prior recoveries, many workers may have simply returned to their previous jobs as the economy recovered rather than shifting to new jobs. Given the high levels of churn typically observed in low-wage labor markets, one would expect greater convergence in employment rates across workers with different initial conditions over time. Still, this evidence suggests that there are significant frictions within labor markets that contributed to persistent impacts of COVID on low-wage workers who happened to work in areas and sectors where demand fell sharply.

Finally, we analyze how these differential shocks to employment affected the consumption of displaced low-wage workers. Returning to the card spending data from Affinity Solutions, we ask whether low-income individuals working in high-rent ZIP codes reduce spending more than those working in low-rent ZIP codes. Figure 11c replicates Figure 11a using spending changes on the y axis, restricting to households living in low-income ZIPs.[34] Low-income individuals living in areas where people tend to work in high-rent ZIP codes cut spending by 40% on average from January to April 2020, compared with 25% for those living in areas where people tend to work in low-rent ZIPs. The relationship remains similar when we compare ZIP codes within the same county and control for rents in the home (residential) ZIP code (Appendix Table 5).

In sum, reductions in spending by high-income households due to concerns about COVID infection led to persistent negative impacts on employment and spending for low-income workers

33. In the Earnin microdata, we find similar results even when comparing workers employed at the *same* firm (e.g., a chain restaurant): people working in high-rent ZIP codes in January remained less likely to have a job (anywhere) in July than their co-workers working in a different establishment of the same firm in lower-rent ZIP codes.

34. We restrict this figure to households living in low-income ZIPs because we cannot disaggregate the Affinity data by individual-level income. Since the Earnin data already represent only low-income workers, we do not restrict to low-income ZIPs in the employment analysis above; however, the patterns are very similar when restricting to low-income ZIPs in the Earnin data.

working in affluent areas.[35]

IV Evaluation of Policy Responses to COVID-19

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

IV.A State-Ordered Reopenings

One direct approach to changing consumer spending and employment is via executive orders. Many states enacted stay-at-home orders and shutdowns of businesses in an effort to limit the spread of COVID infection and later reopened their economies by removing these restrictions. We begin by examining how such executive orders affect economic activity, exploiting variation across states in the timing of shutdowns and reopenings. Throughout this section, we define the reopening date to be the day that a state *began* the reopening process (see Appendix G 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 (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, while New Mexico

35. In future work, this source of geographic variation in displacement rates could potentially be used to study further downstream impacts on a range of other financial and economic outcomes, as in Mian and Sufi (2009).

did not reopen until two weeks later, on May 16. Figure 12a 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.

Figure 12b generalizes the case study in Figure 12a by studying partial reopenings in the five states that issued such orders on or before April 27. For each reopening date (of which there are three: April 20, 24, and 27), we compare the trajectory of spending in treated states to a group of control states that had not reopened at the time that the treated state reopened. We select at least three control states (listed in Appendix Table 6) for each of the reopening dates by matching on pre-period levels of spending (relative to January) during the three weeks prior to reopening. 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 Figure 12b.

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 1.43 percentage point increase in spending. This DD estimate also appears in Table 3, Column 1. Column 2 replicates that specification, but focuses on reopenings where we can go out three weeks after the event before control states begin to reopen; the DD estimate is unchanged with this wider window, at 1.37 percentage points. Figure 12c shows that we also find little impact of reopenings on employment (using the Paychex-Intuit-Earnin data). Finally, Figure 12d shows (using data from Womply) that there was a 3.27 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, and recalculate the variance. Figure 12e 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.[36] 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, Goolsbee and Syverson 2020, Lin and Meissner 2020, Villas-Boas et al. 2020).

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. Consistent with that evidence, spending fell sharply in most states *before* formal state closures (Appendix Figure 12). 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.[37]

IV.B Stimulus Payments to Households

The Coronavirus Aid, Relief, and Economic Security (CARES) Act made direct payments to nearly 160 million people, totaling $267 billion as of May 31, 2020. Individuals earning less than $75,000 received a stimulus payment of $1,200; married couples earning less than $150,000 received a payment of $2,400; and households received an additional $500 for each dependent they claimed. These payments were reduced at higher levels of income and phased out entirely for households with incomes above $99,000 (for single filers without children) or $198,000 (for married couples without children). IRS statistics show that 72% of stimulus payments made in April were direct-deposited on exactly April 15, 2020, while some households received payments on April 14 (Bureau of the Fiscal Service 2020). The goal of these stimulus payments was to increase consumer spending and restore employment. Was the stimulus effective in achieving these goals?

Impacts on Consumer Spending. We estimate the causal effect of the stimulus payments using a regression discontinuity design. Figures 13a-b plot raw daily card spending levels relative to January for low-income (bottom income quartile ZIP codes) and high-income households (top

36. 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).

37. 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.

income quartile ZIP codes) for the month of April. To reduce cyclical patterns, we residualize daily spending levels with respect to day-of-week and first-day-of-month fixed effects, estimated using data for the period January 1, 2019 to May 10, 2019.

Spending levels jumped sharply from April 13th to 15th. Fitting linear control functions to the points on either side of April 15, we estimate that spending levels rose discontinuously on April 15 by 25 pp for low-income households and 8 pp in high-income households.[38] Both effects are statistically significantly different from 0, as well as from each other.[39] Panel A of Table 4 shows that these regression discontinuity estimates remain similar under varying bandwidths. These findings are consistent with contemporaneous work by Baker et al. (2020) and Karger and Rajan (2020), who use individual transaction data on incomes and spending patterns of approximately 15,000 primarily low-income individuals to estimate a large and immediate effect of receiving a stimulus check on spending, especially among low-income households. Together, these results provide empirical support for models that generate excess sensitivity of consumption to anticipated temporary income shocks (e.g., Campbell and Mankiw 1989, Kaplan and Violante 2014).

In Figures 13c-d, we investigate the composition of goods on which households spent their stimulus checks. We pool all households in these figures to maximize precision. Spending on durable goods rose by 21 percentage points following the arrival of the stimulus payments and further increased thereafter, rising well above pre-crisis levels. But spending on in-person services rose by only 7 percentage points, remaining more than 50% below pre-crisis levels. Durable goods accounted for 44% of the recovery in spending levels from the beginning to the end of April, while in-person services accounted for just 18% of the recovery (Appendix Figure 13). The stimulus thus increased the overall level of spending, but did not channel money back to the businesses that lost the most revenue due to the COVID shock. These findings provide empirical evidence for the "broken Keynesian cross" mechanism established in the Guerrieri et al. (2020)'s model, where funds are not recirculated back to the sectors shut down by the pandemic, diminishing multiplier effects.

Impacts on Business Revenue Across Areas. Next, we investigate how the stimulus affected business revenues. In particular, did the businesses that lost the most revenue – those in high-rent areas – gain business as as result of the stimulus? Figures 14a and 14b replicate the analysis above

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

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

using Womply data on small business revenues as the outcome, separately for lowest-rent-quartile and highest-rent-quartile ZIP codes. We see a sharp increase of 18 percentage points in revenues in small businesses in low-rent neighborhoods exactly at the time when households received stimulus payments. In contrast, Panel B shows a small, statistically insignificant increase in revenues of 1 percentage point for small businesses in high-rent areas.

This geographic heterogeneity illustrates another important dimension in which the stimulus did not channel money back to the business that lost the most revenue from the COVID shock. In fact, the stimulus actually *amplified* the difference in small business revenue losses rather than narrowing it across areas. Those in low-rent areas have nearly returned to pre-crisis levels following the stimulus payments, while those in high-rent areas remained more than 20% down relative to January levels (Figure 14c, blue lines). Panel B of Table 4 shows these regression discontinuity estimates remain similar under varying bandwidths.

Impacts on Employment. Finally, we investigate whether the increase in spending induced by the stimulus increased employment rates, as one would expect in a traditional Keynesian stimulus. Here, we do not use the RD design as we do not expect employment to respond immediately to increased spending. Instead, we analyze the evolution of employment of low-wage workers in the Earnin data in low- vs. high-rent ZIP codes over time in Figure 14c (orange lines). In high-rent areas, low-wage employment remained 45% below pre-COVID levels as of the end of April – perhaps not surprisingly, since revenues have not recovered significantly there. But even in low rent areas, employment has recovered only partially, despite the fact that small business revenues have reverted to pre-COVID baseline levels. This result echoes the divergence between employment and revenue at the sectoral level documented above in Figure 7c and again raises the specter of a jobless recovery for low-wage workers.

In summary, the stimulus substantially increased total consumer spending but did not directly undo the initial spending reductions by returning money back to the businesses that lost the most revenue. This empirical impact contrasts with theoretical motivations for stimulus in response to shocks. In particular, Farhi and Werning (2016) show that optimal macroprudential policy involves a stimulus that increases spending in sectors and areas whose demand is depressed. In a frictionless model where businesses and workers could costlessly reallocate their capital and labor to other sectors, the reallocation of spending across sectors and areas might have no consequence for employment levels. But if workers' ability to switch jobs is constrained – e.g., because of job-specific skills that limit switching across industries or costs that limit moving across geographic

areas, as suggested by Yagan (2019) – the ability of the stimulus to foster a uniform recovery in employment to pre-COVID levels is likely to be hampered, perhaps explaining why employment levels remained well below baseline even as total spending recovered after April 15.

IV.C Loans to Small Businesses

We now evaluate the Paycheck Protection Program (PPP), which 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, followed by another $175 billion in a second round beginning on April 27. The program offered loan forgiveness for businesses that maintained sufficiently high employment (relative to pre-crisis levels).

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 (2020a) 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, 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.[40]

We estimate the causal effect of the PPP on employment rates at small businesses using a difference-in-differences research design, comparing trends in employment for firms below the 500 employee cutoff (the treated group) vs. those above the 500 employee cutoff (the control group) before vs. after April 3, when the PPP program began.[41] Figure 15a plots the average change in employment rates (inferred from payroll deposits) relative to January for firms in the Paychex-Earnin data employing 100-500 employees, which were eligible for PPP loans, vs. firms employing

40. 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.

41. 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.

501-800 employees, which were generally ineligible for PPP loans.[42] 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 industries in the below-500 and above-500 employee groups match the overall distribution of industries in January 2020 when computing mean employment rates by firm size. We also residualize employment rates by county x wage quartile x week fixed effects, to account for the differential time patterns of employment rates by county and wage quartile shown in Section III.C.

Before April 3, trends in employment 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 continued to follow a similar trajectory in the treated (<500 employees) and control (>500 employees) groups, with at most a relative increase of 2 pp in the treated group until August, 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.

Figure 15b plots the change in employment from January 4-January 31 to June 1-June 23 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.[43] Appendix Figure 14 shows that we obtain very similar results in the Earnin data alone.[44]

In Table 5, we quantify the impacts of the PPP using OLS regressions of the form:

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

where Emp_{scqit} is the change in employment within each firm size category s × county c × wage quartile q × 2-digit NAICS industry i on week t cell, relative to January 4-January 31, Eligible_s is an indicator variable for whether firm size is 500 or fewer employees in the pre-COVID period, Post-PPP_t is an indicator variable for the date being on or after 3 April 2020, and α_{cqt} represents a county-week-wage quartile fixed effect. We estimate this regression on the sample of firms with 100-800 employees using data from March 11 to August 15. We focus on employment impacts up

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

43. 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.

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

to August 15 because Figure 15a 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 distributions of industries in the below-500 and above-500 employee groups both match the overall distribution of industries in January 2020. We cluster standard errors at the county-by-industry-by-firm-size level to permit correlation in errors across firms and over time within counties and estimate the regression using OLS, weighting by the total number of employees in the cell from January 4-31, 2020.

Column 1 presents the baseline estimate obtained from regression equation (1) of $\beta_{DD} = 1.78$ ($s.e. = 1.99$), an estimate that matches the figure plotted in Figure 15a and is similar to that obtained in confidential ADP data in contemporaneous work by Autor et al. (2020). The mean decline in employment among firms in the control group to August 15 was 18.7%, implying that the PPP saved 9.1% 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 estimates remains statistically indistinguishable from that in Column 1. Columns 3 and 4 replicate the specification in Column 1, using data from Earnin and Kronos, respectively. The estimates from these data sources (which apply to workers in the bottom wage quartile) are similar to our baseline estimate.

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 2018 from the ReferenceUSA database for the Earnin sample and employment levels in 2019 from Dun & Bradstreet data for the Paychex 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 estimate that our reduced-form estimates are attenuated by 35% based on this matched data (see Appendix D for details). Under standard assumptions required to obtain a local average treatment effect in the presence of non-compliance – no direct effect of being classified as having more than 500 workers independent of the PPP and a monotonic treatment effect – we can estimate the LATE of the PPP on employment rates by multiplying the raw estimates reported in Table 5 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 2.41 percentage points.

Costs Per Job Saved. Using SUSB data, we calculate that approximately 53.6 million workers work at firms eligible for PPP assistance, excluding firms in NAICS 72 (for details, see Appendix D). Under the assumption that the PPP's effects on firms with between 100 and 500 employees were the same in percentage terms as the PPP's effects on all eligible firms, our baseline estimates in the combined Paychex-Earnin data (Column 1 of Table 5), adjusted for attenuation bias, imply that the PPP saved $0.02 \times 53.6M = 1.29$ million jobs from April through August 15.[45] 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) of $377,000. Even at the upper bound of the 95% confidence interval for employment impact, we estimate a cost per job saved of $119,000. For comparison, mean annual earnings for workers at PPP-eligible firms are only $45,000.

Why did the PPP have relatively small effects on employment rates despite having a very high takeup rate among small businesses? One potential explanation is that the loans were taken by firms that did not intend to layoff many employees to begin with, i.e. firms that were inframarginal recipients of loans. Consistent with this hypothesis, Granja et al. (2020) show that states and congressional districts that experienced more job losses prior to April 3 actually received *fewer* PPP loans. Moreover, PPP loans also were not distributed to the industries most likely to experience job losses from the COVID crisis. For example, firms in the professional, scientific, and technical services industry received a greater share of the PPP loans than accommodation and food services (SBA 2020). Yet accommodation and food services accounted for half of the total decline in employment between February and March (prior to PPP enactment) in BLS statistics, while employment

45. 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 1.51 million jobs.

in professional, scientific and technical services accounted for less than 5% of the decline.

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

V Conclusion

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

We apply these new data to analyze the economic impacts of COVID-19. We find that COVID-19 induced high-income households to self-isolate and sharply reduce spending in sectors that require physical interaction. This spending shock in turn led to losses in business revenue and layoffs of low-income workers at firms that cater to high-income consumers. Because the root cause of the shock was self-isolation driven by health concerns, there was limited capacity to restore economic activity without addressing the virus itself, at least in the initial months after the pandemic began in mid-March. In particular, state-ordered reopenings of economies had only modest impacts on economic activity; stimulus checks increased spending particularly among low-income households, but very little of the additional spending flows to the businesses most affected by COVID; and loans to small businesses had little impact on employment rates. Our analysis therefore suggests that the most effective approach to mitigating economic hardship in the midst of a pandemic may be to provide benefits to those who have lost their incomes to mitigate consumption losses, while investing in public health measures to restore consumer confidence and ultimately increase spending.

We focused in this paper on the short-run economic consequences of COVID-19. However, such shocks can also have long-lasting scarring effects. Private sector data can be useful in measuring these impacts as well. To illustrate, Figure 16 plots weekly student progress (lessons completed) on Zearn, an online math platform used by nearly one million elementary school students as part of their regular school curriculum. Children in high-income areas temporarily learned less on this platform when the COVID crisis hit and schools shifted to remote instruction, but soon recovered

to baseline levels. By contrast, children in lower-income areas remained 50% below baseline levels through the end of the school year. Although this platform captures only one aspect of education, these findings raise the concern that pandemics may amplify inequality in the long-run by hindering human capital development especially for lower-income children.

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

The advantage of constructing a public database to conduct such analyses rather than working directly with the confidential data held by private sector firms is that it permits a much broader range of work along these lines. For example, the data are now being used by local policymakers to inform local policy responses and forecast tax revenue impacts (e.g., Maine, Missouri, Kansas, and Texas). They are also being used by Congressional staff to design federal policies, e.g. predicting the impacts and costs of policies targeted based on business revenue losses (RESTART Act 2020). And they are being used by other researchers to analyze a broad range of issues: constructing price indices that account for changes in consumption bundles (Cavallo 2020), analyzing the effects of political views on economic outcomes (Makridis and Hartley 2020), estimating the effects of the PPP on smaller firms' employment decisions (Granja et al. 2020), and estimating the impacts of unemployment benefits on aggregate spending (Casado et al. 2020).

Over the 20th century, the Bureau of Economic Analysis built on a prototype developed by Kuznets (1941) to institute surveys of businesses and households that form the basis for today's National Income and Product Accounts. The database built here can be viewed as a prototype for a system of granular, real time national accounts built using transactional private sector data. Even this first prototype yields insights that cannot be obtained from existing data, suggesting that a broader, more refined platform that aggregates data from private companies has great potential for improving our understanding of economic activity and policymaking going forward.[46]

46. We view such data as a complement to rather than replacement for survey-based national accounts. Data from representative surveys will remain essential to obtain consistent statistics on national aggregates and to benchmark private sector data sources; private sector data serve to increase granularity and precision.

References

Abraham, Katharine G, Ron S Jarmin, Brian Moyer, and Matthew D Shapiro (ed.) 2019. *Big Data for 21st Century Economic Statistics.* NBER Book Series Studies in Income / Wealth.

Aladangady, Aditya, Shifrah Aron-Dine, Wendy Dunn, Laura Feiveson, Paul Lengermann, and Claudia Sahm. 2019. "From Transactions Data to Economic Statistics: Constructing Real-time, High-frequency, Geographic Measures of Consumer Spending." *NBER Working Paper No. 26253* (September). https://doi.org/10.3386/w26253. http://www.nber.org/papers/w26253.

Alexander, Diane, and Ezra Karger. 2020. "Do stay-at-home orders cause people to stay at home? Effects of stay-at-home orders on consumer behavior." *Federal Reserve Bank of Chicago Working Paper No. 2020-12* (April). https://doi.org/10.21033/wp-2020-12. https://www.chicagofed.org/publications/working-papers/2020/2020-12.

Allen, Danielle, Sharon Block, Joshua Cohen, Peter Eckersley, and Meredith Rosenthal. 2020. "Roadmap to Pandemic Resilience: Massive Scale Testing, Tracing, and Supported Isolation (TTSI) as the Path to Pandemic Resilience for a Free Society." *Edmond J. Safra Center For Ethics At Harvard University.*

Altonji, Joseph, Zara Contractor, Lucas Finamor, Ryan Haygood, Ilse Lindenlaub, Costas Meghir, Cormac O'Dea, Dana Scott, Liana Wang, and Ebonya Washington. 2020. "Employment Effects of Unemployment Insurance Generosity During the Pandemic." *Yale University Manuscript.*

Angrist, Joshua D, Guido W Imbens, and Donald B Rubin. 1996. "Identification of causal effects using instrumental variables." *Journal of the American statistical Association* 91 (434): 444–455.

Austin, Benjamin A, Edward L Glaeser, and Lawrence H Summers. 2018. "Jobs for the Heartland: Place-based policies in 21st century America." *NBER Working Paper No. 24548* (April). https://www.nber.org/papers/w24548.

Autor, David, David Cho, Leland D Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz. 2020. *An Evaluation of the Paycheck Protection Program Using Administrative Payroll Microdata.* Technical report.

Baker, Scott R, R. A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine" Yannelis. 2020. "Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments." *NBER Working Paper No. 27097.*

Bartik, Alexander W., Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matt Unrath. 2020. "Measuring the labor market at the onset of the COVID-19 crisis." *Brookings Papers on Economic Activity* (June). https://www.brookings.edu/wp-content/uploads/2020/06/Bartik-et-al-conference-draft.pdf.

Bartlett, Robert P, and Adair Morse. 2020. "Small Business Survival Capabilities and Policy Effectiveness: Evidence from Oakland." *NBER Working Paper No. 27629.*

Bennet, M. 2020. "S 3814-RESTART Act." *Senate - Finance Committee.*

Berger, David. 2012. "Countercyclical restructuring and jobless recoveries."

Blanchard, Olivier, and Lawrence Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 1992 (1): 1–61.

Board of Governors of the Federal Reserve System. 2019. *The 2019 Federal Reserve Payments Study.*

Bureau of the Fiscal Service. 2020. *Daily Treasury Statement: Issues: Current and Archive.*

Cajner, Tomaz, Leland D Crane, Ryan A Decker, Adrian Hamins-Puertolas, and Christopher Kurz. 2019. *Improving the Accuracy of Economic Measurement with Multiple Data Sources: The Case of Payroll Employment Data.* Technical report. National Bureau of Economic Research.

Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. "The U.S. Labor Market during the Beginning of the Pandemic Recession." *Working Paper* (May).

Campbell, John Y, and N Gregory Mankiw. 1989. "Consumption, income, and interest rates: Reinterpreting the time series evidence." *NBER Macroeconomics Annual* 4:185–216.

Carnevale, Anthony P, Tamara Jayasundera, and Dmitri Repnikov. 2014. "Understanding online job ads data." *Georgetown University, Center on Education and the Workforce, Technical Report (April).*

Casado, Miguel Garza, Britta Glennon, Julia Lane, David McQuown, Daniel Rich, and Bruce A Weinberg. 2020. "The Effect of Fiscal Stimulus: Evidence from COVID-19."

Cavallo, Alberto. 2020. "Inflation with Covid Consumption Baskets."

Chen, Haiqiang, Wenlan Qian, and Qiang Wen. 2020. "The Impact of the COVID-19 Pandemic on Consumption: Learning from High Frequency Transaction Data." *Working Paper* (April). https://doi.org/http://dx.doi.org/10.2139/ssrn.3568574. https://ssrn.com/abstract=3568574.

Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *The Quarterly Journal of Economics.*

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *The Quarterly Journal of Economics* 129 (4): 1553–1623.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. *Labor Markets During the COVID-19 Crisis: A Preliminary View.* Technical report. National Bureau of Economic Research.

Cox, Natalie, Peter Ganong, Pascal Noel, Joseph Vavra, Arlene Wong, Diana Farrell, and Fiona Greig. 2020. "Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data." *Brookings Papers on Economic Activity.*

Deming, David, and Lisa B. Kahn. 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals." *Journal of Labor Economics* 36 (S1): S337–S369. https://doi.org/10.1086/694106. https://doi.org/10.1086/694106.

Diamond, Peter, and OJ Blanchard. 1989. "The beveridge curve." *Brookings Papers on Economic Activity* 1:1–76.

Dunn, Abe, Kyle Hood, and Alexander Driessen. 2020. "Measuring the Effects of the COVID-19 Pandemic on Consumer Spending Using Card Transaction Data." *National Bureau of Economic Research.*

Ehrlich, Gabriel, John Haltiwanger, Ron Jarmin, David Johnson, and Matthew D Shapiro. 2019. "Re-engineering Key National Economic Indicators." In *Big Data for 21st Century Economic Statistics.* University of Chicago Press.

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt. 2020. *The Macroeconomics of Epidemics.* Technical report. National Bureau of Economic Research.

Elsby, Michael WL, Ryan Michaels, and David Ratner. 2015. "The Beveridge curve: A survey." *Journal of Economic Literature* 53 (3): 571–630.

Farhi, Emmanuel, and Iván Werning. 2016. "A theory of macroprudential policies in the presence of nominal rigidities." *Econometrica* 84 (5): 1645–1704.

Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, Daniel Sullivan, and Joseph Vavra. 2020b. *The Unemployment Benefit Boost.* Technical report. JPMorgan Chase Institute.

Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, and Joseph Vavra. 2020a. *Consumption Effects of Unemployment Insurance during the COVID-19 Pandemic.* Technical report. JPMorgan Chase Institute.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer. 2015. "The next generation of the Penn World Table." *American Economic Review* 105 (10): 3150–82.

Forsythe, Eliza, Lisa B Kahn, Fabian Lange, and David Wiczer. 2020. "Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims." *Journal of Public Economics* 189:104238.

Ganong, Peter, Pascal J Noel, and Joseph S Vavra. 2020. *US Unemployment Insurance Replacement Rates During the Pandemic.* Technical report. National Bureau of Economic Research.

Gindelsky, Marina, Jeremy Moulton, and Scott A Wentland. 2019. "Valuing housing services in the era of big data: A user cost approach leveraging Zillow microdata." In *Big Data for 21st Century Economic Statistics.* University of Chicago Press.

Goolsbee, Austan, and Chad Syverson. 2020. *Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020.* Working Paper, Working Paper Series 27432. National Bureau of Economic Research, June. http://www.nber.org/papers/w27339.

Granja, João, Christos Makridis, Constantine Yannelis, and Eric Zwick. 2020. "Did the Paycheck Protection Program Hit the Target?" *NBER Working Paper No. 27095* (May). https://doi.org/10.3386/w27095. http://www.nber.org/papers/w27095.

Greene, Claire, and Joanna Stavins. 2020. *2019 Diary of Consumer Payment Choice.* Technical report. Federal Reserve Bank of Atlanta.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. *Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?* Working Paper, Working Paper Series 26918. National Bureau of Economic Research, April. https://doi.org/10.3386/w26918. http://www.nber.org/papers/w26918.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108, no. 7 (July): 1737–72. https://doi.org/10.1257/aer.20161570. https://www.aeaweb.org/articles?id=10.1257/aer.20161570.

House Committee on Small Business. 2020. *Oversight of the Small Business Administration and Department of Treasury Pandemic Programs: Hearing Before The House Committee on Small Business, 116th Cong. (Testimony of Steven Mnuchin).*

Hubbard, Glenn, and Michael R Strain. 2020. "Has the Paycheck Protection Program Succeeded?" *Brookings Institution.*

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993. "Earnings Losses of Displaced Workers." *The American Economic Review,* 685–709.

Jaimovich, Nir, and Henry E Siu. 2020. "Job polarization and jobless recoveries." *Review of Economics and Statistics* 102 (1): 129–147.

Jaravel, Xavier, and Martin O'Connell. 2020. "Real-time price indices: Inflation spike and falling product variety during the Great Lockdown." *Journal of Public Economics* 191.

Kaplan, Greg, and Giovanni L Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica* 82 (4): 1199–1239.

Karger, Ezra, and Aastha Rajan. 2020. "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments." *FRB of Chicago Working Paper.*

Kurmann, André, Etienne Lalé, and Lien Ta. 2020. "The Impact of COVID-19 on U.S. Employment and Hours: Real-Time Estimates with Homebase Data" (May). http://www.andrekurmann.com/hb_covid.

Kuznets, Simon. 1941. *National Income and Its Composition, 1919-1938.* New York: National Bureau of Economic Research.

Lazear, Edward P., Kathryn L. Shaw, and Christopher Stanton. 2016. "Making Do with Less: Working Harder during Recessions." *Journal of Labor Economics* 34 (S1): S333–S360. https://doi.org/10.1086/682406. eprint: https://doi.org/10.1086/682406. https://doi.org/10.1086/682406.

Lin, Zhixian, and Christopher M Meissner. 2020. "Health vs. Wealth? Public Health Policies and the Economy During Covid-19." *NBER Working Paper No. 27099* (May). https://doi.org/10.3386/w27099. http://www.nber.org/papers/w27099.

Makridis, Christos, and Jonathan Hartley. 2020. "The Cost of Covid-19: A Rough Estimate of the 2020 US GDP Impact."

Mathy, Gabriel. 2020. *The COVID-19 Epidemic will be the First Services Recession and it Could be a Bad One.*

Mian, Atif, and Amir Sufi. 2009. "The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis." *The Quarterly Journal of Economics* 124 (4): 1449–1496.

Morath, Eric. 2020. "Federal Government Sent Workers Nearly 250$Billionin$600-a-Week Jobless Aid." *The Wall Street Journal.*

Paychex. 2020. *Small Business Employment Watch.* https://www.paychex.com/employment-watch/#!/, April.

Petev, Ivaylo, Luigi Pistaferri, and Itay Saporta Eksten. 2011. *Consumption and the Great Recession: An analysis of trends, perceptions, and distributional effects.*

Pew Research Center. 2020. *2020 Pew Research Center's American Trends Panel Wave 68 June 2020 Final Topline.*

PwC. 2020. *PwC US CFO Pulse Survey.* Technical report. June. https://www.pwc.com/us/en/library/covid-19/pwc-covid-19-cfo-pulse-survey.html.

Romer, Paul. 2020. "Roadmap to Responsibly Reopen America." *roadmap.paulromer.net,* http://roadmap.paulromer.net.

Small Business Administration. 2020a. *Joint Statement by SBA Administrator Jovita Carranza and U.S. Treasury Secretary Steven T. Mnuchin Regarding Enactment of the Paycheck Protection Program Flexibility Act.*

———. 2020b. *Paycheck Protection Program (PPP) Report.* Technical report. May. https://www.sba.gov/document/report--paycheck-protection-program-ppp-report.

Sullivan, Daniel, and Till von Wachter. 2009. "Job Displacement and Mortality: An Analysis Using Administrative Data." *The Quarterly Journal of Economics* 124 (3): 1265–1306.

Summers, Robert, and Alan Heston. 1984. "Improved International Comparisons of Real Product and its Composition: 1950–1980." *Review of Income and Wealth* 30 (2): 207–219.

———. 1991. "The Penn World Table (Mark 5): an expanded set of international comparisons, 1950–1988." *The Quarterly Journal of Economics* 106 (2): 327–368.

U.S. Bureau of Economic Analysis. 2020. *National Income and Product Accounts.* Data retrieved from U.S. Bureau of Economic Analysis, National Income and Product Accounts.

U.S. Department of Commerce. 2020. *U.S. Census Bureau News (Tuesday, 18 August, 2020).* Technical report.

Villas-Boas, Sofia B, James Sears, Miguel Villas-Boas, and Vasco Villas-Boas. 2020. "Are We #StayingHome to Flatten the Curve?" *UC Berkeley: Department of Agricultural and Resource Economics CUDARE Working Papers* (April). https://escholarship.org/uc/item/5h97n884.

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

Online Appendix

A Automated Data Processing Pipeline

This appendix describes the automated pipeline we built to ingest raw data, process it to construct aggregate statistics, and then release those statistics publicly. This automated pipeline allows us to typically post updated statistics within one business day of receiving the raw data. By automating the data processing to the extent possible, we aim to post data as close to real-time as possible, while maintaining the quality of the data and minimizing the manual upkeep required. The primary source of lags in the posted data is therefore driven by lags in the underlying data generating processes: for example, card transactions can take up to a week to settle and employment income is typically paid in 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. In order to flexibly accommodate diverse data sources, with varying secure file transfer methods and update frequencies, we operate a server in the cloud that pulls updated data from each source on a regular interval. We receive data updates from private companies on a daily, weekly or monthly cadence. Many companies have unique policies and requirements for securing data transfers, so we write scripts to intake this data using a variety of secure file transfer services (e.g. Amazon S3 buckets and SFTP servers). We also download or scrape a variety of publicly available statistics from the web, such as unemployment insurance claims and COVID-19 case counts.

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

pointer in the Git repository.

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

This data processing step generates two outputs: (1) a set of CSV files that contain all the data to be posted publicly and (2) a quality control report. The quality control report is a document that allows analysts to quickly assess any notable deviations in the data and determine whether the updated data 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 (www.tracktherecovery.org). While doing so, we also update the "last updated" and "next expected update" dates on the website. Second, we upload the CSV files containing all the updated data to our "data downloads" page. The updated visualizations and data downloads are then both immediately available for public use.

B Consumer Spending Series Construction

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

Structure of Raw Data. We receive data from Affinity Solutions in cells corresponding to the intersection of (i) county by (ii) income quartile by (iii) industry by (iv) day and week, where 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. ZIP code income quartile and county are both determined by the cardholder's residence. We crosswalk from ZIP codes to counties using the geographic definitions described in Appendix G to aggregate the series to the county and state levels.

We adjust the raw data we receive from Affinity Solutions to address three challenges: (1) changes in the customer base over time, (2) a data quality issue which creates spurious increases in consumer spending, and (3) seasonal fluctuations.

Changing Customer 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 transacting cards at a weekly level on the date separately for each year-by-county, and then implementing a Supremum Wald test for a structural break at an unknown break point.

950 counties have a structural break where the p-value of the test is less than 5×10^{-8}. For counties with only one break below this threshold, we correct our estimates as follows. We first compute the state-level week-to-week percent change in spending excluding all counties with a structural break (using the national series for DC and states for which all counties have a structural break). If we identify a structural break in week t, we impute spending levels in weeks $t-1$, t, and $t+1$, 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). When there is a change in coverage we adjust the series to be in line with the lower level of coverage. 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 $(0.92) \times (1.155n) = 1.0626n$. At this point, state-level fluctuations no longer impact the series, and we use the reported percent change each week to adjust this number for card coverage. We omit 98 counties with multiple structural breaks from our series. We do not remove any counties where the structural break occurred between March 10th and March 31st of 2020 because the consumer spending response to the COVID-19 was so strong that in many places it could be classified as a structural break. For the week around Christmas, we impute the number of cards by averaging the preceding and succeeding week, since holiday spending spikes are also sometimes classified as a structural break.

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

Spurious Changes in Consumer Spending. There is an implausibly large spike in consumer spending between January 15th, 2019 and January 17th, 2019 that is not found in other data series. This spike in national consumer spending is not driven by specific regions nor sectors. We treat this spike as a data quality issue, and respond by replacing each impacted day with the average spending on $t-7$, $t+7$, and $t+14$, where t is the impacted day. A similar problem arises in the "Accommodations and Food Services" sector in Richmond City County, Virginia where spending increases by over 80 times on May 23rd, 2019 relative to to nearby days. We implement a similar procedure replacing the impacted day with the average spending on $t-14$, $t-7$, $t+7$, and $t+14$, where t is the impacted day.

Seasonal Adjustment. 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.

Definition of Categories of Goods. In parts of our analysis, we distinguish between four cat-

egories of goods and services: durable goods; non-durable goods; remote services; and in-person services. We define durable goods as the following MCC groups: motor vehicles, sporting and hobby goods, home improvement centers, consumer electronics, and telecommunications equipment. Non-durable goods include wholesale trade, agriculture, forestry and hunting, general merchandise, apparel and accessories, health and personal care stores, and grocery stores. Remote services include utilities, professional/scientific services, public administration, administration and waste services, information, construction, education, and finance and insurance. In-person services include real estate and leasing, recreation, health care services, transportation and warehousing services, and accommodation and food services, as well as barber shops, spas, and assorted other services.

C Small Business Revenue and Small Businesses Open Series Construction

This appendix details our methodology for constructing the Small Business Revenue and Small Businesses Open series, using data from Womply.

Initial Construction. We receive Womply data on total revenue and number of open businesses at the date x ZIP code x firm category level. The Womply data are limited to small businesses, defined as businesses with annual revenue below Small Business Administration thresholds.

To reduce the influence of outliers, firms outside twice the interquartile range of firm annual revenue within this sample are excluded and the sample is further limited to firms with 30 or more transactions in a quarter and more than one transaction in 2 out of the 3 months. We convert Womply's firm categories to two-digit NAICS codes using an internally generated Womply category-NAICS crosswalk, and then aggregate to NAICS supersectors. We measure small business revenue as the sum of all credits (generally purchases) minus debits (generally returns). We define small businesses as being open on a given day if they have at least one transaction in the previous three days.

We crosswalk from ZIP codes to counties using the geographic definitions described in Appendix G to aggregate the series to the county and state levels. We then collapse the Womply data to aggregate spending and total small businesses open within each day x NAICS supersector x geography x ZIP income quartile, creating ZIP income quartiles as described in Appendix B. We take a seven-day look-back moving average of each series, and norm each series relative to its average level over the period January 4-31. We then seasonally adjust the data by following the procedure described in Appendix B.

Masking. To preserve the privacy of firms in the data and to avoid displaying noisy estimates for small cells, we mask Womply series that report less than $250,000 in total revenue during the base period of January 4-31. In addition, 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 minimize the impact of this masking by processing data at the highest level of aggregation available. When possible, we correct for this imputation by comparing similar datasets.

Anomalous Data. Our quality-control process checks for anomalous variations in the Womply raw data. There are several cases of single-day spikes of positive or negative revenue within a given firm category x ZIP code. We treat these cases as outliers, and replace the revenue value with the revenue for the same category x ZIP code from the previous week, unless that is also an outlier or a holiday, in which case we substitute zero revenue.[47]

Delayed Processing of Payments. Due to differences in the speed at which data providers share their data with Womply, data for the most recent dates as of a given data refresh are typically incomplete. If left unaddressed, there would appear to be a decline in small business revenue and small businesses open in the most recent data. We generally receive Womply's data a week after the reported transactions. As a conservative approach, we exclude the four most recent days in the data we publish.

D Employment Series Construction

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

Paychex Employment Series. We receive Paychex data at the county x industry x 2019 hourly wage quartile x 2019 firm size bin x pay frequency x week of payroll processing level. Salaried employees' wages are translated to hourly wages by dividing weekly pay by 40 hours. Since we seek to measure private sector employment, we exclude workers employed in public administration and those with an unclassified industry (which each represent 0.8% of workers as of January 2020). We

47. More generally, negative revenue may appear in the Womply data due to returns and refunds. There are a number of cases of observed negative revenue, especially during March 2020, due to consumers seeking returns or refunds on certain products. We include these cases in the Womply series, but exclude large single-day occurrences of negative revenue.

restrict the sample to workers with weekly, bi-weekly, semi-monthly or monthly pay frequencies; these workers represent over 99% of employees in the Paychex data.

We begin by creating a daily series of paychecks processed on each date by linearly interpolating daily values between each week in each county x 2-digit NAICS code x 2019 hourly wage quartile x 2019 firm size bin x pay frequency cell. In order to construct a series of employment as of each date, rather than paychecks being processed as of each date, we take two steps.

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

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

Finally, we take steps to prevent the introduction of new Paychex clients from artificially creating breaks in the employment series at smaller levels of geography. We begin by calculating the share of employment in January 2020 accounted for by each industry x firm size bin within each county x wage quartile cell. Next, we calculate the change in employment relative to January for each county x wage quartile x industry x firm size bin, and multiply this change by the share of total employment in the corresponding county x wage quartile cell, creating an employee-weighted employment series for each county x wage quartile x industry x firm size bin cell. We denote a county x wage quartile x industry x firm size bin cell as an "influential cell" if the county contains 100 or fewer than unique county x quartile x industry x firm size bin cells, and the cell accounts for over 10% of employment

in the corresponding quartile x wage quartile at any date in 2020, or if the county contains greater than 100 unique county x wage quartile x industry x firm size bin cells, and the cell accounts for over 5% of employment in the county x wage quartile at any date in 2020. We drop influential cells that record a change in employment relative to January 2020 of at least +50% on any date, on the basis that such a trend likely arises due to changes in Paychex's client base rather than true employment changes. Around 4% of county x wage quartile x industry x firm size bin cells are affected by this procedure.

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

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

We use external data sources to gather further information on firm size and industry. To obtain information on industry, we use a custom-built crosswalk created by Digital Divide Data which contains NAICS codes for each employer in the Earnin data with more than ten Earnin users. To obtain information on firm size, we crosswalk Earnin employers to ReferenceUSA data at the firm location level by spatially matching Earnin employers to ReferenceUSA firms. We begin by geocoding Earnin addresses to obtain latitudes and longitudes for each Earnin employer. We then remove common prefixes and suffixes of firm names, such as "inc" and "associated". Next, we compute the trigram similarities between firm names for all Earnin and ReferenceUSA firms within twenty-five miles of another. We then select one "match" for each Earnin firm within the ReferenceUSA data, among the subset of firms within one mile. We first match Earnin employers to ReferenceUSA firms if the firms are within one mile of one another, and share the same firm

name. Second, where no such match is available, we choose the geographically closest firm (up to a distance of one mile) among all firms with string similarities of over 0.6. Third, where no such match is available, we match an Earnin employer to the ReferenceUSA employer within twenty-five miles with the highest trigram string similarity, provided that the 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.

Earnin data are observed at the ZIP code level. We crosswalk from ZIP codes to counties using the geographic definitions described in Appendix G to aggregate the series to the county and state levels.

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. To account for the delay in receipt of paychecks, we shift the Earnin series back by one week. We then take the count of employed individuals across the Earnin sample as our measure of employment. We take a 7-day moving average to form our Earnin employment series, and express the series as a change relative to January 4-31.

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

Because the Earnin data are available at the worker level, we do not observe whether a business as a whole does not process payroll every week. However, under the assumption that all workers

within a business are paid on the same paycycle, our worker-level approach of distributing paychecks uniformly over the paycycle matches the approach in Cajner et al. (2020) of imputing employment based on the next week in which paychecks are observed. The two primary differences between our treatment of paycycles and the treatment in Cajner et al. (2020) are that we use a 7-day moving average, whereas Cajner et al. (2020) use a 14-day moving average, and that we treat that the last date of the employment period as seven days prior to the receipt of the paycheck, whereas Cajner et al. (2020) observe the pay period directly. The seven-day lag accounts for delays between the end of a worker's pay period, which is the event observed in Cajner et al. (2020), and the date on which paychecks are received by workers, which is the event observed in the Earnin data.

Because the Paychex data are not available at the firm level, we are not able to directly implement the approach in Cajner et al. (2020) of imputing employment using the count of paychecks in the "missing" week for firms that do not process payroll on a weekly basis. Instead, we make the conceptually similar assumption that workers are employed throughout the full duration of their paycycle, such that we can infer the full set of dates on which an individual worked by observing the last date of each of their pay periods and their pay frequency. Under the assumptions that all workers within a given firm are paid according to the same paycycle, our approach of inferring employment based on last date of pay period matches the approach in Cajner et al. (2020) of imputing employment based on the next week in which paychecks are observed. A further difference is that pay period is observed in Cajner et al. (2020); by contrast, in the Paychex data, pay periods are imputed using payroll processing date and the distribution of (payroll processing date – last date in pay period). Finally, Cajner et al. (2020) use a 14-day moving average, whereas we use a 7-day moving average.

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

Because Earnin best represents workers in the bottom wage quartile, we combine Earnin data with Paychex data to construct employment estimates for the bottom wage quartile. To do so, we first calculate total employment levels within each two-digit NAICS code by firm size by geography cell by summing employment levels for bottom-wage-quartile Paychex workers and Earnin workers.[48] We place the majority of the weight on Paychex, with greater weight on Earnin in geo-

48. We convert the weekly Paychex data to daily measures of employment by assuming that employment is constant

graphic areas and in NAICS codes where it has greater coverage; the exact weights are undisclosed to protect privacy. These combined Paychex-Earnin values are used to assess the effects of the Paycheck Protection Program. In order to create the other analysis datasets, we then collapse across firm sizes and compute mean levels of employment for bottom-wage-quartile workers by geography by taking a weighted average of the NAICS-by-geography combined estimates, weighting by the January Paychex NAICS shares for bottom-wage-quartile workers in each geography.

Next, we combine Intuit with the Paychex-Earnin data. Intuit provides us with overall national industry shares as of 2019, but does not release data broken down by wage level or industry. We therefore must effectively impute the Intuit data to wage-industry cells in order to combine it with the Paychex data. To do so, we assume that any differences in employment between Intuit and Paychex are constant (in percentage terms, relative to the January baseline) by industry and wage quartiles within a given geography and month. We reweight the Paychex data to match the national Intuit industry distribution and compute the percentage difference between the employment decline in the reweighted Paychex data and the Intuit data in each geography-month cell. We then apply this correction factor to each wage-industry cell in the Paychex data to obtain imputed values by wage and industry for the relevant industries covered by Intuit. For instance, if Intuit exhibits a 5% larger employment decline than the reweighted Paychex series in Manhattan in April, we would impute a value for each wage-by-industry cell covered in the Intuit data that is 1.05 times the decline observed in Paychex for that cell. When constructing the series that we use to analyze the effects of the Paycheck Protection Program, we exclude the Intuit data, since Intuit primarily consists of small firms.

Finally, we take a weighted average of the Paychex data and the imputed Intuit data in each industry to compute the final combined series. We place the majority of the weight on Paychex, with greater weight on Intuit in sectors where it has greater coverage; the exact weights are undisclosed to protect privacy.

The preceding steps yield combined data at the industry by wage quartile for each geography (county, state, and national). We construct aggregate estimates across industries, wage quartiles, and overall by aggregating these estimates using Paychex January employment weights.[49] We

within each week.

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

report seven-day moving averages of these series, expressed as a percentage change relative to January 4-31. We construct a series for average total earnings analogously, using total earnings instead of total employment.

To construct employment predictions for the most recent weeks, we regress the combined employment series for each quartile of workers on the Kronos series for the same week, the corresponding quartile of the Paychex weekly series for the same week, as well as the three prior weeks of Paychex weekly data. We then use these regression coefficients combined with the most recent Kronos and Paychex weekly data to create a series of predicted employment rates for workers in each wage quartile. Appendix Figure 15 assesses the accuracy of this forecasting approach by comparing out-of-sample predictions of employment to realized employment trends. Predicted employment is generally similar to realized employment for both high-wage (top quartile) and low-wage (bottom quartile) workers.

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

Assessing Mismeasurement of Firm Sizes using SBA data. We assess the degree of misclassification of PPP eligibility in our sample by merging publicly available data on PPP recipients from the SBA to data on firm sizes from both ReferenceUSA and Dun & Bradstreet, which form our measures of firm size in the Earnin and Paychex data, respectively. To construct SBA data on PPP recipients, 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, ReferenceUSA, and 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 ReferenceUSA firms, and all SBA and Dun & Bradstreet firms within twenty-five miles of another. We then select one "match" for each PPP recipient from both the ReferenceUSA and Dun & Bradstreet data, among the subset of firms within twenty-five miles, following the procedure described above in our merge of Earnin data to ReferenceUSA data. For firms with loans of above $150,000, exact loan size is not observed; we impute loan size as the midpoint of loan range.

We use the merged SBA-ReferenceUSA and SBA-Dun & Bradstreet data to estimate the first-stage of our difference-in-differences design, i.e. how much more PPP assistance firms classified

as having 100-500 employees in our sample received relative to those classified as having more than 500-800 employees. To do so, we stack the datasets and use the same weights used when constructing the combined employment series. 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 100-500 employees and $586 per worker for firms with 500-800 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 100-500 employees in our sample received $\frac{\$2,303}{0.62\times0.73} = \$5,090$ of PPP assistance per worker, while firms with 500-800 employees received $\frac{\$586}{0.62\times0.73} = \$1,290$ in PPP assistance per worker.[50] We calculate that PPP assistance to eligible firms with between 100 and 500 employees (excluding NAICS 72) is $\$5,092$ per worker on average.[51] Hence, firms with 501-800 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 100-500 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.

Calculating PPP Expenditures Per Worker. 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). To compute total PPP expenditure, we first use publicly released data on loan recipients to calculate that 92.1% of total PPP expenditure was received by non-NAICS 72 firms. We then multiply this share by total PPP expenditure as of August 8 to reach an estimate of $486 billion in non-NAICS 72 firms.

50. 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.

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

E Educational Progress Data Construction

In this appendix, we provide additional details about how we define Zearn indices of math progress and engagement.

Masking. The data we obtain are masked such that any county with fewer than two districts, fewer than three schools, or fewer than 50 students on average using Zearn Math during the pre-period of January 6 to February 7 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 from January 6 to February 7.

School Breaks. To reduce the effects of school breaks, we replace the value of any week for a given school that reflects a 50% decrease (increase) greater than the week before or after it with the mean value for the three relevant weeks.

F Public Data Sources

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

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

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

last received the benefits.

COVID-19 Data. We report the number of new COVID-19 cases and deaths each day using publicly available data from the New York Times available at the county, state and national level.[52] We also report daily state-level data on the number of tests performed per day per 100,000 people from the COVID Tracking Project.[53] For each measure - cases, deaths, and tests – we report two daily series per 100,000 people: a seven-day moving average of new daily totals and a cumulative total through the given date. 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, unless there was also a spike on that day.

Google Mobility Reports. We use data from Google's COVID-19 Community Mobility Reports to construct measures of daily time spent at parks, retail and recreation, grocery, transit locations, and workplaces.[54] We report these values as changes relative to the median value for the corresponding day of the week during the five-week period from January 3rd - February 6, 2020. Details on place types and additional information about data collection is available from Google. We use these raw series to form a measure of time spent outside home as follows. We first use the American Time Use survey to measure the mean time spent inside home (excluding time asleep) and outside home in January 2018 for each day of the week. We then multiply time spent inside home in January with Google's percent change in time spent at residential locations to get an estimate of time spent inside the home for each date. The remainder of waking hours in the day provides an estimate for time spent outside the home, which we report as changes relative to the mean values for the corresponding day of the week in January 2020. Finally, we report each series as a seven-day moving average.

G Dates and Geographic Definitions

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

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

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

54. Google Mobility trends may not precisely reflect time spent at locations, but rather "show how visits and length of stay at different places change compared to a baseline." We call this "time spent at a location" for brevity.

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

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

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

- Schools closed statewide: Sourced from COVID-19 Impact: School Status Updates by MCH Strategic Data, available here. Compiled from public federal, state and local school information and media updates.
- Nonessential businesses closed: Sourced from the Institute for Health Metrics and Evaluation state-level data (available here), who define a non-essential business closure order as: "Only locally defined 'essential services' are in operation. Typically, this results in closure of public spaces such as stadiums, cinemas, shopping malls, museums, and playgrounds. It also includes restrictions on bars and restaurants (they may provide take-away and delivery services only), closure of general retail stores, and services (like nail salons, hair salons, and barber shops) where appropriate social distancing measures are not practical. There is an enforceable consequence for non-compliance such as fines or prosecution."
- Stay-at-home order goes into effect: Sourced and verified from the New York Times reopening data, available here, and hand-collection from local news and government sources where needed.
- Stay-at-home order ends: Sourced and verified from the New York Times reopening data, available here, and hand-collection from local news and government sources where needed. Defined as the date at which the state government lifted or eased executive action or other policies instructing residents to stay home. We code "regional" and "statewide" expiry of stay-at-home orders separately. A "regional" expiration of a stay-at-home orders occurs when a stay-at-home order expires in one region within a state, but not everywhere within the state. A "statewide" expiration of a stay-at-home order occurs when a stay-at-home order

first expired throughout a whole state, either due to a statewide change in policy, or due to the stay-at-home order in each county having expired.

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

Geographic Definitions. For many of the series we convert from counties to metros and ZIP codes to counties. We use the HUD-USPS ZIP code Crosswalk Files to convert from ZIP code to county. When a ZIP code corresponds to multiple counties, we assign the entity to the county with the highest business ratio, as defined by HUD-USPS ZIP Crosswalk. We generate metro values for a selection of large cities using a custom metro-county crosswalk, available in Appendix Table 7. We assigned metros to counties and ensured that a significant portion of the county population was in the metro of interest. Some large metros share a county, in this case the smaller metro was subsumed into the larger metro. We use the Uniform Data Systems (UDS) Mapper to crosswalk from ZIP codes to ZCTAs.

Table 1

Changes in Consumer Spending on Debit and Credit Cards by Sector and Income Quartile

Dep. Var.:	Change in Consumer Debit and Credit Card Spending Per Day Relative to January 4-31 2020 ($ Billions)			Level of Consumer Spending Per Day ($ Billions)
	Change as of April 8-14	Change as of June 8-14	Change as of July 8-14	Level as of January 4-31 2020
	(1)	(2)	(3)	(4)
Panel A: Card Spending by Income Quartile				
Pooled, All Income Quartiles	-$7.5	-$3.1	-$3.1	$22.0
Bottom Quartile	-$0.9	-$0.2	-$0.3	$3.3
	(12.5%)	(6.8%)	(8.5%)	(15.2%)
Second Quartile	-$1.5	-$0.6	-$0.6	$4.9
	(20.1%)	(17.5%)	(19.4%)	(22.4%)
Third Quartile	-$2.0	-$0.8	-$0.8	$6.0
	(26.6%)	(26.6%)	(26.8%)	(27.1%)
Top Quartile	-$3.1	-$1.5	-$1.4	$7.8
	(40.9%)	(49.1%)	(45.3%)	(35.3%)
Panel B: Card Spending by Sector				
Durable Goods	-$0.8	+$1.0	+$0.9	$4.9
	(10.6%)	(-42.5%)	(-45.6%)	(22.7%)
Non-Durable Goods	-$0.6	+$0.1	+$0.2	$4.9
	(8.0%)	(-3.0%)	(-9.9%)	(22.3%)
Remote Services	-$1.2	-$0.2	-$0.2	$4.5
	(15.0%)	(6.9%)	(11.5%)	(20.4%)
In-Person Services	-$5.1	-$2.9	-$3.0	$6.9
	(65.1%)	(127.3%)	(156.1%)	(31.8%)
Panel C: In-Person Services Sub-Sector Decomposition				
Hotels & Food	-$2.0	-$1.1	-$1.2	$2.6
	(38.3%)	(39.6%)	(38.6%)	(37.6%)
Transportation	-$1.5	-$1.1	-$1.1	$1.7
	(29.4%)	(37.3%)	(36.2%)	(25.0%)
Health Care	-$0.5	-$0.1	-$0.2	$0.9
	(10.0%)	(3.9%)	(4.8%)	(12.3%)
Recreation	-$0.4	-$0.3	-$0.3	$0.5
	(8.4%)	(11.3%)	(10.9%)	(7.3%)
Other In-Person Services	-$0.7	-$0.2	-$0.3	$1.2
	(13.9%)	(8.0%)	(9.4%)	(17.8%)

Notes: This table describes how national consumer spending on credit and debit cards changes over time, relative to its level in January 2020. To compute these changes, we begin by calculating total daily spending in the Affinity Solutions data for each day in 2019 and 2020. We then inflate these numbers to estimate total card spending for the full U.S. population by multiplying by the ratio of January 2020 (or 2019) total spending for components of PCE that are likely captured in credit/debit card spending (shown in the last bar of Figure 1a) to the January 2020 (or 2019) total spending in the Affinity data. In Column (1), we calculate the change in card spending as ((Spending for April 8 through April 14 2020) - (Spending for April 8 through April 14 2019)) - ((Spending for January 4 through January 31 2020) - (Spending for January 4 - January 31 2019)). Columns (2) and (3) replicate Column (1) for the June 8-14 and July 8-14 periods respectively, instead of April 8-14. Column (4) shows mean daily national spending over the period of January 4-31 2020. Panel A shows changes in total spending for the full sample and separately by income quartile, imputed using median household income within the cardholder's residential ZIP code in the 2014-2018 ACS. The share of the national decline accounted for by each quartile is shown in parentheses in each row. The declines in card spending among bottom-quartile and top-quartile ZIP codes in Panel A Columns (1) and (2) appear in the annotations on Figure 2a. Panel B replicates Panel A, disaggregating by categories of goods instead of income quartiles. For definitions of these categories, see Appendix B. The shares of the total change in consumer spending summed across these four categories in Panel B do not necessarily add up to 100% because of spending on uncategorized goods or services. Panel C disaggregates the change in consumer spending within in-person services across five sub-categories, which add up to 100% of in-person services. Data source: Affinity Solutions.

Table 2

Association Between Rent and Changes in Business Revenue and Employment, by ZIP Code

Panel A: Business Revenue

Dep. Var.:	Change in Small Business Revenue (%)		
	(1)	(2)	(3)
Median 2BR Rent (per thousand dollars)	-13.47	-14.00	-10.55
	(0.35)	(0.67)	(0.69)
Log of Density of High Wage Workers			-2.26
			(0.12)
County FEs		X	X
Observations	18378	18378	17159

Panel B: Low-Wage Employment

Dep. Var.:	Change in Low-Wage Employment at Small Businesses (%)		
	(1)	(2)	(3)
Median 2BR Rent (per thousand dollars)	-13.93	-9.70	-6.89
	(0.53)	(0.95)	(0.99)
Log of Density of High Wage Workers			-1.80
			(0.18)
County FEs		X	X
Observations	15685	15685	14502

Notes: This table presents estimates from population-weighted OLS regressions at the ZIP-code level, regressing percentage changes in small business revenue (using Womply data) and small business low-wage employment (using Earnin data) on average median two-bedroom rent within the ZIP-code (as measured in the 2014-2018 ACS). Standard errors are reported in parentheses. The dependent variable is scaled from 0 to 100, so that, for example, the coefficient of -13.47 in the first row of Column (1), Panel A implies that a $1000 increase in monthly two-bedroom rent is associated with a 13.47 percentage point larger decline in total revenue relative to the January level. The dependent variable in Panel A is the change in small business revenue between January 4-31 and March 25-April 14. The dependent variable in Panel B is the change in low-wage employment at small businesses between January 4-31 and April 8-28, defining small businesses as firms with fewer than 500 employees. In both panels, the dependent variable is winsorized at the 99th percentile of the (population-weighted) ZIP-level distribution. Column (1) shows the baseline regression without any controls: this specification corresponds to the estimated slope coefficient and standard error reported in Figure 5c (small business revenue) and Figure 9a (low-wage employment at small businesses). Column (2) adds county fixed effects and Column (3) further adds the log of the density of high wage workers as a control (which is observed using the Census LODES for 92% of ZIP codes representing 99% of the U.S. population). Data sources: Panel A: Womply; Panel B: Earnin.

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

Dep. Var.:	Spending (%)		Employment (%)		Low-Wage Employment (%)	High-Wage Employment (%)	Small Businesses Open (%)		Time Spent Outside Home (%)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
DD Estimate of	1.43	1.37	0.65	1.04	-0.30	1.26	3.27	4.44	0.76	0.92
Effect of Reopening:	(0.51)	(0.53)	(0.51)	(0.97)	(0.85)	(0.88)	(1.26)	(1.85)	(0.56)	(0.98)
State-Week Observations:	200	312	208	258	248	248	244	324	112	138
Analysis Window (weeks on either side of	2	3	2	3	2	2	2	3	2	3
Mean Decline in Outcome (Jan-April):	-28.2		-18.3		-29.2	-10.5	-40.5		-21.1	

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 and April 27. Each state that reopens is matched to multiple control states (listed in Appendix Table 6) 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 of the change in the outcome variable relative to January 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-difference (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-difference estimate for the effect of reopening on consumer spending over a two-week horizon is a 1.45 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 4-31 to the period March 25-April 14 — except in columns (7) and (8) which use January 4-31 instead of January 4-31 as the January reference period, and columns (9) and (10) which use January 3-February 6 as the January reference period. Column (1) shows the estimated effect of reopening on consumer spending using data from Affinity Solutions. Consumer spending is expressed as a percentage change relative to its level over the period January 4-31, 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-Earnin data and expressed as a percentage change relative to its level over the period January 4-31. Columns (5) and (6) replicate column (3) for employment in the bottom wage quartile and top wage quartile respectively. Columns (7) and (8) replicate columns (1) and (2) respectively using number of small businesses open as the dependent variable, calculated using Womply data and expressed as a percentage change relative to its level over the period January 4-31. Columns (9) and (10) replicate columns (1) and (2) respectively using time spent away from home as the dependent variable, calculated using Google Mobility data and expressed as a percentage change relative to its level over the period January 3-February 6. Columns (1), (3), and (7) correspond to the specifications displayed in Figures 12B, 12C, and 12D respectively. Data sources: Affinity Solutions, Paychex, Intuit, Earnin, Womply, Google.

Table 4

Causal Effect of Stimulus Payments on Spending and Small Business Revenue:

Regression Discontinuity Estimates

Panel A: Impact of Stimulus Payments on Consumer Spending

Dep. Var.:	Change in Consumer Spending (%)			
	Bottom Income Quartile ZIP Codes		Top Income Quartile ZIP Codes	
	(1)	(2)	(3)	(4)
RD Effect of Stimulus:	25.15	36.97	8.45	15.83
	(7.15)	(9.81)	(3.83)	(5.14)
Window:	April 1 - April 30	April 7 - April 21	April 1 - April 30	April 7 - April 21

Panel B: Impact of Stimulus Payments on Small Business Revenue

Dep. Var.:	Change in Small Business Revenue (%)			
	Bottom Rent Quartile ZIP Codes		Top Rent Quartile ZIP Codes	
	(1)	(2)	(3)	(4)
RD Effect of Stimulus:	17.92	20.83	1.20	-7.54
	(9.59)	(16.76)	(6.27)	(10.45)
Window:	April 1 - April 30	April 7 - April 21	April 1 - April 30	April 7 - April 21

Notes: This table shows regression discontinuity estimates of changes in spending and business revenue around the date of stimulus payments on April 15, 2020. Panel A shows estimated effects of stimulus payments on consumer spending. To construct the estimates, we first express consumer spending on each day as a percentage change relative to mean daily consumer spending over the period January 4-31 in the corresponding calendar year. We then residualize these daily percentage changes with respect to day of week and first day of the month fixed effects, which we estimate using data from January 1, 2019, to May 10, 2019. We then compute OLS regressions of the residualized outcome variable on an indicator variable for the date being on or after April 15 2020, using a linear control function before and after April 15, and excluding the partially treated date of April 14. The first row shows the coefficient on the indicator variable for the date being on or after April 15, which we refer to as the RD effect of stimulus; standard errors are reported in parentheses. The dependent variable is scaled as a percentage change from January so that, for example, the first row of Column (1) indicates that stimulus payments increased consumer spending by 25.15 percentage points in bottom income quartile ZIP codes relative to the January 2020 level of spending. In columns (1) and (2), we compute daily changes in spending restricted to cardholders in ZIP codes in the bottom quartile of the distribution of ZIP code median household income (based on data from the 2014-2018 ACS). Columns (3) and (4) replicate columns (1) and (2), computing daily changes in spending restricted to cardholders living in the top income quartile of ZIP codes. The coefficient and standard error in columns (1) and (3) of Panel A correspond to the specifications displayed in Figures 13b and 13c. Panel B shows regression discontinuity estimates for the effect of stimulus payments on small business revenue using data from Womply. We first express small business revenue on each day relative to mean daily small business revenue over the period January 4-31 of the corresponding year. We then residualize daily changes in small business revenue as in Panel A, and compute OLS regressions as in Panel A. Columns (1) and (2) restrict to ZIP codes in the bottom quartile of the distribution of ZIP code median rent for a two bedroom apartment (based on data from the 2014-2018 ACS). Columns (3) and (4) restrict to businesses in the top rent quartile ZIP codes. The coefficient and standard error in columns (1) and (3) of Panel B correspond to the specifications displayed in Figures 14a and 14b. In both panels, columns (1) and (3) include all of April 2020 in the regression specification, while columns (2) and (4) restrict to a narrower bandwidth, within one week on either side of the stimulus payment date. Data sources: Panel A: Affinity Solutions; Panel B: Womply.

Table 5

Causal Effect of Paycheck Protection Program on Employment: Difference-in-Difference Estimates

Dep. Var.:	Change in Employment (%)			
	Combined Paychex and Earnin Data		Earnin Data	Kronos Data
	(1)	(2)	(3)	(4)
	Baseline Estimate (100-800 Employees)	Smaller Bandwidth (300-700 Employees)	Baseline Estimate (100-800 Employees)	Baseline Estimate (100-800 Employees)
DD Estimate	1.78	1.62	1.01	-0.22
	(1.99)	(2.68)	(0.94)	(2.11)

Notes: This table shows difference-in-difference (DD) estimates of the effect of PPP eligibility (defined as the parent firm having 500 or fewer employees) on employment. The outcome variable is employment at the county x 2-digit NAICS x income quartile x PPP eligibility x week level, excluding the Accommodations and Food Services Sector (NAICS 72), expressed as a percentage change relative to a pre-period of January 4-31, 2020. Columns (1) and (2) present regressions in combined Paychex-Earnin data. In the baseline estimate in column (1), which corresponds to the difference-in-difference estimate in Figure 15a, we begin by restricting to firms with between 100 and 800 employees. We then reweight firms on 2-digit NAICS codes such that the (worker-weighted) distribution of 2-digit NAICS codes within Paychex firms in each size bin matches the national distribution of 2-digit NAICS codes among Paychex firms in the period January 4-31 2020. We do the same for Earnin firms. Next, to combine the datasets, we reweight such that the (worker-weighted) share of each dataset is constant in each firm size bin. We then sum employment across datasets at the county x 2-digit NAICS x income quartile x eligibility x week level. Finally, we report estimates from regression equation (1): an OLS regression of changes in employment on county x worker income quartile x week fixed effects, an indicator for PPP eligibility (firm size <500 employees), an indicator for dates after April 3, and an interaction term between PPP eligibility and the date being after April 3 (the DD estimate). The sample for this regression is limited to weeks ending between March 11 and August 15; we restrict the sample to these weeks because Figure 15a indicates that there is little evidence of a persistent effect on employment for businesses that were eligible for PPP assistance after mid-August. 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 level, and winsorize the dependent variable at the 99th percentile (weighted by reweighted employment over the period January 4-31). Column (2) replicates Column (1), restricting to firms with between 300 and 700 employees. Columns (3) and (4) replicate Column (1) in the Earnin data alone and the Kronos data alone, respectively. Column (3) corresponds to the specification for the Earnin-only difference-in-difference estimate reported in Appendix Figure 15a. As we treat all Earnin workers as belonging to the first quartile, Column (3) uses county x week fixed effects, rather than county x worker income quartile x week fixed effects. Data sources: Paychex, Earnin, Kronos.

Appendix Table 1

Industry Employment Shares Across Datasets

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

Notes: This table compares the industry (2-digit NAICS) composition of four private employment-based 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-Earnin 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 Earnin as the number of worker-days in each 2-digit NAICS code. We then calculate combined Paychex-Earnin employment in each 2-digit NAICS code as a weighted sum of Paychex and Earnin, where the weights are undisclosed to meet privacy protection requirements. Column (4) shows the industry composition of Intuit data in January 2020. Column (5) shows the industry composition of Kronos in January 2020, constructed by calculating the total number of "punches" (i.e. worker-days) in each 2-digit NAICS code. Column (6) shows the industry composition of Homebase in January 2020. To construct Column (6), we crosswalk Homebase industry classifications to 2-digit NAICS codes and calculate the number of unique employee x employer pairs with a positive number of recorded hours in January 2020 in each 2-digit NAICS code. Where a dataset has no coverage in a given NAICS code, we present the employment share in that dataset as blank. Data sources: Homebase, Paychex, Earnin, Kronos, Intuit.

Appendix Table 2

Hourly Pre Tax Wage Rates By Industry Across Datasets

		Occupational Employment Statistics (OES)	Mean in Private Datasets		
		May 2019 OES	Homebase	Paychex - Earnin	Intuit
NAICS Code	NAICS Description	(1)	(2)	(3)	(4)
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	$12.47	$21.07	
48-49	Transportation and Warehousing	$24.33	$14.63	$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	$14.20	$34.37	
55	Management of Companies and Enterprises	$42.59		$24.06	
56	Administrative Support	$20.50		$23.51	
61	Educational Services	$28.34	$12.57	$24.78	
62	Health Care and Social Assistance	$26.98	$15.87	$25.47	
71	Arts, Entertainment, and Recreation	$19.18	$12.26	$22.47	
72	Accommodation and Food Services	$13.65	$11.11	$16.62	
81	Other Services (except Public Administration)	$21.58	$14.80	$22.50	
	All	$25.72	$12.11	$25.34	$26.27
	Industry-Weighted Average of BLS Mean Wages		$17.24	$26.00	$27.79

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 Homebase in January 2020. In Homebase data, mean wages are measured as pre-tax wages recorded by the employer. Mean wages are computed weighting by worker-days, excluding workers whose reported mean wage is zero. Column (3) 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. Column (4) shows mean wages in Intuit in January 2020. As we do not observe wages within each industry in Intuit data, we report only the average wage across all industries. Observed wages in Intuit data represented pre-tax wages as recorded by the employer. The last row of Columns (2) - (4) displays BLS mean wages, reweighted to match the 2-digit NAICS composition within each private sector dataset (as constructed in Appendix Table 1). Data sources: Earnin, Homebase, Paychex, Intuit.

Appendix Table 3

Demographic Characteristics of Zearn Users

	Zearn Users (1)	U.S. Population (2)
Panel A: Income		
ZIP Median Household Income		
25th Percentile	43,355	45,655
Median	54,941	57,869
75th Percentile	71,485	77,014
Number of ZIP codes	6,529	33,253
Number of People	925,978	322,586,624
Panel B: School Demographics		
Share of Black Students		
25th Percentile	1.2%	1.5%
Median	5.2%	5.8%
75th Percentile	21.3%	19.1%
Share of Hispanic Students		
25th Percentile	4.3%	5.6%
Median	11.4%	15.0%
75th Percentile	33.5%	40.6%
Share of Students Receiving FRPL		
25th Percentile	35.7%	28.2%
Median	56.9%	50.1%
75th Percentile	80.4%	74.8%
Number of Schools	11,400	88,459
Number of Students	887,592	49,038,524

Notes: This table reports demographic characteristics for Zearn schools vs. the U.S. population. Panel A compares income characteristics of ZIP codes in which Zearn schools are located vs. all ZIP codes. We define Zearn to have coverage in a ZIP code if at least five students at schools located in the ZIP code used Zearn over the period January 6-February 7. Column (1) shows income characteristics of ZIP codes in which Zearn schools are located. The first three rows in Panel A display the 25th percentile, median, and 75th percentile of ZIP-level median household income in ZIP codes in which Zearn has coverage. Household income percentiles are calculated using the 2014-2018 median household income in each school's ZIP code, as measured in the ACS. The fourth and fifth rows of Panel A display the number of ZIP codes in which Zearn has coverage, 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) display the total number of ZIP codes in the U.S. and the total population of the U.S., respectively. Panel B compares the demographic composition of Zearn schools vs. the demographic composition of all U.S. schools. We calculate the demographic characteristics in Panel B using school-level data from the Common Core data set constructed by MDR Education, a private education data firm. The first three rows of Panel B, Column (1) show the 25th percentile, median, and 75th percentile of share of Black students among schools with coverage in Zearn. 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. Row 10 of Panel B displays the number of Zearn schools matched to the Common Core data, and the number of students in those schools, respectively. Column (2) replicates Column (1) using data for all U.S. schools, rather than restricting to Zearn schools. Data source: Zearn.

Appendix Table 4

Cities with Largest Small Business Revenue Losses Following COVID Shock

City	State	Change in Small Bus. Revenue
(1)	(2)	(3)
New Orleans	Louisiana	-77%
Washington	District of Columbia	-74%
San Francisco	California	-69%
New York City	New York	-68%
Boston	Massachusetts	-65%
Honolulu	Hawaii	-65%
Charlotte	North Carolina	-64%
Philadelphia	Pennsylvania	-63%
San Jose	California	-62%
Baltimore	Maryland	-59%

Notes: This table shows the ten cities with the largest small business revenue declines as measured in the Womply data (among the fifty largest cities in the U.S.). Columns (1) and (2) display the name of the city and the state in which it is located. Column (3) shows the decline in small business revenue, computed as the change in net small business revenue in Womply data between January 4-31 2020 and March 25 2020-April 14 2020, seasonally adjusted using 2019 values of net revenue. Data source: Womply.

Appendix Table 5

Association Between Changes in Consumer Spending in ZIP Code and Workplace Rent

Dep. Var.:	Change in Consumer Spending (%)		
	(1)	(2)	(3)
Mean Workplace 2BR Rent	-12.61	-9.55	-16.49
(per thousand dollars)	(0.69)	(1.47)	(4.38)
Median Home 2BR Rent		-4.80	
(per thousand dollars)		(1.93)	
Controls:			
County Fixed Effects			X
Observations	8837	6624	8837

Notes: This table shows OLS regressions of average percentage changes in consumer spending by ZIP code (using data from Affinity Solutions) on average workplace ZIP code median two-bedroom rent. Standard errors are reported in parentheses. In each regression, we restrict the sample to ZIP codes in the bottom quartile of median household income. We compute ZIP-level changes in consumer spending in the Affinity data as the percentage change in seasonally-adjusted spending from the period January 4-31 to the period March 25-April 14, winsorizing at the 99th percentile of (population-weighted) ZIP-level changes in employment. We construct the average workplace rent variable by combining data on the matrix of home residence by workplace ZIP codes taken from Census' LEHD Origin-Destination Employment Statistics (LODES) for low-income workers (workers earning below $1,250 per month) with data on median two bedroom monthly rents from the 2014-2018 ACS. In particular, we assign median rents from the ACS to each ZIP code of workplace in the LODES data and then compute mean workplace rent in each home ZIP code, weighting by the number of low-wage jobs in each workplace ZIP code. The dependent variable is scaled from 0 to 100 and the independent variable is expressed in thousands such that, for example, the coefficient of -12.61 in Column (1) implies that a $1000 increase in monthly workplace rent is associated with a 12.61% larger drop in total spending. Column (1) shows the baseline regression without any controls. This specification corresponds with the slope coefficient and standard error in Figure 11c. Column (2) adds median home two bedroom rent. Column (3) adds county fixed effects. Data source: Affinity Solutions.

Appendix Table 6

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

Date	States that Re-Opened	Consumer Spending Controls (1)	Employment Controls (2)	Small Businesses Open Controls (3)	Mobility Controls (4)
April 20th, 2020	South Carolina	California, Connecticut, Delaware, Florida, Hawaii, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Missouri, Nebraska, New Mexico, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	California, Connecticut, Delaware, District Of Columbia, Florida, Hawaii, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Oregon, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	Indiana, Missouri, Nebraska, New Mexico, Oregon, South Dakota
April 24th, 2020	Alaska, Georgia	California, Connecticut, Delaware, Florida, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Missouri, Nebraska, New Mexico, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	California, Connecticut, Delaware, District Of Columbia, Florida, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Missouri, Nebraska, New Jersey, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	Delaware, Indiana, Louisiana, Missouri, Nebraska, New Mexico, South Dakota, Virginia, Wisconsin
April 27th, 2020	Minnesota, Mississippi	Illinois, Nebraska, New Jersey, Pennsylvania, Virginia	California, Connecticut, Delaware, District Of Columbia, Illinois, Maryland, Nebraska, New Mexico, South Dakota, Virginia, Washington, Wisconsin	California, Connecticut, Delaware, District Of Columbia, Illinois, Maryland, Nebraska, New Mexico, New York, Pennsylvania, South Dakota, Virginia, Washington, Wisconsin	Delaware, Illinois, Nebraska, New Mexico, Pennsylvania, South Dakota, Virginia, Wisconsin

Notes: This table lists the treatment and control states for the analysis of state reopenings in Figures 12b-d and Table 3. Column (1) displays the control states that are compared to the treatment states in the two-week event horizon in the event study of consumer spending (as measured in the Affinity data) described in Figure 12b-d and Table 3; for details, see notes to Figure 12b-d and Table 3. Column (2) replicates Column (1) for employment (as measured in the Paychex-Intuit-Earnin data). Column (3) replicates Column (1) for number of small businesses open (as measured in the Womply data). Column (4) replicates Column (1) for time away from home (as measured in the Google Mobility data).

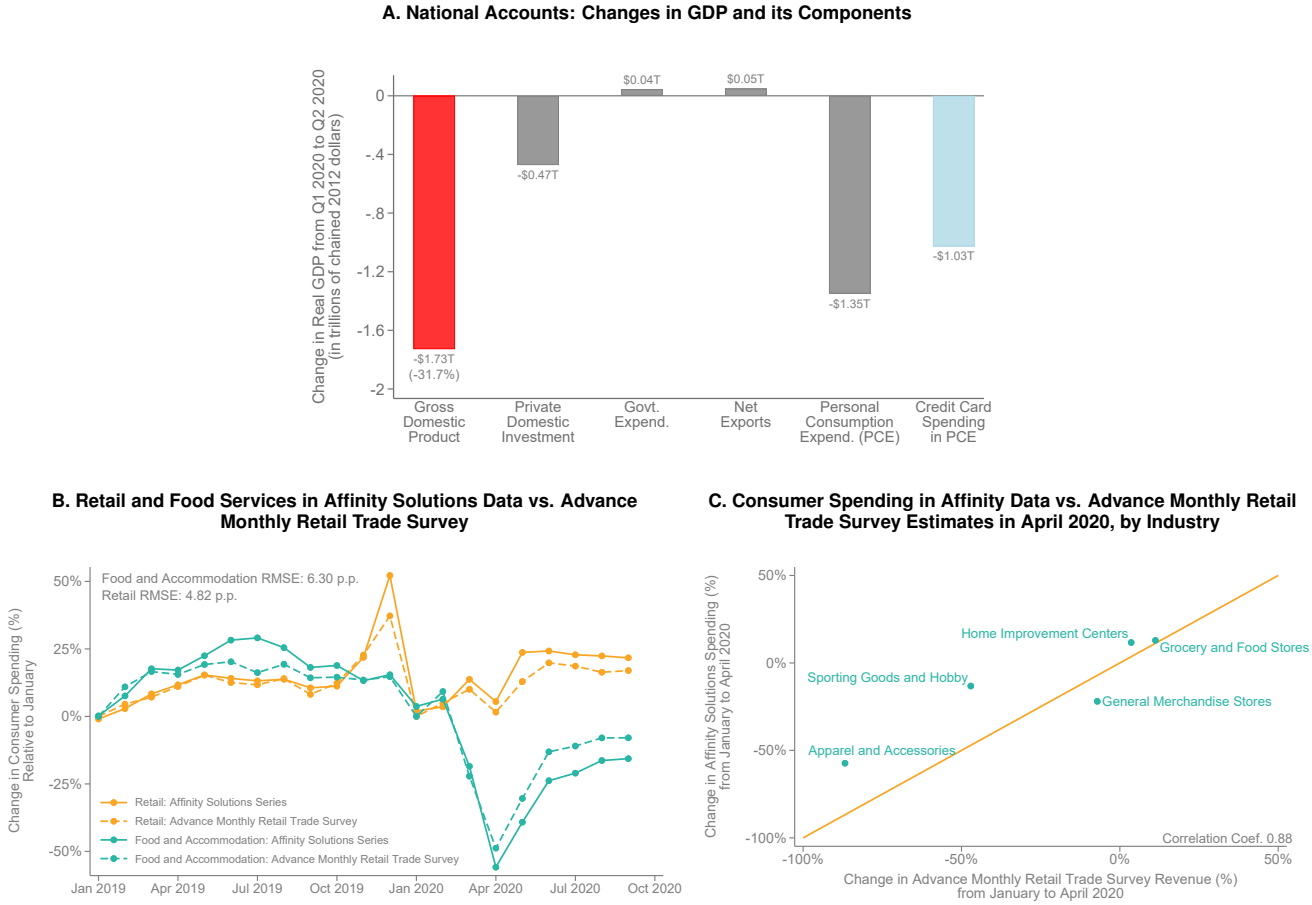
Appendix Table 7

City to County Crosswalk

City Name	State Name	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

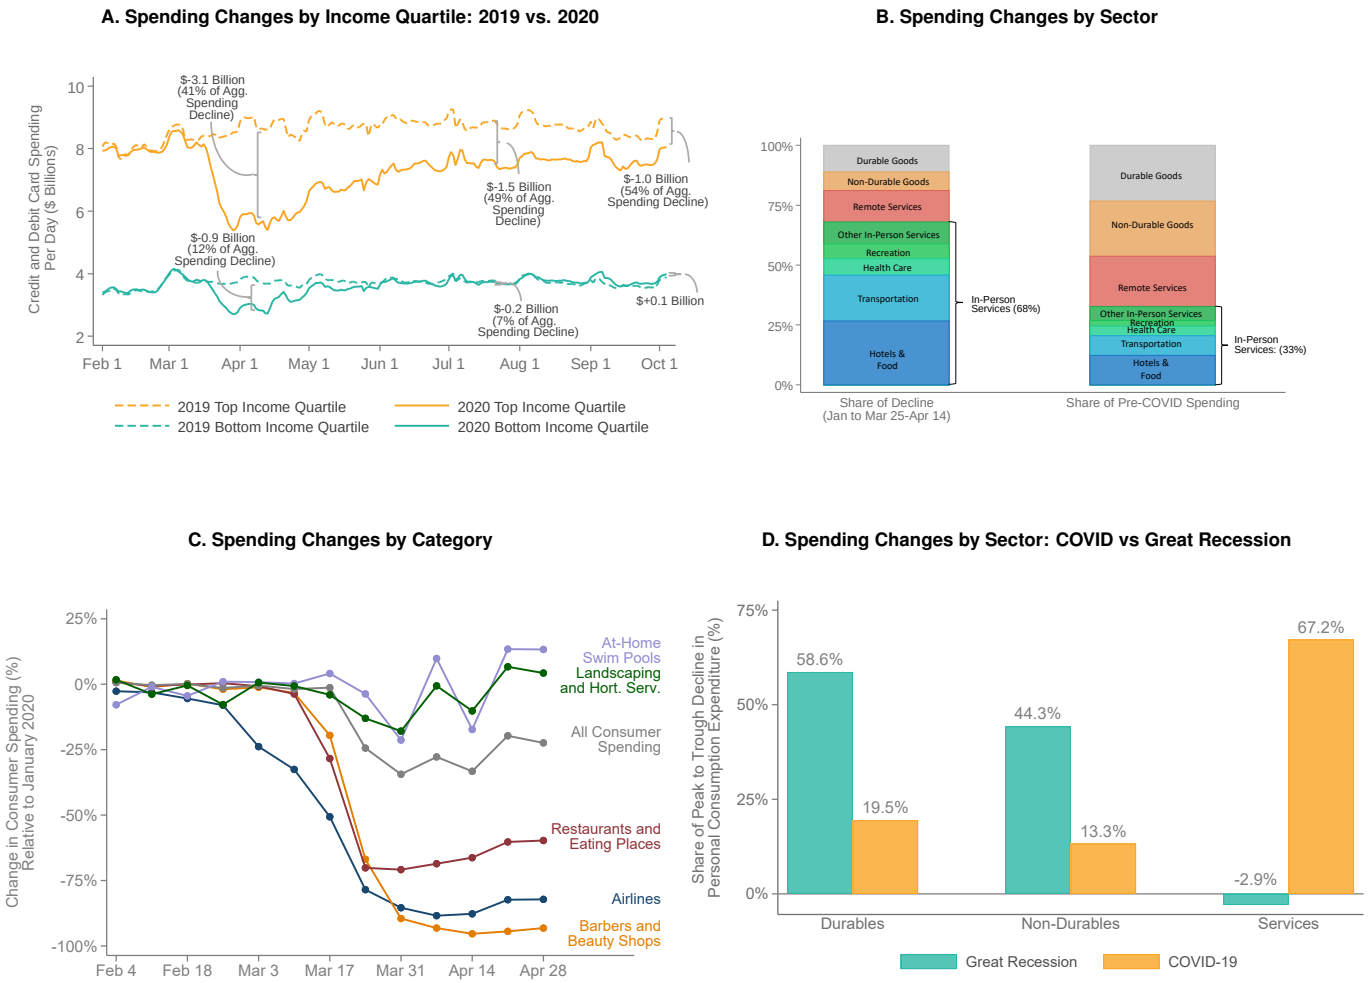
Notes: This table shows our metro area (city) to county crosswalk. We assigned metros to counties and ensured that a significant portion of the county population was in the metro of interest. Some large metros share a county; in this case the smaller metro was subsumed into the larger metro.

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



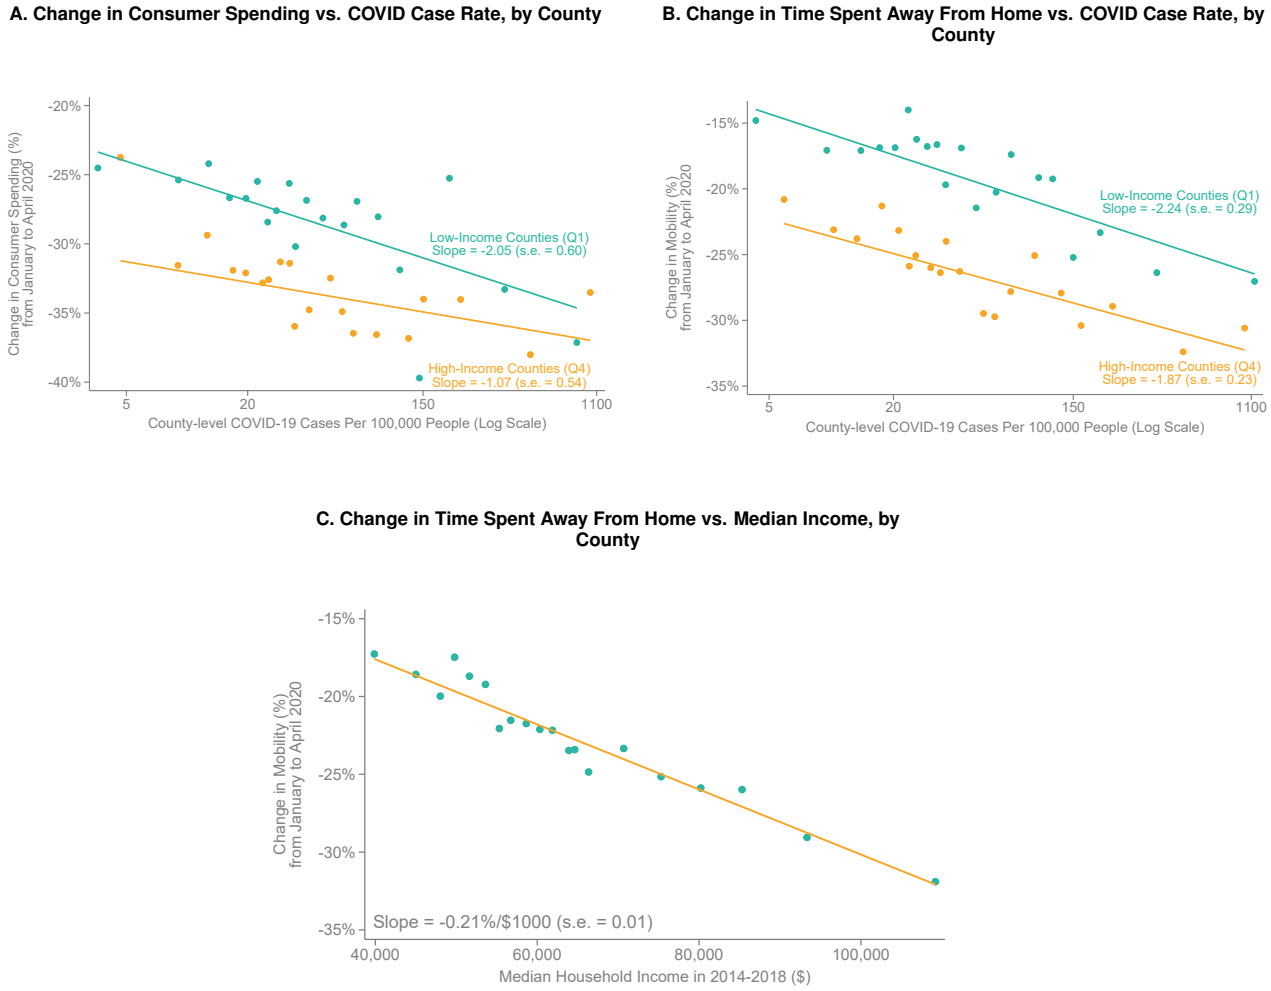
Notes: This figure compares changes in consumer spending measured in Affinity Solutions data on debit and credit card expenditures to the National Income and Product Accounts (NIPA) data and Advance Monthly Retail Trade Survey (MARTS) data. Panel A examines the change in GDP from Q1-2020 to Q2-2020 using NIPA data (Tables 1.1.2, 1.1.6 and 2.3.2 from the Second Estimate of 2nd Quarter 2020 Gross Domestic Product, released on August 27 2020). The first bar shows the seasonally-adjusted change in real GDP (-$1.73T). In parentheses under the first bar we report the compound annual growth rate corresponding to this one-quarter change in real GDP (-31.7%). Bars two through five show the contribution to the change in real GDP of its components, estimated using NIPA Table 1.1.2. The final bar shows the contribution of components of Personal Consumption Expenditures (PCE) that are likely to be captured in credit card spending (-$1.03T), 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, e.g. 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 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 retail series in Panel B restricts to retail trade sectors (NAICS 44-45) excluding motor vehicles (NAICS 441) and gas (NAICS 447). The MARTS series is constructed by dividing the total spending in each category by the number of days in that month, and then indexing the average daily spending to January of the corresponding year. The Affinity series is constructed by taking the monthly average of the seven-day moving average series indexed to January of the respective year. We also report the root mean squared error (RMSE) corresponding to the difference between indexed MARTS monthly spending and indexed Affinity monthly spending. Panel C 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. Data source: Affinity Solutions.

FIGURE 2: Changes in Consumer Spending During the COVID Crisis



Notes: This figure disaggregates spending changes by income and sector using debit and credit card data from Affinity Solutions. Panel A plots a weekly series of average daily consumer spending, separately for cardholders residing in ZIP codes in the top and bottom (population-weighted) quartiles of median household income (measured using the 2014-2018 ACS). We inflate the spending observed in Affinity data to estimate total spending for the full U.S. population by multiplying by the ratio of January 2020 (or 2019) total spending for components of PCE that are likely captured in credit/debit card spending (described in the notes to Figure 1) to the January 2020 (or 2019) total spending in the Affinity data. The annotations in Panel A display the change in spending during the periods of April 8-14, July 8-14 and Sept 28-Oct 4 respectively (calculated as described in Table 1). Panel B disaggregates seasonally-adjusted spending changes (left bar) and pre-COVID levels (right bar) by sector. In the left bar of Panel B, the total decline being decomposed is defined as ((Spending in March 25 through April 14 2020) - (Spending in March 25 through April 14 2019)) - ((Spending in January 4 through January 31 2020) - (Spending in January 4 through January 31 2019)). The right bar of Panel B shows the sectoral shares of spending during January 4-31 2020. See Appendix B for the definitions of these sectors. Panel C compares seasonally-adjusted trends in weekly consumer spending for five specific categories of goods and pooled consumer spending. Panel D decomposes the change in personal consumption expenditures (PCE) in the Great Recession and the COVID-19 Recession using NIPA Table 2.3.6U. PCE is defined here as the sum of services, durables and non-durables in seasonally adjusted, chained (2012) dollars. The peak to trough declines are calculated from December 2007 to June 2009 for the Great Recession and from January 2020 to April 2020 for the COVID-19 Recession. Data source: Affinity Solutions.

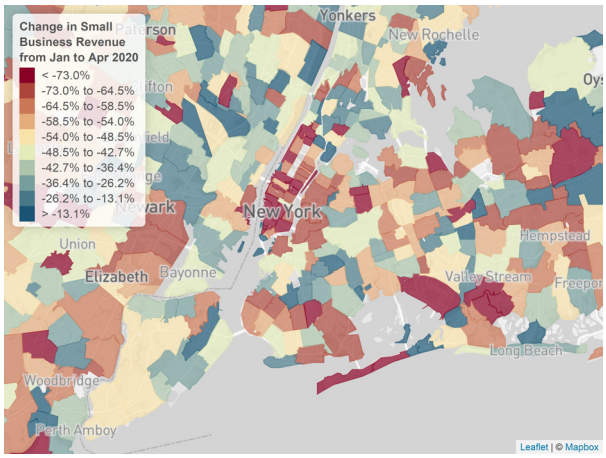
FIGURE 3: Association Between COVID-19 Incidence, Spending, and Time Away From Home



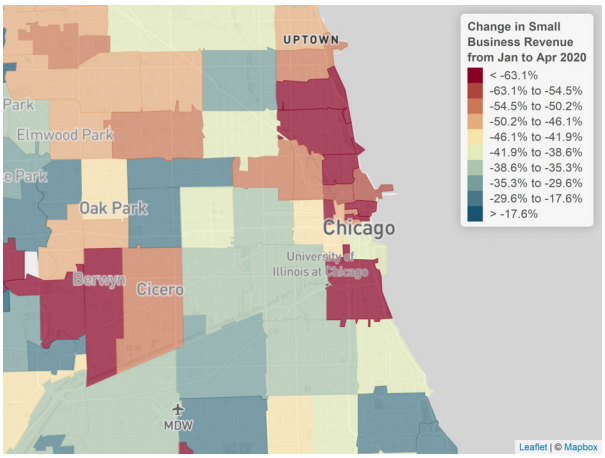
Notes: This figure presents three county-level binned scatter plots. To construct each binned scatter plot, we divide the data into twenty equal-sized bins, ranking by the x-axis variable and weighting by the county's population, and plot the (population-weighted) means of the y-axis and x-axis variables within each bin. Panel A presents the change in seasonally-adjusted consumer spending from the base period (Jan 4-31) to the three-week period of March 25-April 14. Panel B presents the change in time spent away from home from the base period (Jan 3-Feb 6) to the three-week period of March 25-April 14. We exclude weekends when calculating the change in time spent outside home. In both Panels A and B, the x-axis variable is the logarithm of the county's mean daily cumulative COVID case rate per capita over the period March 25-April 14. These panels plot values separately for counties in the top and bottom quartiles of median household income (measured using population-weighted 2014-2018 ACS data). Panel C replicates Panel B with county median household income on the x-axis (measured in 2014-2018 ACS data). Data sources: Affinity Solutions, Google Community Mobility Reports.

FIGURE 4: Changes in Small Business Revenues by ZIP Code

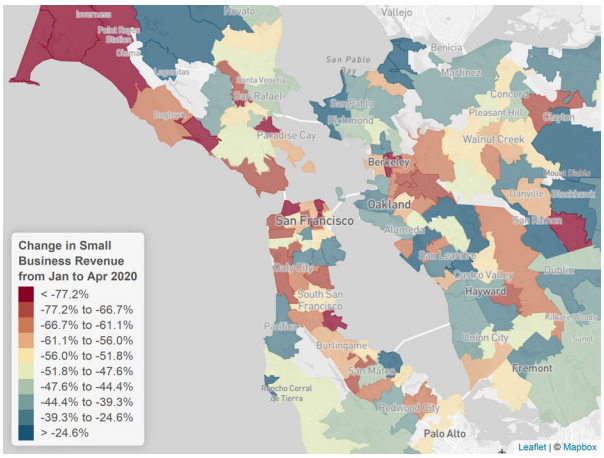
A. New York



B. Chicago

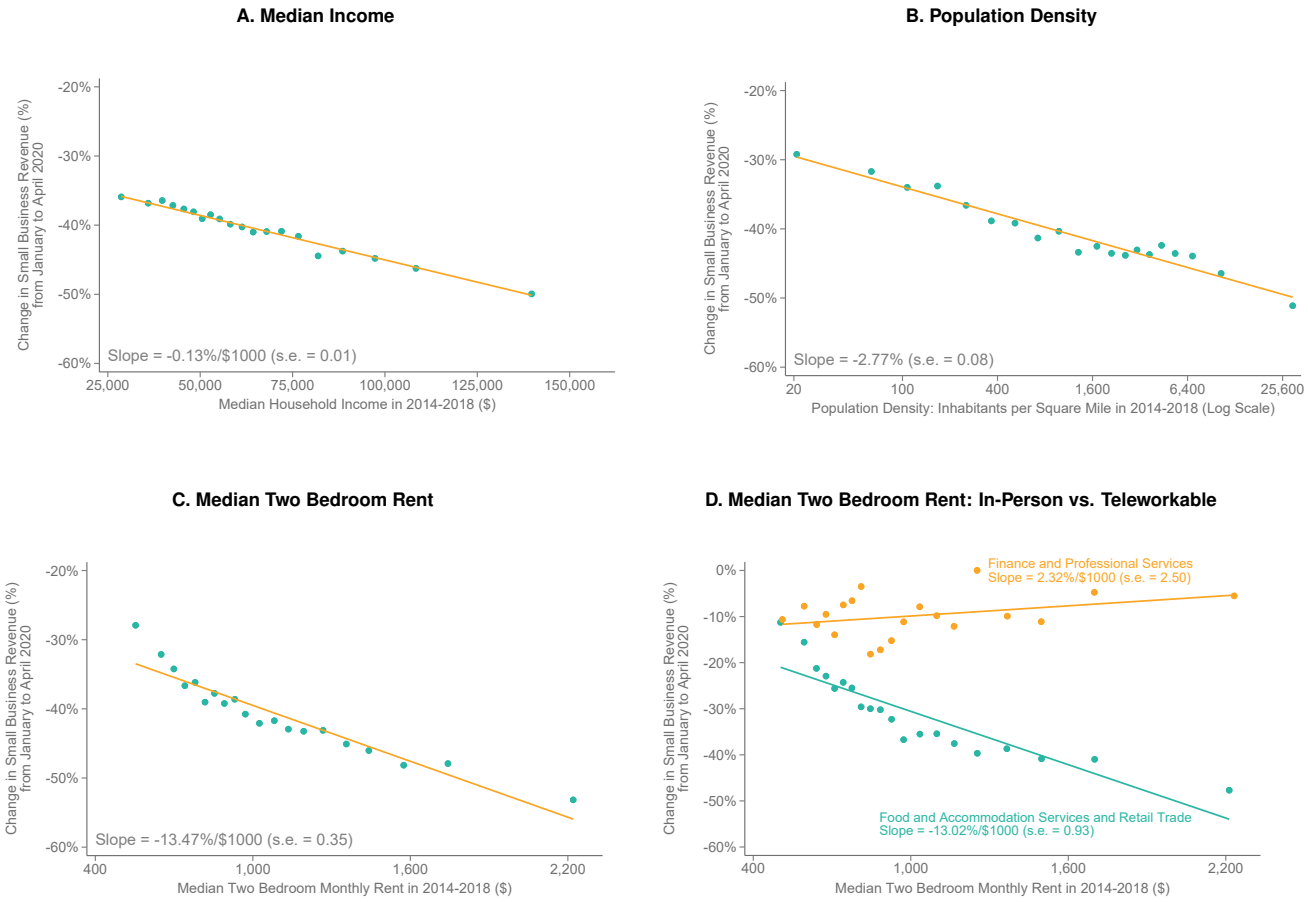


C. San Francisco



Notes: This figure plots seasonally-adjusted changes in small business revenue by ZIP code in the MSAs corresponding to New York-Newark-Jersey City, NY-NJ-PA (Panel A), Chicago-Naperville-Elgin, IL-IN-WI (Panel B), and San Francisco-Oakland-Hayward, CA (Panel C). The changes are measured during March 25 to April 14 relative to the period from January 4 to January 31. 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. We calculate the signal-to-noise ratio by regressing seasonally-adjusted normalized weekly revenue on an indicator variable for whether the week is after March 9, 2020, within each ZIP code, denoting the coefficient and standard error on this indicator variable in each ZIP code as β_z and SE_z, respectively. We then calculate the signal-to-noise ratio as $1 - (\sum SE_z^2/n)/\mathrm{Var}(\beta_z)$. The signal variance to total variance ratios for the panels are 0.79 (New York), 0.78 (Chicago), and 0.78 (San Francisco). These maps must be printed in color to be interpretable; dark red colors represent areas with larger revenue declines, while dark blue colors represent areas with smaller declines. Data source: Womply.

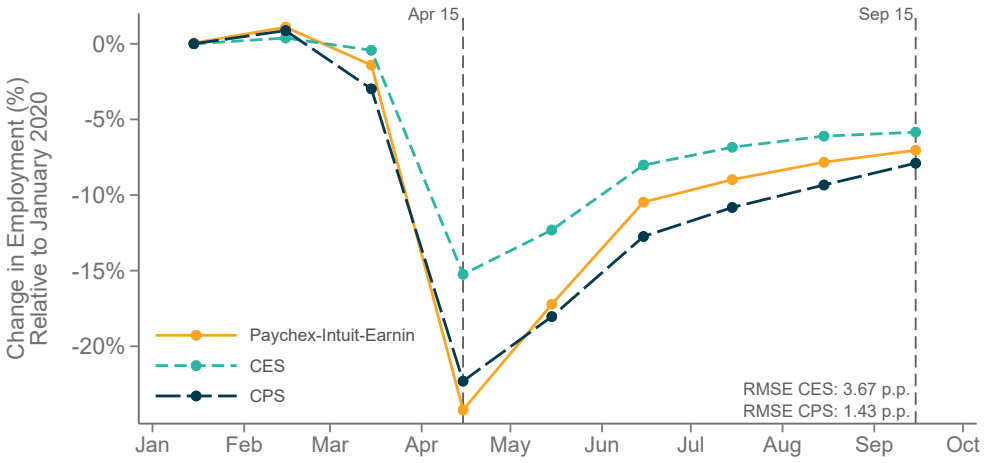
FIGURE 5: Changes in Small Business Revenues vs. ZIP Code Characteristics



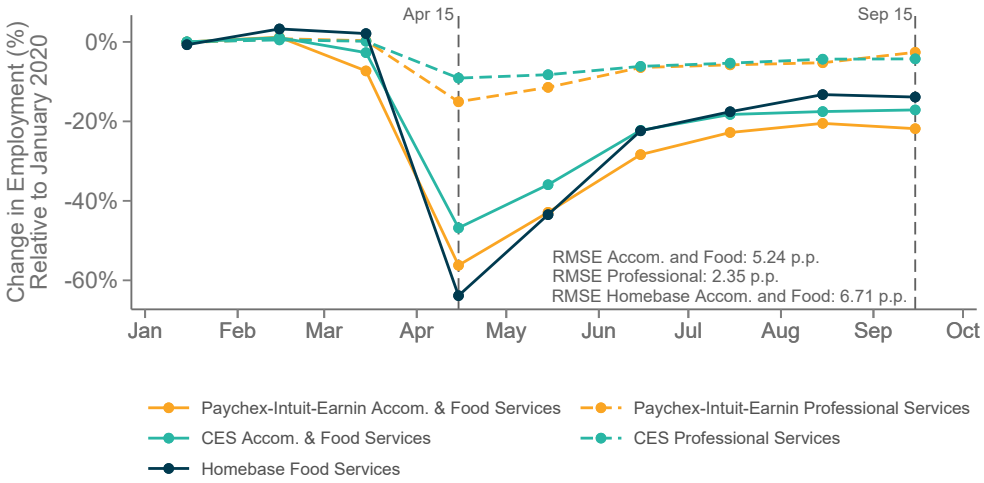
Notes: This figure presents binned scatter plots showing the relationship between changes in seasonally-adjusted small business revenue in Womply data vs. various ZIP code-level characteristics. The binned scatter plots are constructed as described in Figure 3. In each panel, the changes in small business revenue are measured during March 25 to April 14 relative to the period from January 4 to January 31 and seasonally adjusted as defined in Figure 4, winsorizing at the 99th percentile of the (population-weighted) ZIP-level distribution. We exclude data from ZIP codes in which changes are larger than 200% or where the variance of normalized revenue exceeds 900%. We also exclude ZIP code-by-industry cells with average weekly revenue of less than $4,250 during the base period of January 4 to January 31. 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. In Panel C, the x-axis variable is the ZIP code median rent for a two-bedroom apartment in the 2014-2018 ACS. Panel D replicates Panel C for two distinct sectors: in-person services defined as the combination of Food and Accommodation (NAICS 72) and Retail Trade (NAICS 44 and 45), and teleworkable services defined as Finance and Professional Services (NAICS 52 and NAICS 54). Data source: Womply.

FIGURE 6: Changes in Employment Rates Over Time

A. Pooling All Industries

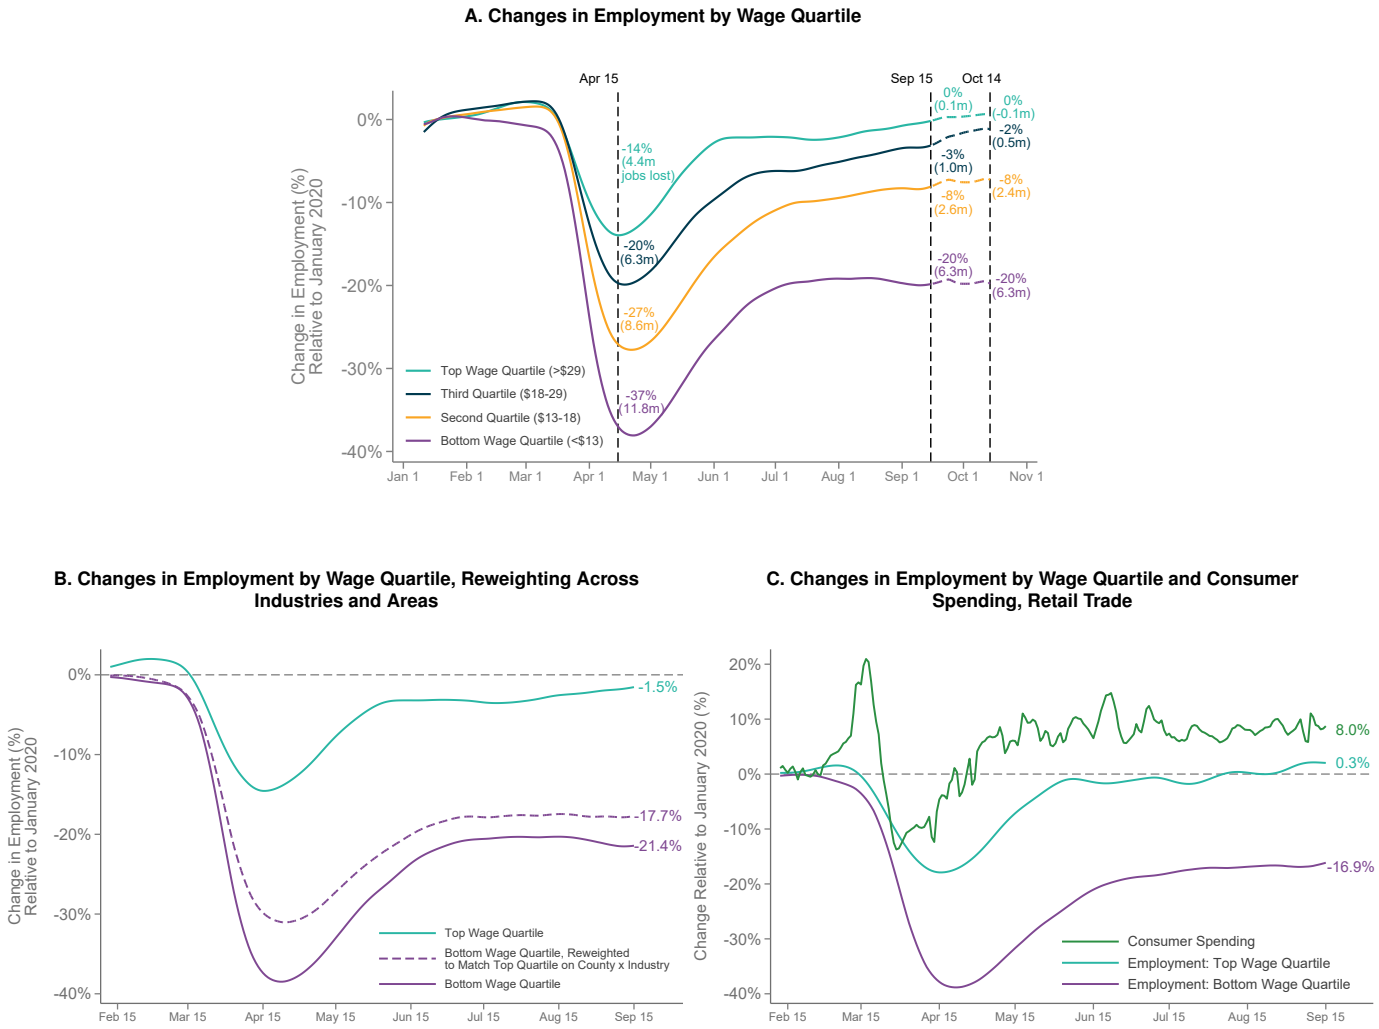


B. Accommodation and Food Services vs. Professional Services



Notes: This figure compares employment changes relative to January 2020 within various datasets. The combined Paychex-Intuit-Earnin data covers private non-farm employment in the United States, constructed as described in Section II.C. We plot the values of the combined Paychex-Intuit-Earnin employment series as of the 15th of each month, relative to the period January 4-31. The Current Employment Statistics (CES) and the Current Population Survey (CPS) data are available monthly, so we plot changes in each month relative to January 2020. The CES is a monthly survey of firms at the establishment level. The CPS is a monthly survey of households, which we then adjust to match a payroll definition of employment by accounting for multiple jobholders. The CES reports employment for the pay period including the 12th of each month, and the CPS is fielded during the week of the 19th of each month. In Panel B, we present employment series restricted to two specific sectors: Accommodation and Food Services (NAICS sector 72) and Professional and Business Services (NAICS supersector 60). Panel B also presents employment changes measured in the Homebase data at small businesses in Accommodation and Food Services (NAICS sector 72). Data sources: Paychex, Intuit, Earnin, Homebase.

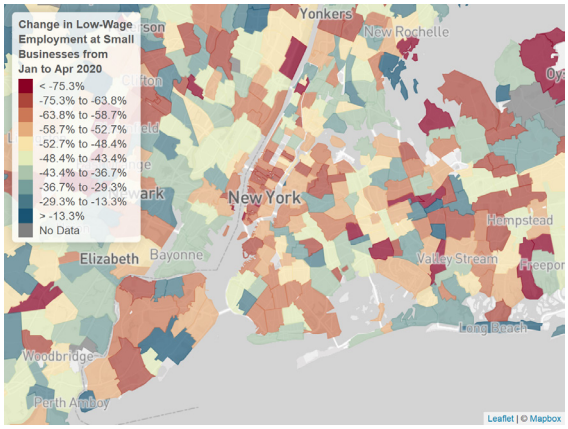
FIGURE 7: Changes in Employment by Wage Quartile



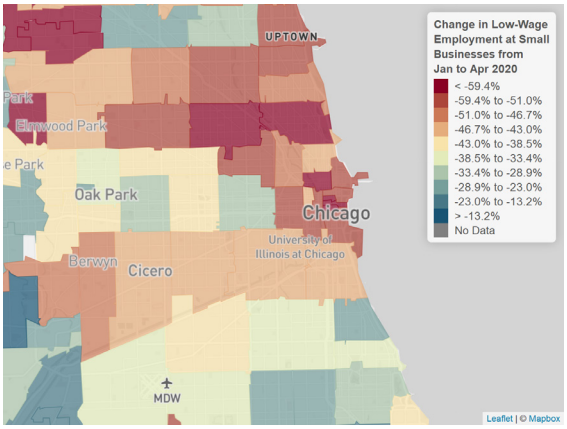
Notes: This figure plots changes in employment by wage quartile relative to January 4-31 2020. In each panel, we show a daily series of private non-farm employment constructed by combining Paychex, Intuit and Earnin data, as described in Section II.C. We separate the sample into wage quartiles based on fixed thresholds of the 2019 hourly wage distribution. In Panel A, the solid portion of each line represents the combined Paychex-Intuit-Earnin data, while the dashed portion of the line is forecasted using Kronos data and Paychex data from firms with weekly paycycles. To construct this forecast, we regress the combined Paychex-Intuit-Earnin series on the de-seasonalized Kronos series for the same date (t), the Paychex weekly series for the same date (t), and the Paychex weekly series for three prior weeks $(t-7), (t-14), (t-21)$. We then use the resulting coefficients to predict the value of combined Paychex-Intuit-Earnin employment before the combined data is available. In Panel B, we reweight Paychex-Intuit-Earnin data to examine whether differences in employment trends between the top and bottom wage quartile are explained by differences in industry and geographic composition. We restrict the sample in Panel B to county x industry (2-digit NAICS) cells which have nonzero first-quartile and fourth-quartile employment in the period January 4-31 2020; this sample restriction excludes 0.9% of worker-days from the sample. We then calculate the daily change in employment since January 4-31 2020 in each county x industry x income quartile cell, winsorizing at the 99th percentile (weighted by total employment in the period January 4-31 2020). For the top and bottom wage quartiles, we estimate overall daily employment by taking the weighted mean of employment changes in every county x industry cell, weighting by the level of employment in January 4-31. These series are similar but not identical to Panel A because of the sample restriction and winsorization. Finally, we estimate bottom quartile employment reweighted to match the industry and geographic composition of top quartile employment by performing the same procedure, but weighting bottom quartile employment changes by top quartile employment levels in January 4-31 in each cell. In Panel C, we restrict to the retail trade sector (NAICS 44-45) and present the change relative to January 4-31 in the Paychex-Intuit-Earnin employment series and the Affinity consumer card spending series. Data sources: Paychex, Intuit, Earnin, Kronos, Affinity Solutions.

FIGURE 8: Changes in Low-Wage Employment Rates by ZIP Code

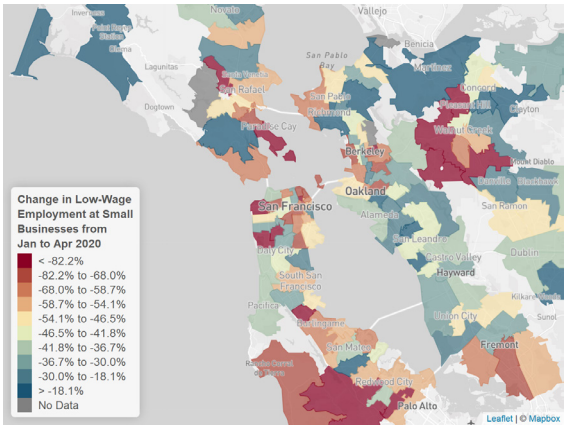
A. New York



B. Chicago

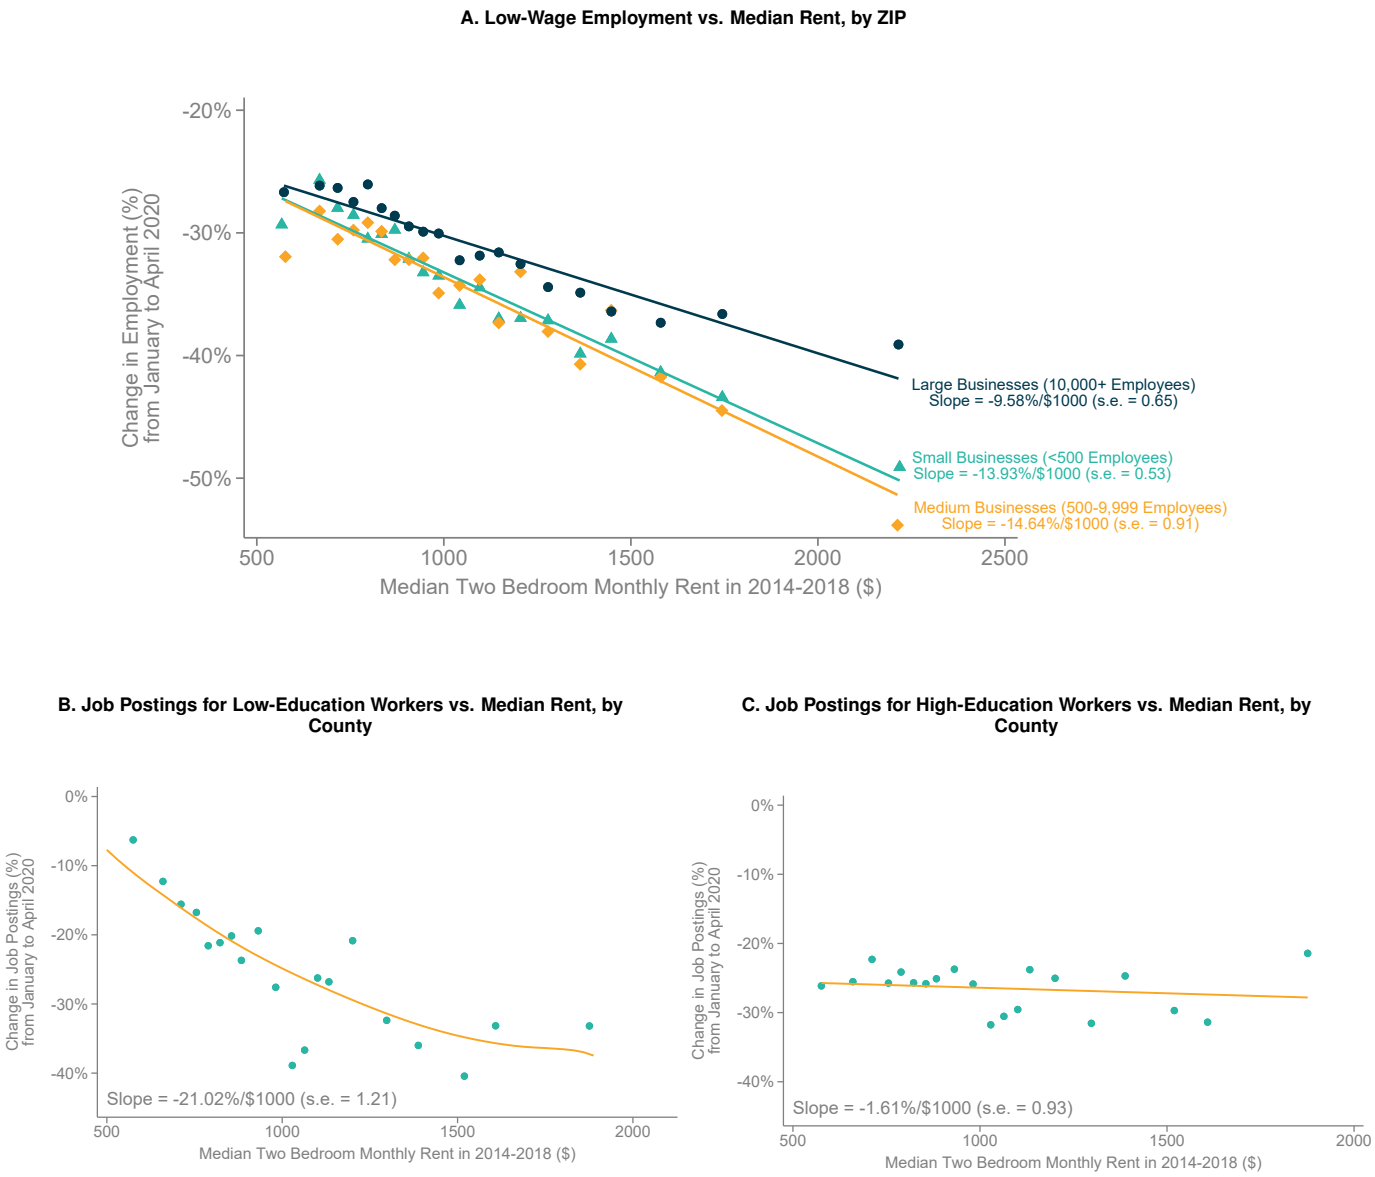


C. San Francisco



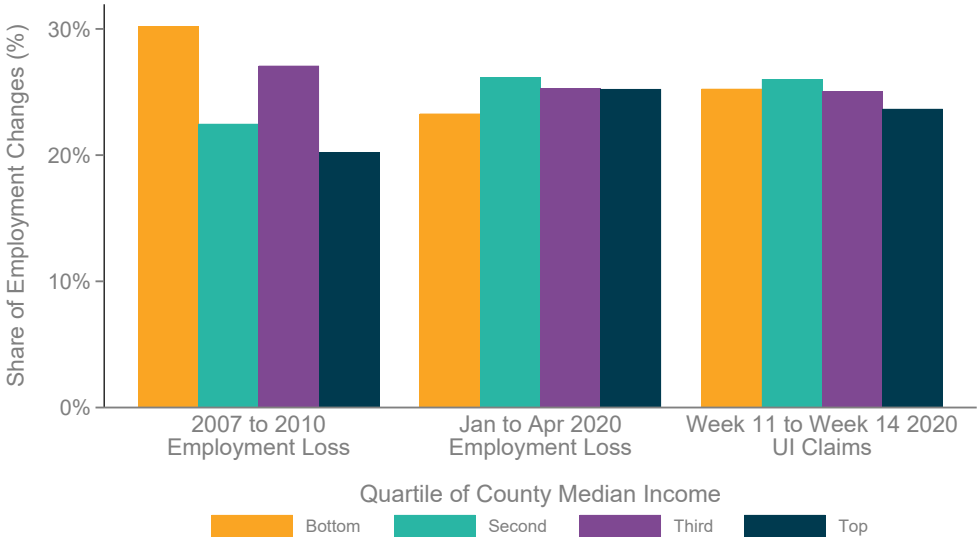
Notes: This figure replicates Figure 4 using Earnin data on changes in employment among low-wage workers, plotted by employer ZIP code. We focus on small and medium-sized businesses, defined as firms with at most 500 employees, measured using linked data from ReferenceUSA. For users whose employer cannot be linked to ReferenceUSA data on firm sizes, we restrict to users whose employer is in the fourth decile or below of firms in the Earnin data, in terms of number of Earnin users working for the firm; the median firm size for the fourth decile of Earnin employers is roughly 300 employees (among employers matched to ReferenceUSA data). We measure the change in employment as total average weekly employment during the period of April 8-28 divided by total average weekly employment in the period of January 4-31, 2020. We calculate the signal-to-noise ratios as in Figure 4; these ratios are 0.79 in New York, 0.59 in Chicago, and 0.67 in San Francisco. These maps must be printed in color to be interpretable; dark red colors represent areas with larger employment declines, while dark blue colors represent areas with smaller declines. Data source: Earnin.

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



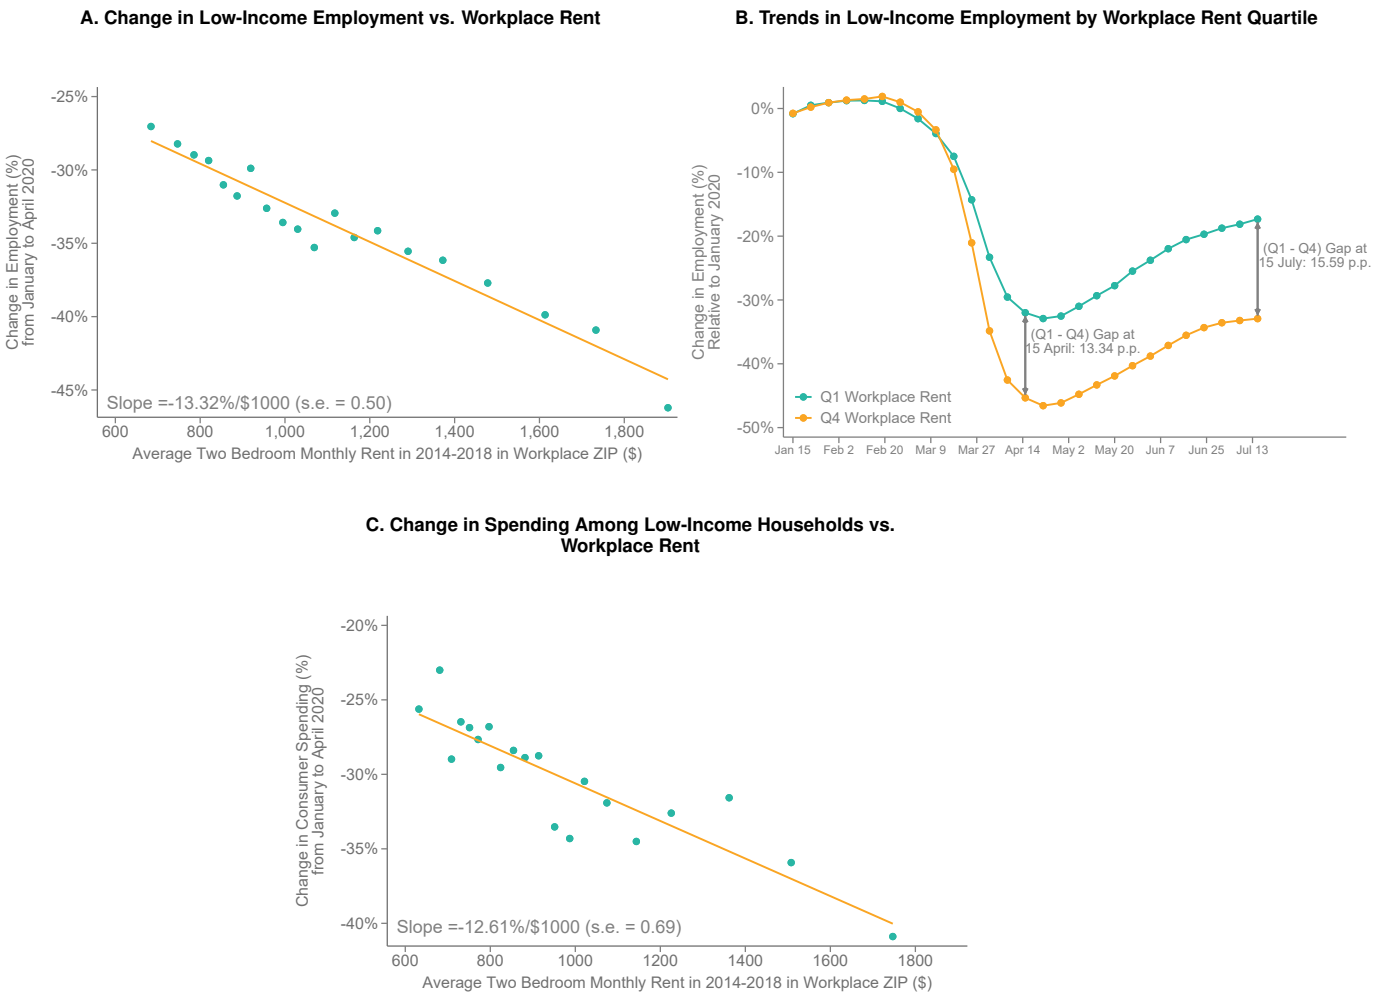
Notes: This figure shows binned scatter plots of the relationship between median rents and changes in employment rates (Panel A) or changes in job postings (Panels B and C). The binned scatter plots are constructed as described in Figure 3. In each panel, the x-axis variable is the median rent (within a county or ZIP code) for a two-bedroom apartment in the 2014-2018 ACS. Panel A presents the change in employment rates among low-income workers from January 4-31 to April 8-28 2020, measured by employer ZIP code separately by firm size in Earnin data linked to ReferenceUSA firm sizes. We winsorize at the 99th percentile of (population-weighted) ZIP-level changes in employment for each firm size category. Panel B presents the county-level change in job postings for workers with minimal or some education from January 8-28 to March 25-April 14 2020 in the Burning Glass data. Panel C replicates Panel B with job postings for workers with moderate, considerable, or extensive education. To construct Panels B and C, we compute the change in job posts in each county x worker education cell and winsorize at the 99th percentile of (population-weighted) county-level changes in employment within each education group. Solid lines are best-fit lines estimated using OLS, except in Panel B, where we use a lowess fit. Each panel also displays the slope coefficient and standard error of the corresponding linear OLS regression. Data sources: Earnin, Burning Glass Technologies.

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



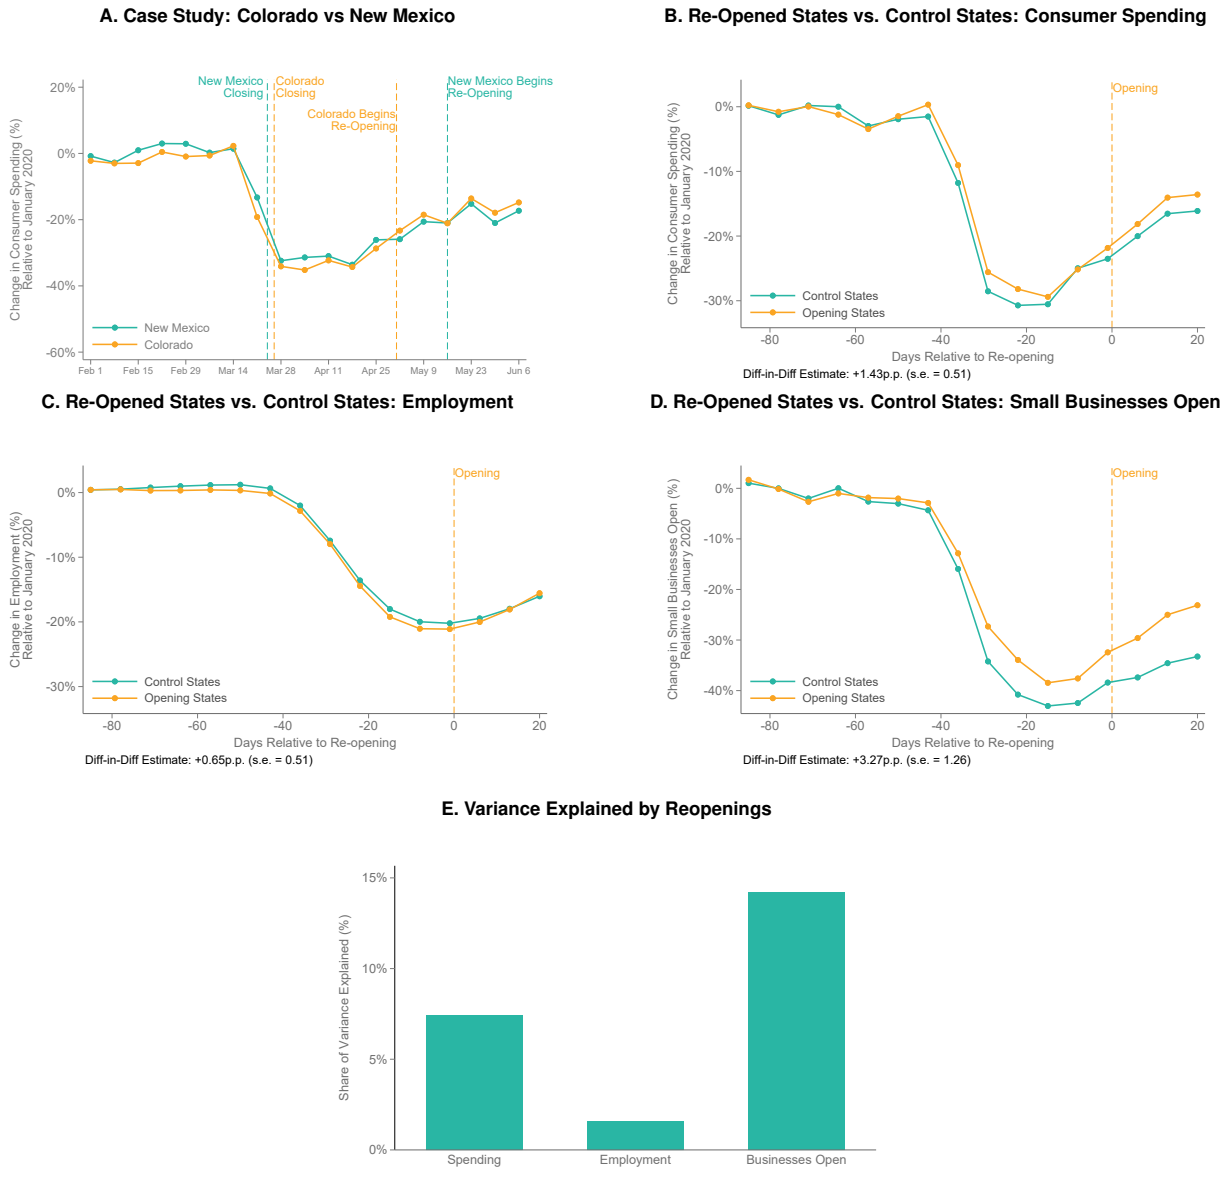
Notes: This figure displays the share of job losses occurring in low vs. high income counties during the Great Recession and the COVID Recession. We split counties into (population-weighted) quartiles by median household income in the 2006 ACS for the Great Recession (Panel A) and the 2014-2018 ACS for the COVID Recession (Panels B and C). 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 2020 to April 2020. The third set of bars reports the share of total initial UI claims within each county income quartile between March 15 (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 comprise 53% of the U.S. population.

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



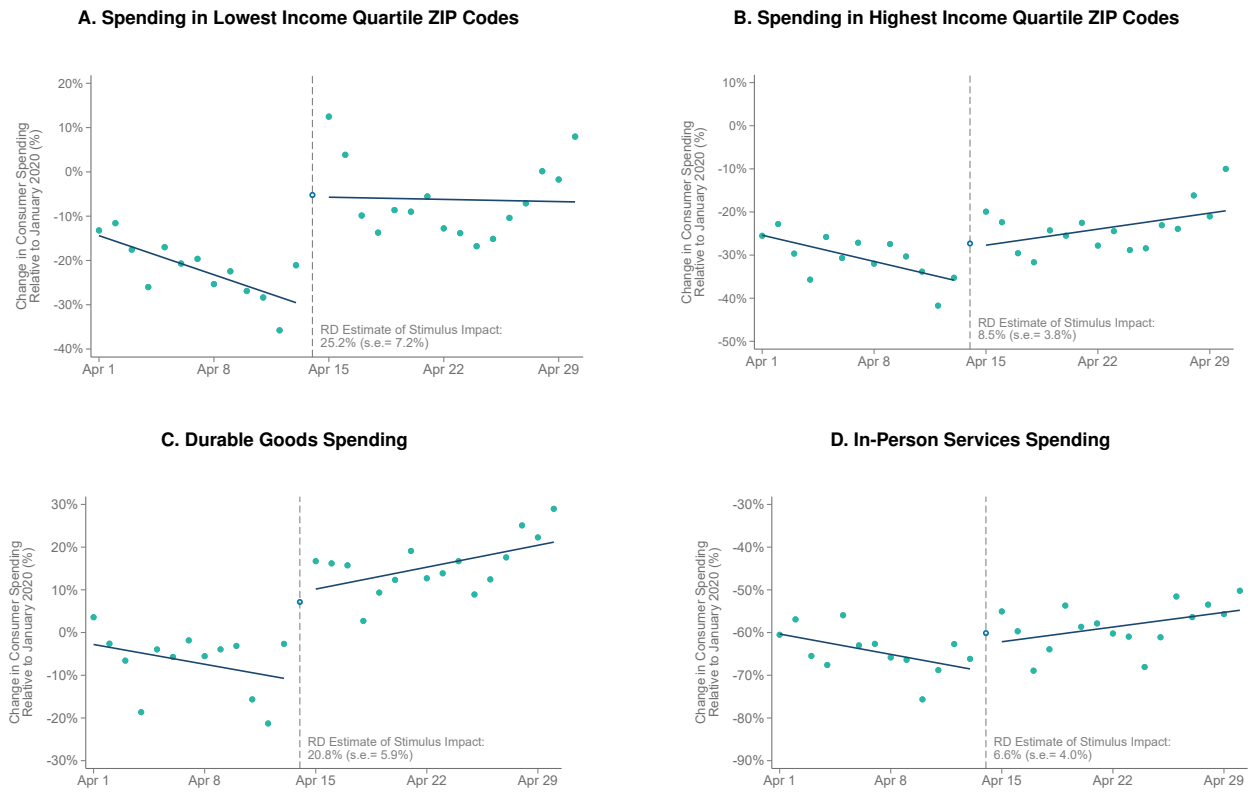
Notes: This figure compares changes in low-income employment (Panels A and B) and seasonally-adjusted consumer spending (Panel C) by home ZIP code to average rent in the *workplace* ZIP codes of low-income workers who live in each home ZIP code. We construct the average workplace rent for each home ZIP code in two steps. First we measure the distribution of workplace ZIP codes for low-income workers (earning below $1,250 per month) using the matrix of home ZIP codes by workplace ZIP codes in the Census LEHD Origin-Destination Employment Statistics (LODES). Then, for each home ZIP code, we use these distributions to construct a weighted mean over the median rent of a two-bedroom apartment in *workplace* ZIP codes (measured in the 2014-2018 ACS), weighting by the share of low-income workers in each workplace ZIP code. Panel A presents the change in employment rates among low-income workers from January 4-31 to April 8-28 2020, measured by home ZIP code in Earnin data, winsorizing at the 99th percentile of (population-weighted) ZIP-level changes in employment. Panel B uses the same Earnin data on employment rates by home ZIP code to present the trends over time in employment of low-income workers relative to the base period of January 4-31, separately for home ZIP codes in the top and bottom (population-weighted) quartiles of workplace rent. Panel C replicates Panel A using the change in seasonally-adjusted consumer spending from January 4-31 to March 25-April 14, measured by home ZIP code in the Affinity Solutions data and restricted to ZIP codes in the bottom quartile of the household income distribution. The binned scatter plots in Panels A and C are constructed as described in Figure 3. Data sources: Earnin, Affinity Solutions.

FIGURE 12: Effects of Reopenings on Economic Activity: Event Studies



Notes: This figure analyzes the causal effects of state reopenings. Panel A presents a time series plot of the change in seasonally-adjusted consumer spending relative to the base period of January 4-31 for New Mexico and Colorado. Colorado partially reopened non-essential businesses on May 1, while New Mexico did not do so until May 16. Panel B presents an event study plot of the same outcome variable, showing average spending changes in five states (SC, AK, GA, MN and MS) that partially reopened non-essential businesses between April 20 and April 27. Each reopening state is matched to multiple control states (listed in Appendix Table 6) 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 reopening day and then stack the resulting event studies to align the events. Panel C replicates Panel B with the change in employment relative to January 4-31 measured using combined Paychex-Intuit-Earnin data (as described in Section II.C). Panel D replicates Panel B with the seasonally-adjusted change in small businesses open measured using Womply data. In Panels B to D, we report the coefficient from a difference-in-differences regression comparing treated vs. untreated states in the two weeks following vs. the two weeks prior to the partial reopening (also reported in Table 3). Panel E reports the share of variance in outcomes explained by reopenings as of May 18. 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-difference 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 May 18). We then recalculate the variance in this counterfactual in which all states had reopened. The share of variance explained by reopenings for each outcome is defined as 1-(counterfactual variance/actual variance). Data sources: Affinity Solutions, Paychex, Intuit, Earnin, Womply.

FIGURE 13: Effect of Stimulus Payments on Consumer Spending: Regression Discontinuity Estimates



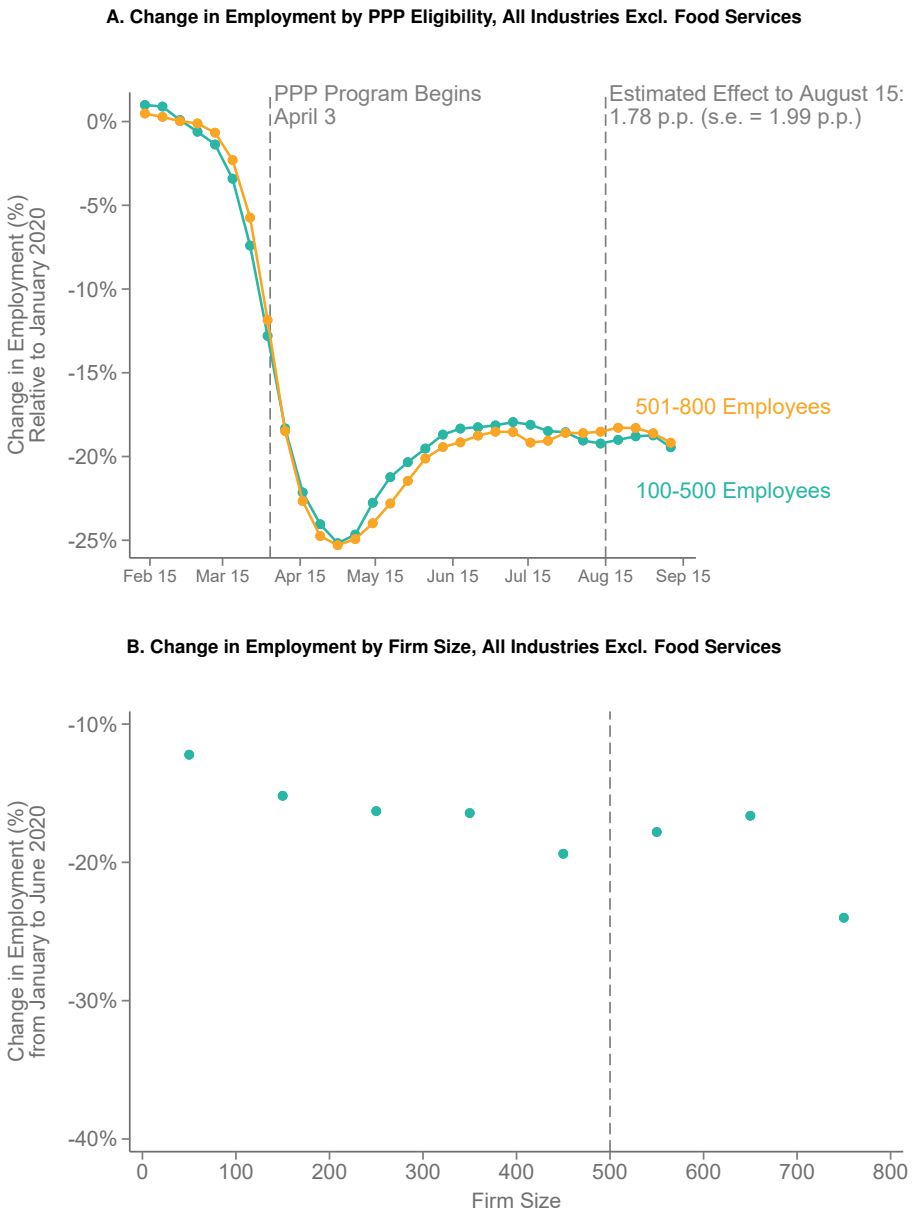
Notes: This figure analyzes the effect of the CARES Act stimulus payments deposited on April 15, 2020 on card spending by plotting daily spending levels in the Affinity data for the month of April. In each panel, we first calculate spending on each day relative to mean spending over January 4-31, and then residualize daily spending with respect to day of week and first of the month fixed effects, which we estimate using data from January 1 2019 to May 10 2019. Each panel reports regression discontinuity estimates of the jump in spending on April 15, using a linear control function before and after April 15 (shown by the solid best fit lines), excluding the partially treated day of April 14, shown by the hollow point and demarcated by the dashed vertical line. Panel A restricts the sample to cardholders living in ZIP codes in the lowest quartile of (population-weighted) median household income measured using the 2014-2018 ACS. Panel B restricts the sample to cardholders in highest income quartile ZIP codes. Panels C and D pool all cardholders and restrict to spending on durable goods (Panel C) and in-person services (Panel D), as defined in Appendix B. Data source: Affinity Solutions.

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



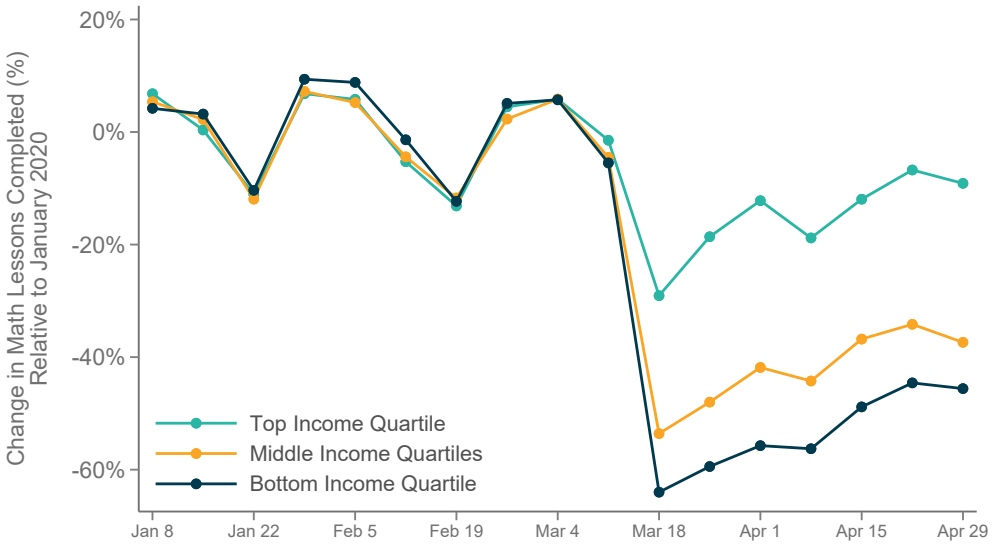
Notes: This figure analyzes the effect of the CARES Act stimulus payments deposited on April 15, 2020 on small business revenues and low-wage employment. In Panels A and B, we first calculate small business revenue on each day relative to mean spending over January 4-31, and then residualize daily revenue with respect to day of week and first of the month fixed effects, which we estimate using data from January 1 2019 to May 10 2019. Panels A and B report regression discontinuity estimates of the jump in spending on April 15, using a linear control function before and after April 15 (shown by the solid best fit lines), excluding the partially treated day of April 14, shown by the hollow point and demarcated by the dashed vertical line. Panel A restricts the sample to businesses in ZIP codes in the lowest quartile of (population-weighted) median two-bedroom rent measured using the 2014-2018 ACS. Panel B restricts the sample to cardholders in highest rent quartile ZIP codes. Panel C plots the 7-day moving average of the changes in seasonally-adjusted small business revenue from Womply data and low-income employment at small businesses from Earnin data. The employment data is restricted to businesses whose parent firm has at most 500 employees (measured in ReferenceUSA data linked to Earnin data). Data sources: Womply, Earnin.

FIGURE 15: Effects of Paycheck Protection Program on Employment



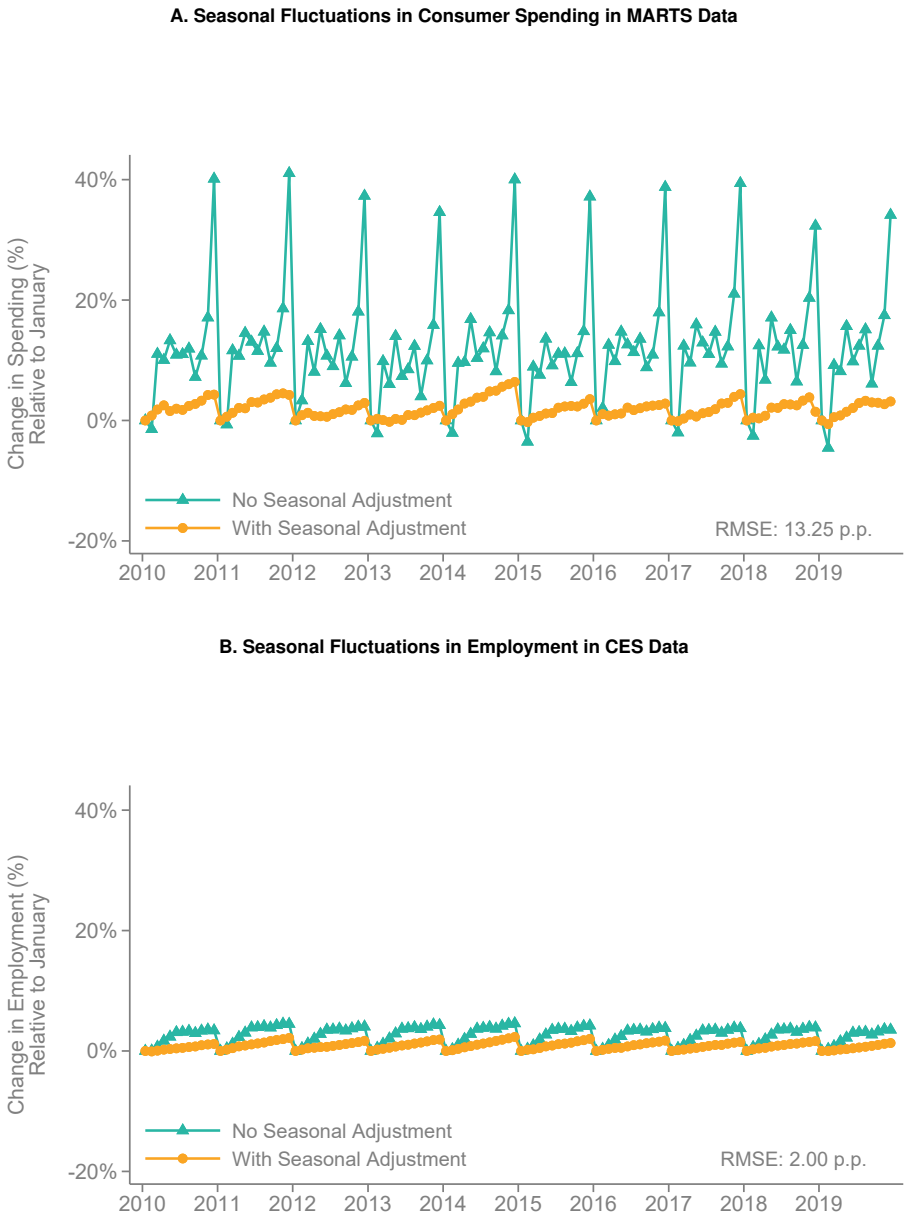
Notes: This figure analyzes the effects of the Paycheck Protection Program on employment using the threshold in eligibility at 500 employees. We pool all industries except Accommodation and Food Services (NAICS 72), which was subject to different eligibility rules (discussed in Section IV.C). Panel A compares employment trends measured in Paychex and Earnin data among firms with 100-500 employees (generally eligible for PPP loans) to firms with 501-800 employees (generally ineligible for PPP loans). To construct these employment trends, we begin by calculating weekly employment changes relative to January 4-31 2020 disaggregated by data source, county, industry (2-digit NAICS), wage quartile and firm size bin. We reweight these cells so that the composition in each firm size bin matches the pooled distribution of industry and data source over the period January 4-31 2020. We plot the "control" series (firms with 501-800 employees) directly as the mean weekly value of the reweighted employment series. We plot the "treated" series (firms with 100-500 employees) as the sum of the control series and the weekly difference between control and treated firms after residualizing on interacted county and wage quartile fixed effects. For visual clarity, we recenter each series so the mean change in employment is 0% over the pre-period (February 12 to March 18). The difference between these two series corresponds to the regression estimate in Column 1 of Table 5, which is also reported in the figure. Panel B presents a binned scatter plot of changes in reweighted employment (defined as in Panel A) from January 4-31 to June 1-23 vs. firm size. To construct changes in employment by firm size, we first classify firms in bins of size 50 according to their parent employer size. Then we calculate the mean change in employment among firms in each bin and plot this mean change against the midpoint in each bin. Data sources: Paychex, Earnin.

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



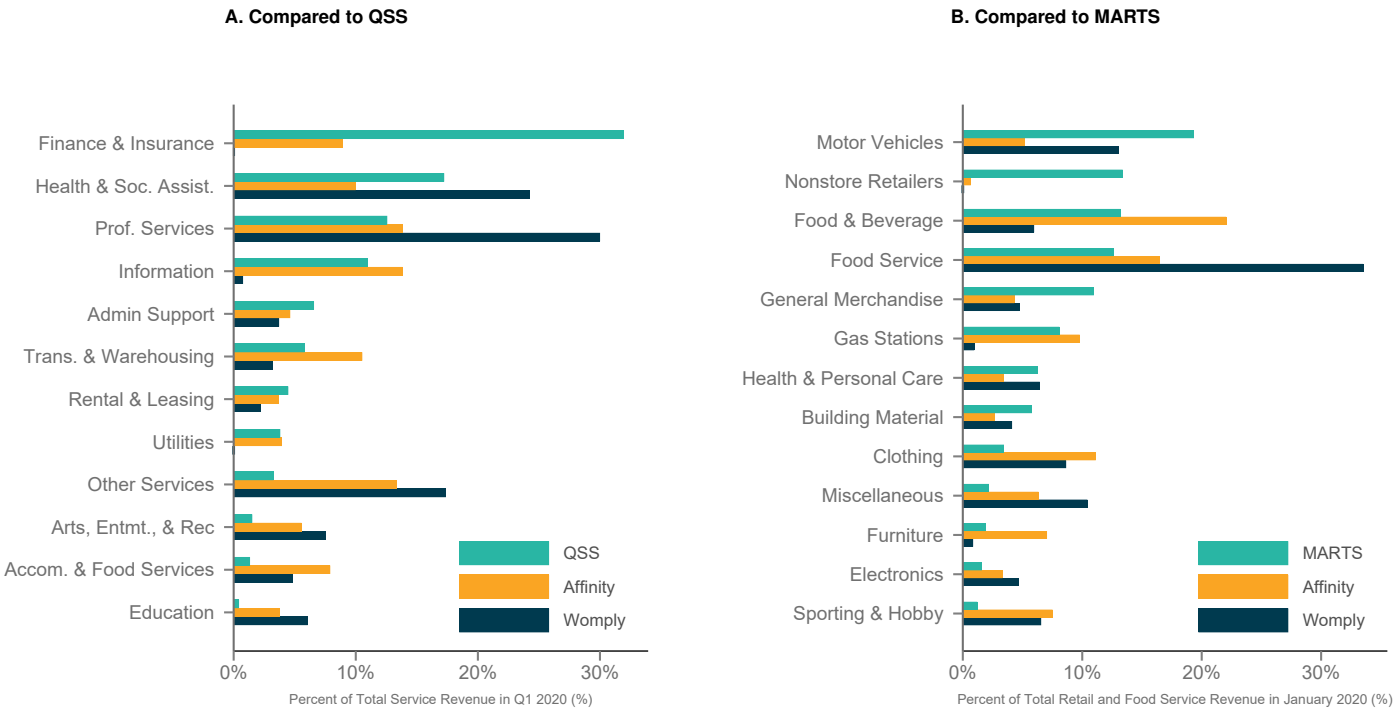
Notes: This figure plots a time series of students' educational progress on the Zearn Math online platform, splitting schools into quartiles based on the share of students eligible for Free and Reduced Price Lunch. We measure educational progress as the number of accomplishment badges earned in Zearn Math in each week, relative to the mean value of badges earned during the reference period of January 6 to February 7 2020. The sample is restricted to classes with more than 10 students using Zearn during the reference period and at least five users in every week during the reference period. We measure the share of students eligible for Free and Reduced Price Lunch in each school using demographic data from the Common Core data set from MDR Education, a private education data firm. Data source: Zearn.

APPENDIX FIGURE 1: Seasonal Fluctuations in Consumer Spending vs. Employment



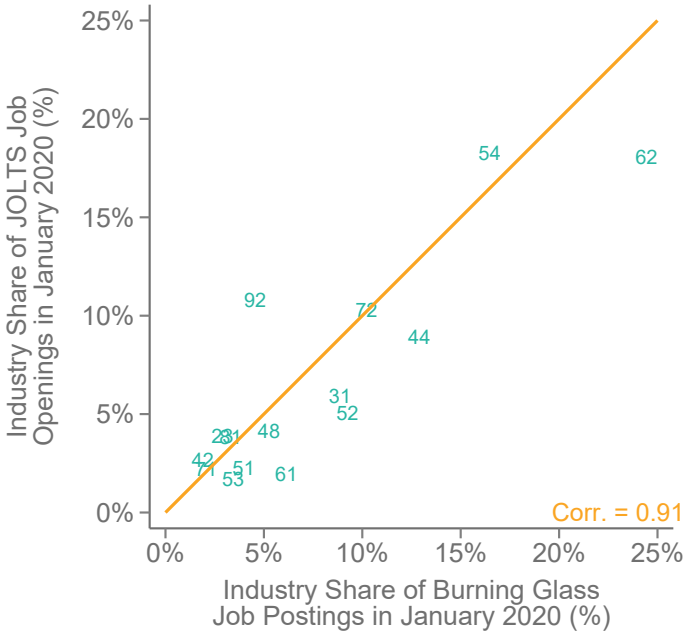
Notes: This figure compares seasonal fluctuations in Advance Monthly Retail Trade Survey (MARTS) data on consumer spending on retail sales and food services (excluding motor vehicle and gas) vs. Current Employment Statistics (CES) data on private sector non-farm employment. Panel A shows seasonal fluctuations 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, expressed in percentage terms. Panel B replicates Panel A using CES data on private sector, non-farm employment.

APPENDIX FIGURE 2: Industry Shares of Consumer Spending and Business Revenues Across Datasets



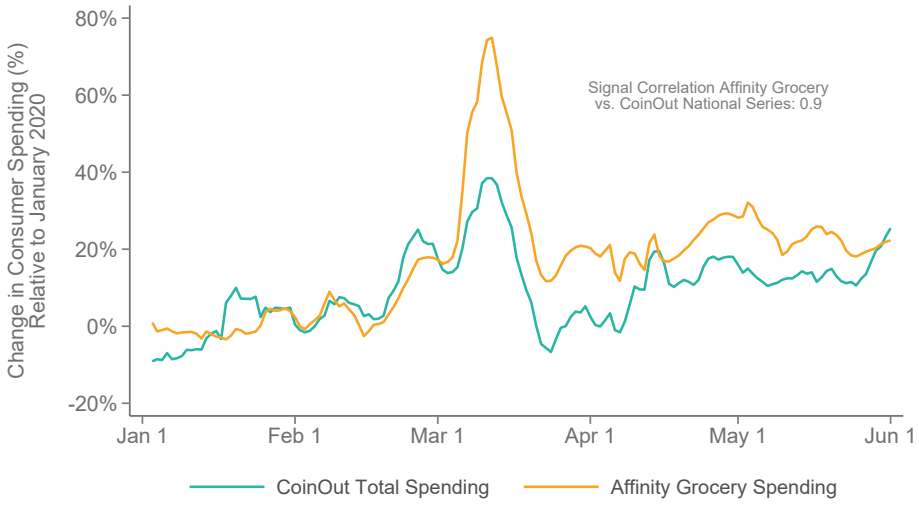
Notes: This figure compares the industry composition of spending in private sector datasets to the industry composition of spending in representative survey datasets. Panel A shows the NAICS two-digit industry mix for transactions in the Affinity Solutions and Womply datasets compared with the Quarterly 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 three-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.

APPENDIX FIGURE 3: Industry Shares of Job Postings in Burning Glass and Job Openings in Job Openings and Labor Turnover Survey (JOLTS)



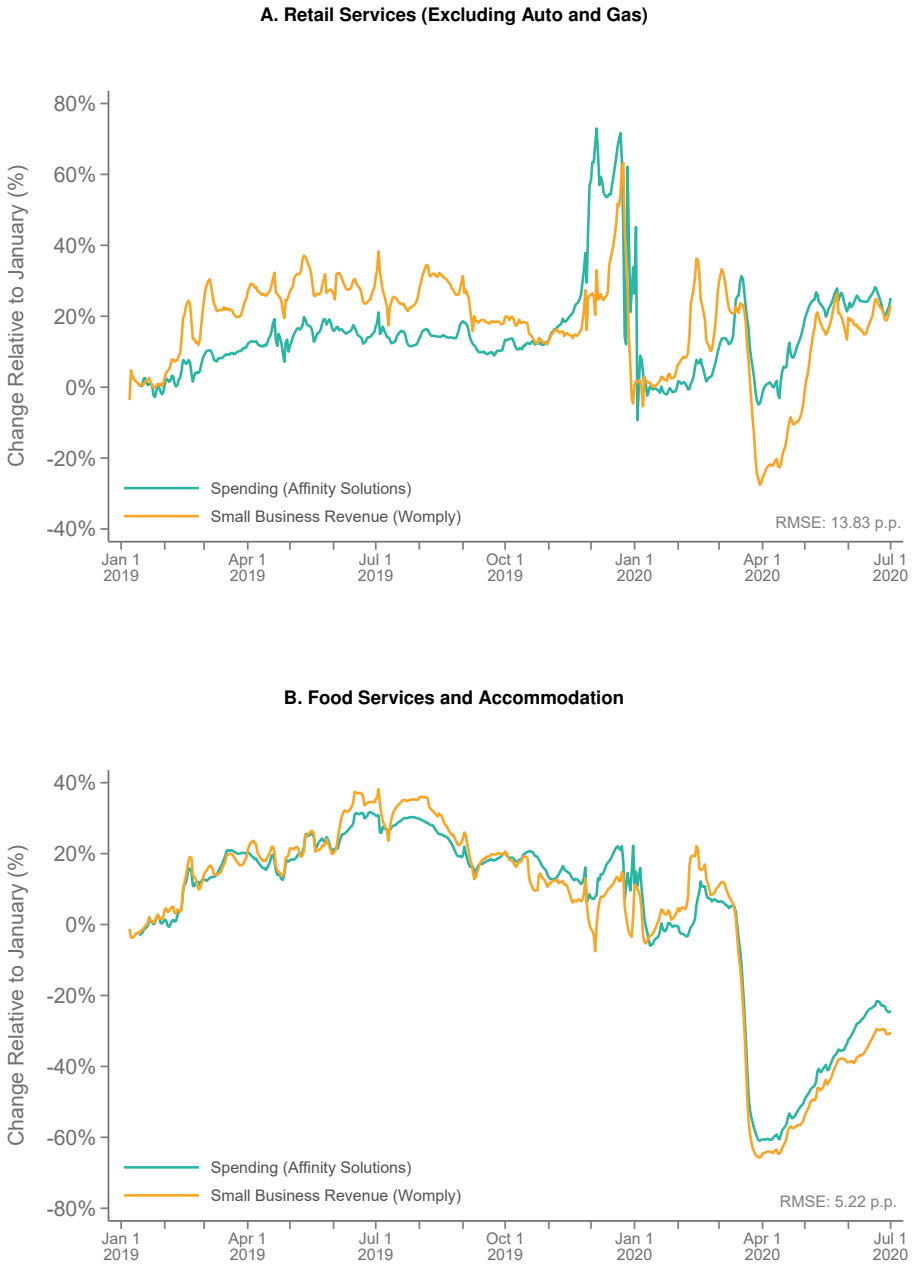
Notes: This figure presents a scatter plot showing the industry share of each 2-digit NAICS code of job postings in the Job Openings and Labor Turnover Survey (JOLTS) data in January 2020 vs. the corresponding industry share in job postings in Burning Glass data in January 2020. The solid line is a 45 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. Burning Glass data in January 2020, excluding NAICS 92 (Public Administration), and weighting according to total job openings in each NAICS code in JOLTS in January 2020. Data source: Burning Glass.

APPENDIX FIGURE 4: Cash Spending in CoinOut Transactions Data vs. Card Spending



Notes: This figure compares 7-day moving averages of national trends in cash transactions in CoinOut data vs. card spending on groceries in Affinity Solutions data between January 1 2020 and June 1 2020. The signal correlation between the two datasets at the national level is 0.90 at the weekly level. We compute this correlation by collapsing both datasets to the national weekly level, where values in each week are expressed as the percentage change from the January average. To adjust for measurement error at the weekly level, we calculate the series-specific reliability as the week-on-week correlation within each dataset. We then divide the raw weekly correlation between datasets by the square root of the product of the reliabilities to get the signal correlation. Data sources: CoinOut, Affinity Solutions.

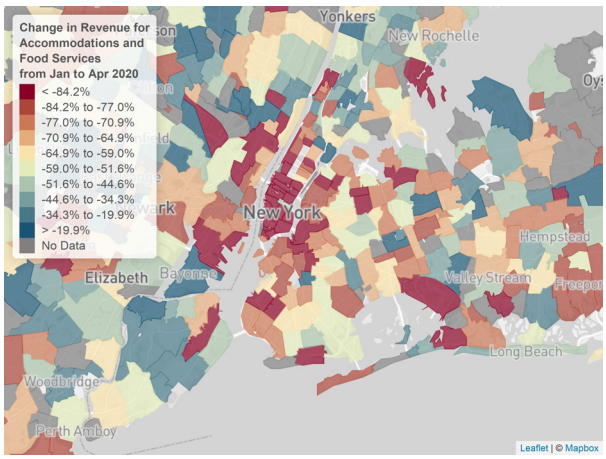
APPENDIX FIGURE 5: Small Business Revenue Changes vs. Consumer Spending Changes



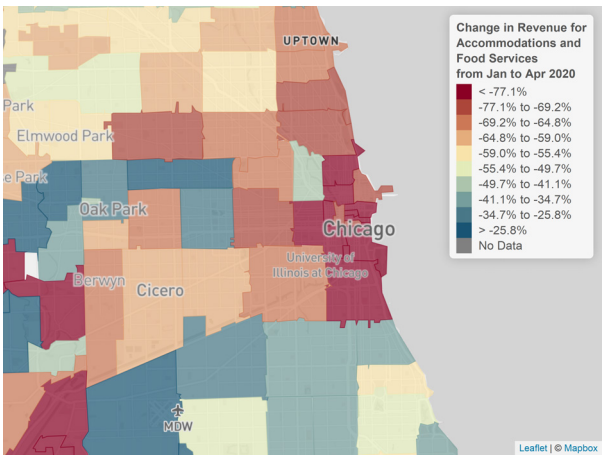
Notes: This figure compares seven-day moving averages of total consumer spending (from Affinity Solutions data) and small business revenue (from Womply data) for the period January 1 2019 to June 30 2020. Each series is expressed as a percentage change relative to the January 4-31 level in each calendar year. 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. Following the sectors defined in the Advance Monthly Retail Trade Survey (MARTS), Panel A restricts to specifically retail trade sectors (NAICS code 44-45) excluding motor vehicles (NAICS code 441) and gas (NAICS code 447), and Panel B restricts specifically to food services and accommodation (NAICS code 72). The bottom right corner of each panel displays the root mean squared error (RMSE) corresponding to the difference between the two lines. Data sources: Affinity Solutions, Womply.

APPENDIX FIGURE 6: Changes in Small Business Revenues by ZIP Code for Food and Accommodation Service Businesses

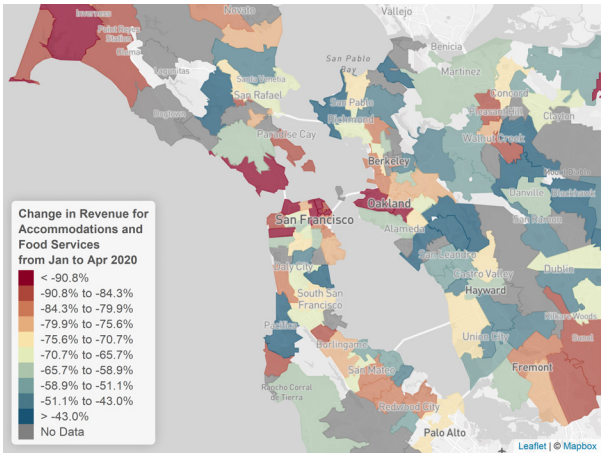
A. New York City



B. Chicago



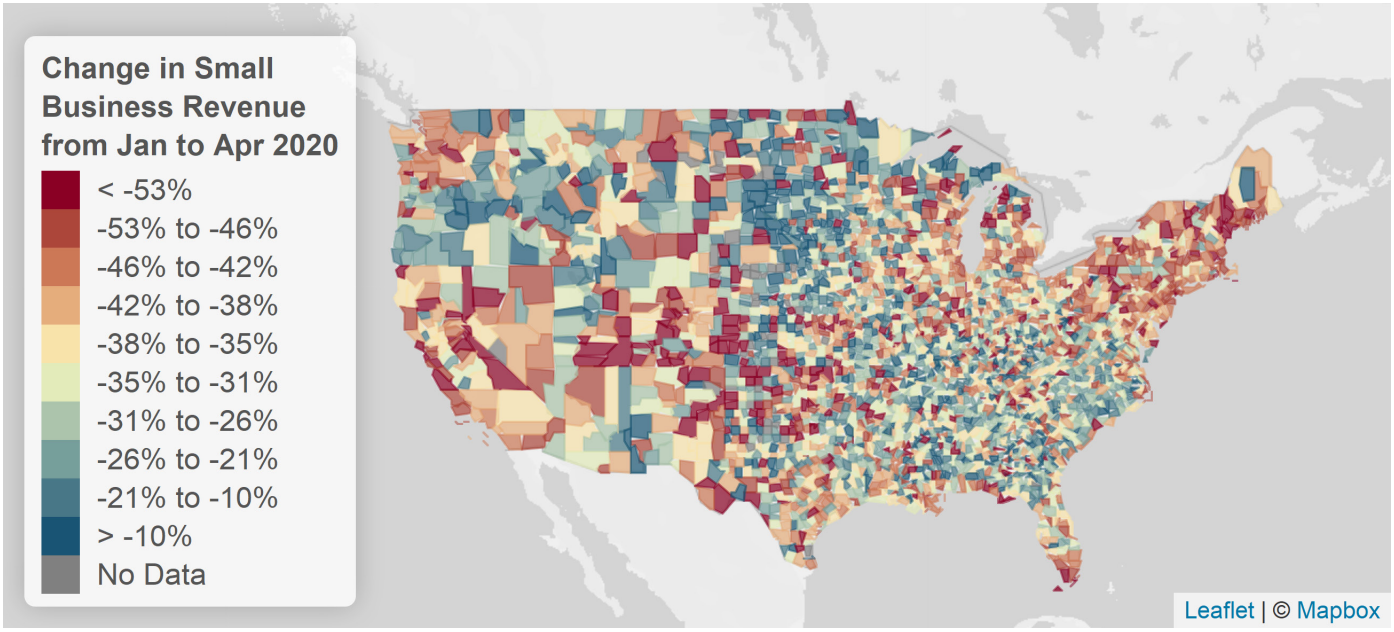
C. San Francisco



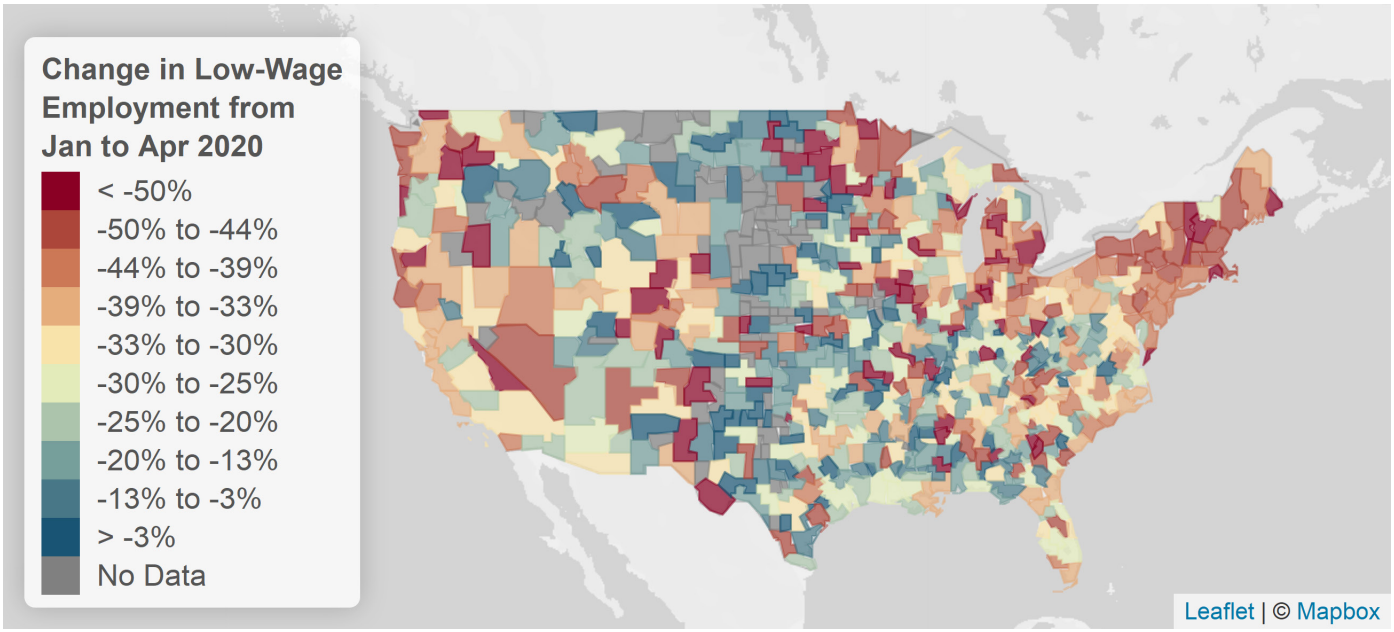
Notes: This figure replicates Figure 4 for small businesses in the food and accommodation service sector (NAICS 72), showing the change in revenue levels by ZIP code from January 4-31 to March 25-April 14. For further details, see the notes to Figure 4. The signal variance to total variance ratios for the panels are 0.83 (New York), 0.88 (Chicago), and 0.69 (San Francisco). Data source: Womply.

APPENDIX FIGURE 7: National Maps of Changes in Small Business Revenues and Low-Income Employment

A. Changes in Small Business Revenues, by County

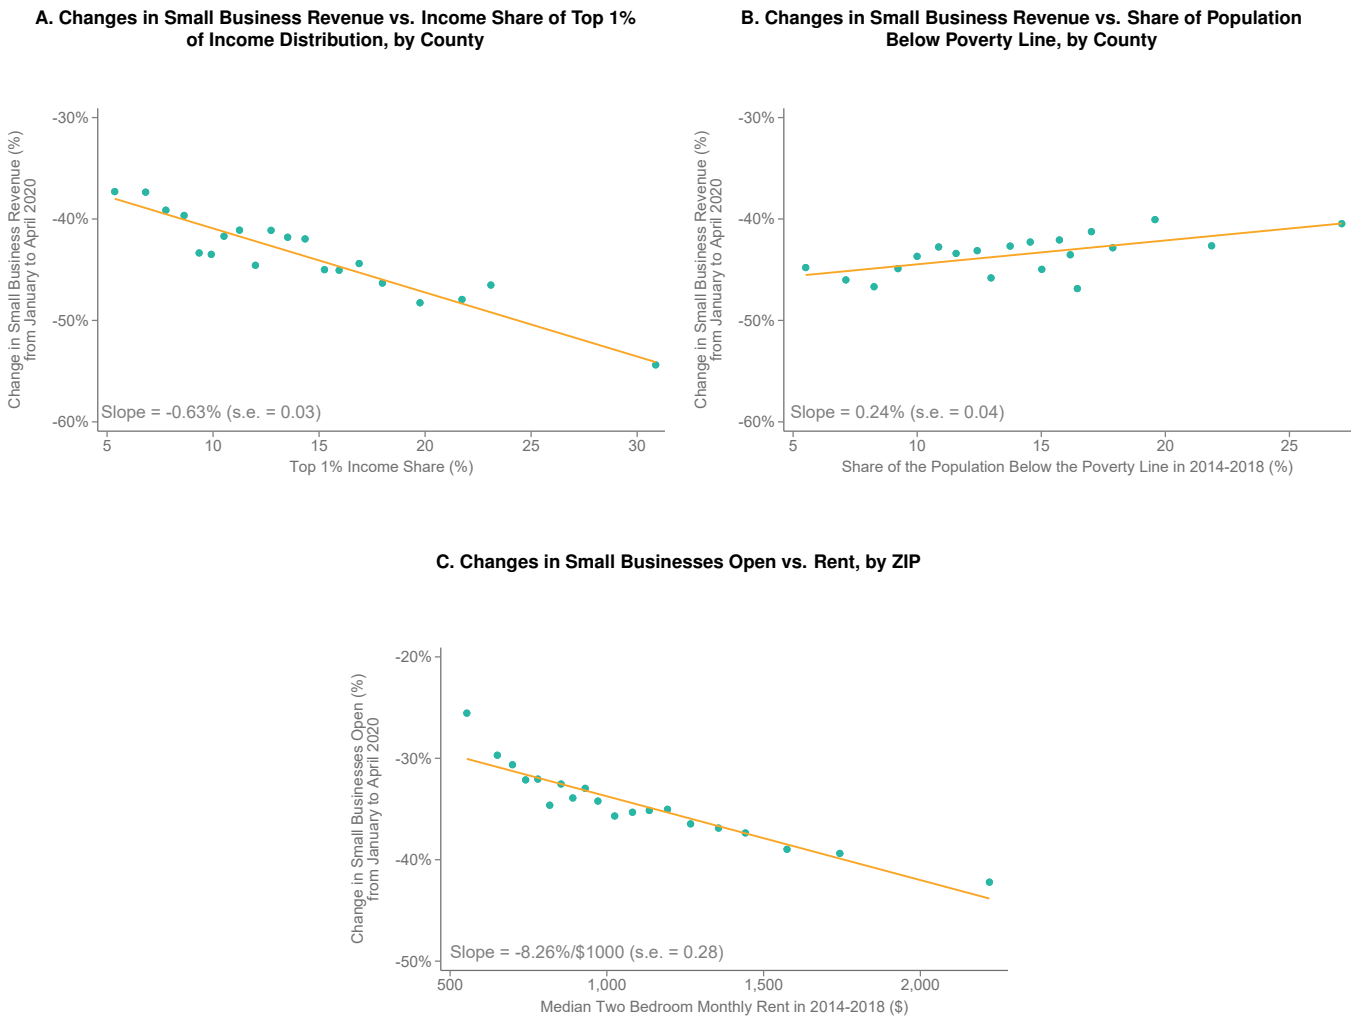


B. Changes in Low-Wage Employment, by CZ



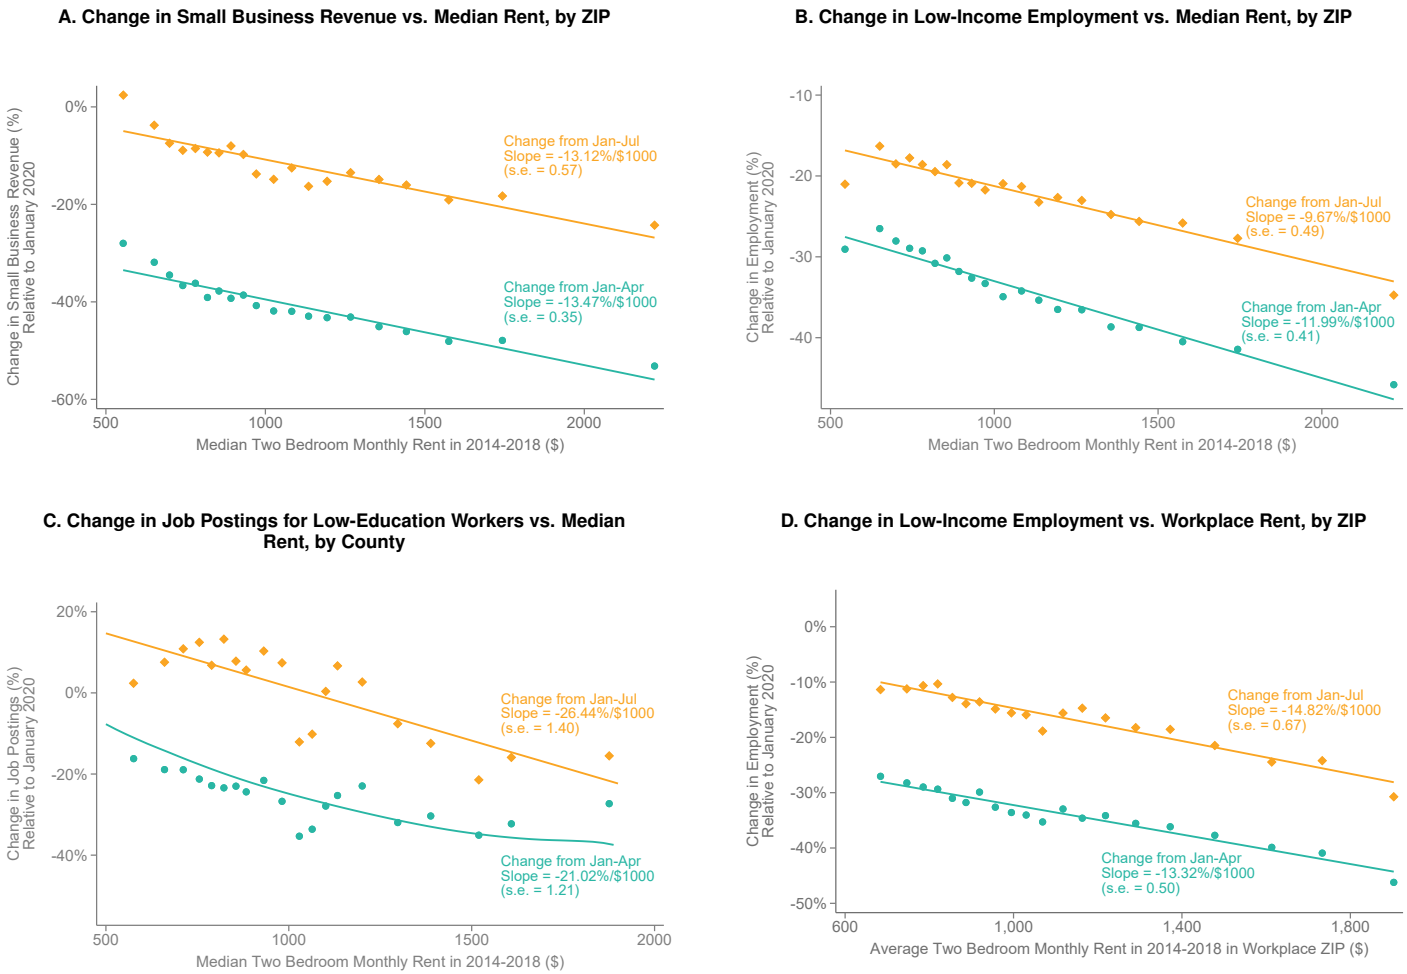
Notes: This figure presents national maps of changes in small business revenues (Panel A) and low-income employment (Panel B). Panel A replicates Figure 4 for the entire United States instead of a single MSA, showing the change in small business revenue from January 4-31 to March 25-April 14 in each county (rather than ZIP code, as in Figure 4). See the notes to Figure 4 for further details. Panel B replicates Figure 8 at the commuting zone (CZ) level for the entire United States instead of a single MSA, showing the change in employment from January 4-31 to April 8-28 in the Paychex-Intuit-Earnin combined data on employment in the bottom wage quartile (rather than Earnin data alone, as in Figure 8) in each CZ (rather than ZIP code, as in Figure 8). See the notes to Figure 8 for further details. Data sources: Panel A: Womply; Panel B: Paychex, Intuit, Earnin.

APPENDIX FIGURE 8: Changes in Small Business Outcomes vs. ZIP and County Characteristics



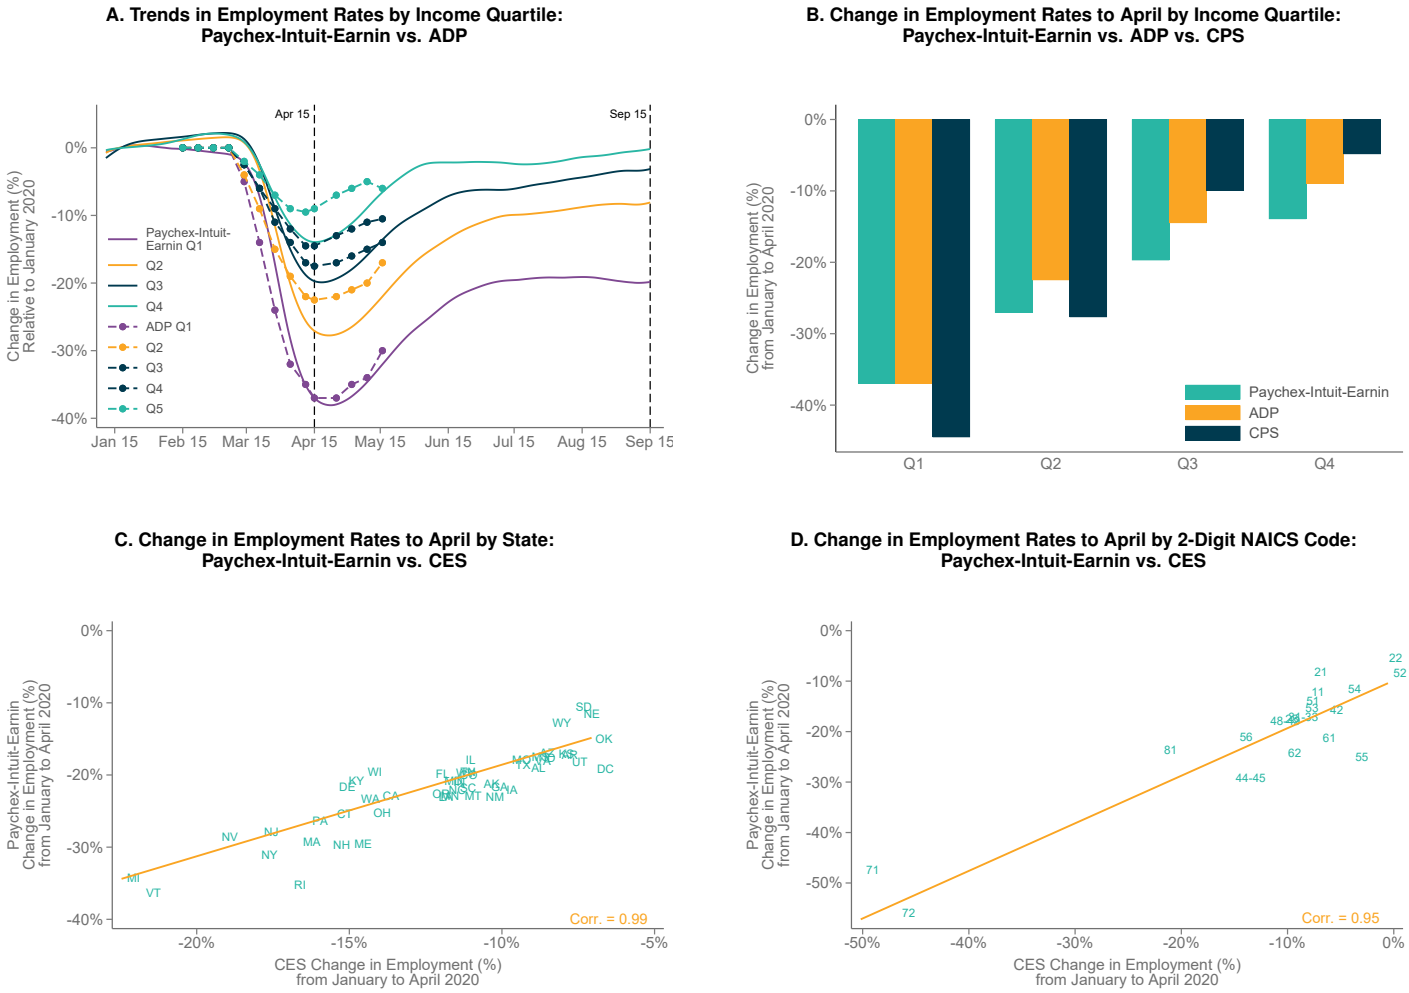
Notes: This figure shows the association between ZIP- or county-level characteristics and changes in small business outcomes between January 4-31 and March 25-April 14 2020, as measured in Womply data. The binned scatter plots are constructed as described in the notes to Figure 3. Panels A-B replicate Figure 5 but compare the declines in small business revenue with various measures of the distribution of income at the county level. Panel A 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, as 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 B 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. Panel C replicates Figure 5c using the change in the number of small businesses open, rather than the change in small business revenue, as the outcome variable. See notes to Figure 5 for details. Data source: Womply.

APPENDIX FIGURE 9: Changes in Small Business Revenue, Employment, and Job Postings From January to July vs. Rent



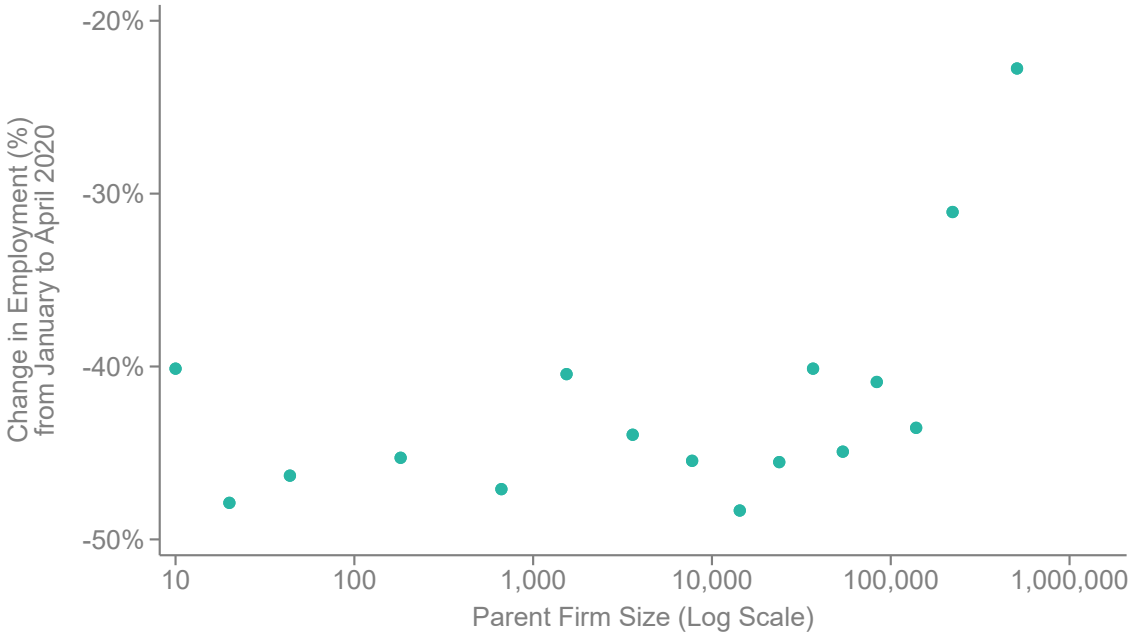
Notes: This figure presents binned scatter plots showing the association between changes in various economic measures vs. ZIP- and county-level median rent levels, contrasting the patterns in April vs. July 2020. See the notes to Figure 3 for more details on the construction of binned scatter plots. Panel A replicates Figure 5c, adding a second series showing ZIP-level changes in small business revenue from January to July 4-31 2020 vs. ZIP median rent. Panel B replicates Figure 9a, pooling all firm sizes, and then adding a second series showing ZIP-level changes in low-income employment (from Earnin) from January to July 4-31 2020 vs. ZIP median rent. Panel C replicates Figure 9b, adding a second series showing county-level changes in job postings from January to July 4-31 2020 vs. county median rent. Panel D replicates Figure 11a, adding a second series showing ZIP-level changes in low-income employment (from Earnin) from January to July 4-31 vs. workplace rent in ZIP. See the notes to Figure 11 for more details on the construction of workplace rent in ZIP. Data sources: Panel A: Womply; Panel B: Earnin; Panel C: Burning Glass; Panel D: Earnin.

APPENDIX FIGURE 10: Employment in Paychex-Intuit-Earnin Data vs. ADP, CPS, and CES



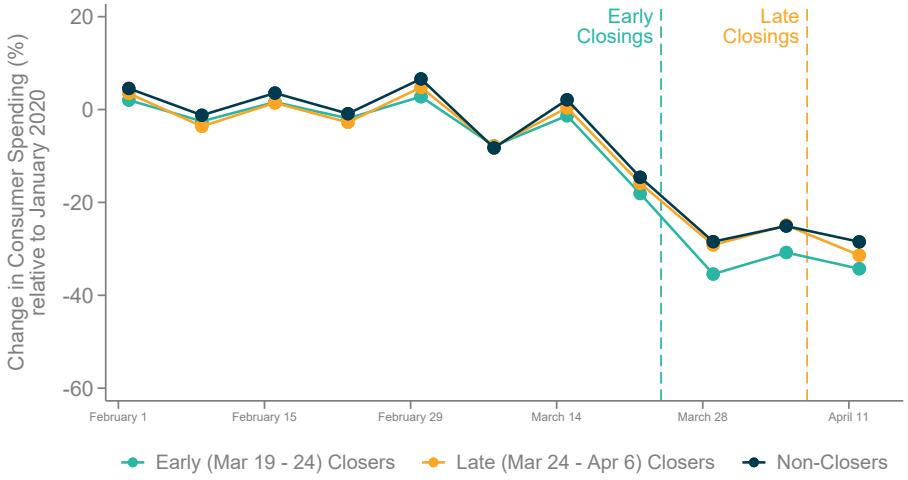
Notes: This figure benchmarks the Paychex-Intuit-Earnin combined employment series to the Current Population Survey (CPS), the Current Employment Statistics (CES), and estimates based on ADP data in Cajner et al. (2020). Panel A shows employment trends in the Paychex-Intuit-Earnin combined data (solid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

APPENDIX FIGURE 11: Changes in Low-Wage Employment by Firm Size



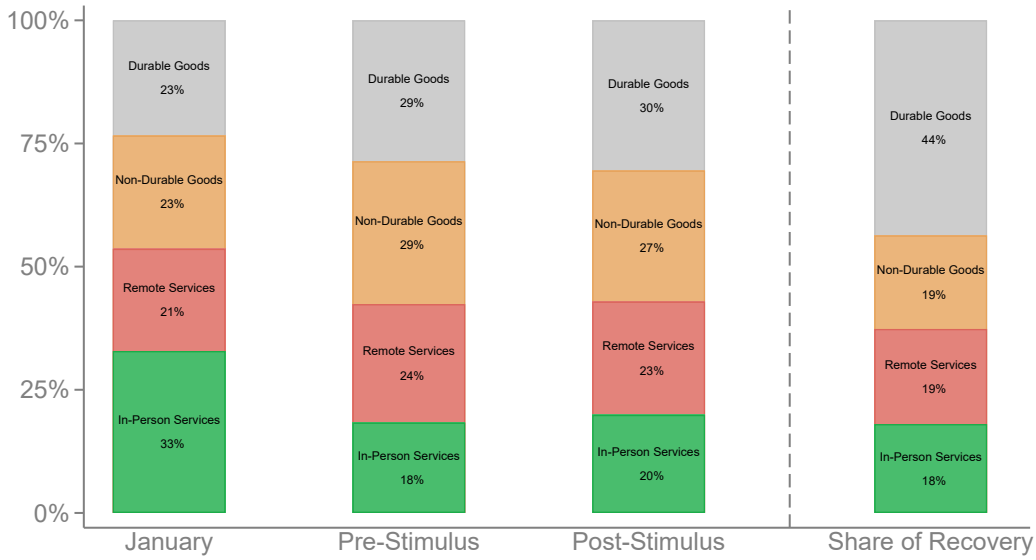
Notes: This figure displays a binned scatter plot of average percent declines in employment in the Earnin data at firms of different sizes. Binned scatter plots are constructed as described in notes to Figure 3. We calculate the change in employment from the period January 4-31 2020 to the period April 8-28 2020, and weight the binned scatter plot by employment in each firm over the period January 4-31 2020. We estimate the size of firms by matching Earnin employer names and locations to employer names and locations in ReferenceUSA data. Data source: Earnin.

APPENDIX FIGURE 12: Event Studies of Consumer Spending Around State-Ordered Business Closures



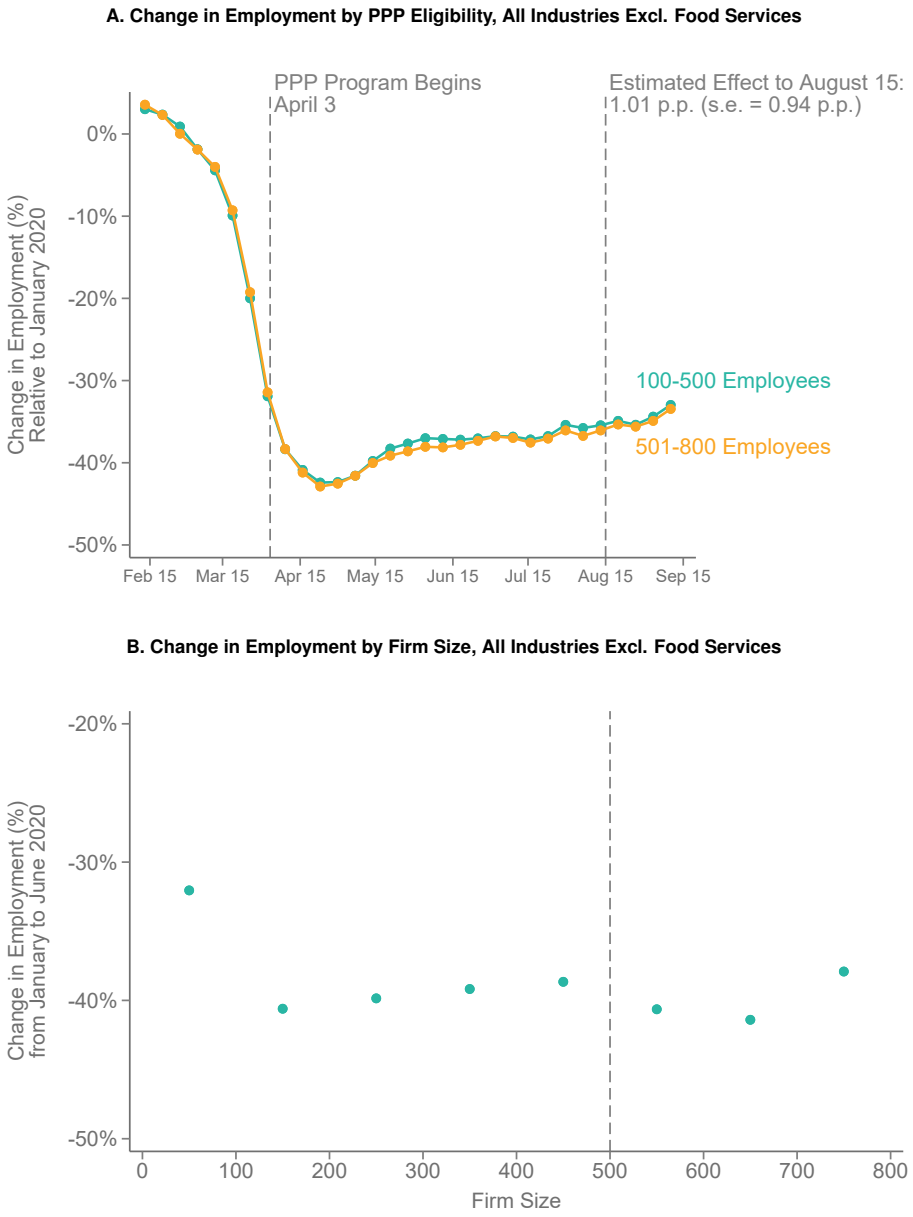
Notes: This figure displays trends in seasonally-adjusted consumer spending in the Affinity Solutions data, pooling states by the date on which a state-wide order closed non-essential businesses. States are aggregated into three groups: "Early" (state-wide closure order issued between March 19 and March 24), "Late" (state-wide closure order issued between March 30 and April 6), and "Non-Closers" (no state-wide closure order issued by April 6). Dashed lines denote the first date on which state-wide orders closing non-essential businesses were issued by "Early Closers" (March 19) and "Late Closers" (March 30). Data source: Affinity Solutions.

APPENDIX FIGURE 13: Effects of Stimulus Payments on Composition of Consumer Spending



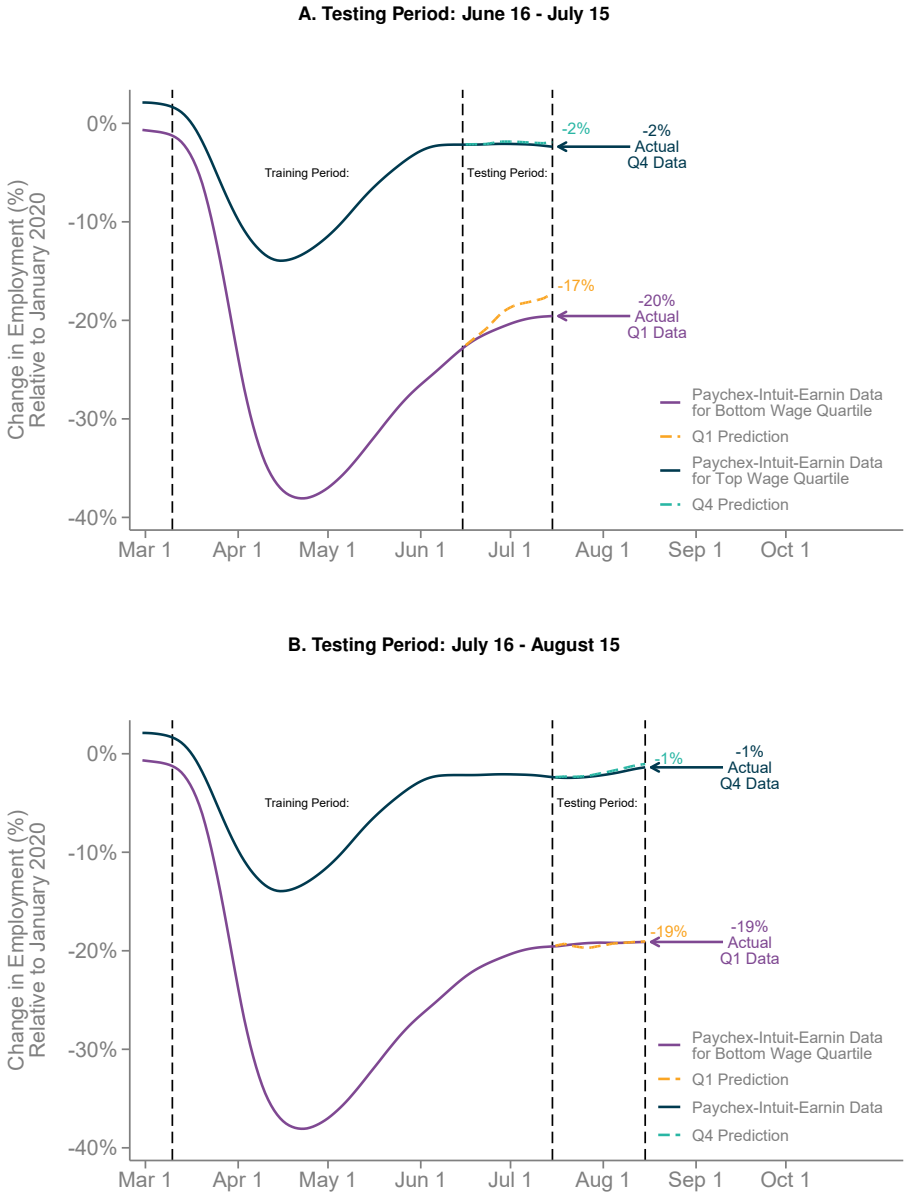
Notes: This figure presents statistics on the distribution of card spending across categories during various periods. The first bar replicates the right bar in Figure 2b, showing the composition of the level of spending for the period January 4-31 2020. The second and third bars replicate this distribution for the post-COVID, pre-stimulus period (March 25-April 14) and the post-COVID, post-stimulus period (April 29-May 5), respectively. The fourth bar replicates the left bar in Figure 2b, except decomposing the change during the recovery (the pre-stimulus to post-stimulus periods) rather than the decline. We define each spending category using Merchant Category Codes (MCCs), see Appendix B for details. Data source: Affinity Solutions.

APPENDIX FIGURE 14: Effect of Paycheck Protection Program on Employment in Earnin Data



Notes: This figure replicates Figure 15, using Earnin data rather than combined Paychex-Earnin data. For details, see notes to Figure 15. Data source: Earnin.

APPENDIX FIGURE 15: Out-Of-Sample Fit of Advance Employment Series



Notes: This figure compares out-of-sample predictions for employment to realized employment series. We construct predicted values for Paychex-Intuit-Earnin employment using Kronos data and Paychex data for firms with weekly paycycles; see notes to Figure 7 for details. Panel A compares out-of-sample predictions to realized values from June 16-July 15 2020. Panel B compares the out-of-sample prediction to realized values from July 16-August 15 2020. The root mean squared error (RMSE) for the difference between the prediction model and the true values across the top, middle, and bottom quartiles in the first testing period is 0.946 percentage points, while the RMSE across the top, middle, and bottom quartiles in the second testing period is 0.042 percentage points. Data sources: Paychex, Intuit, Earnin, Kronos.