Title

Fiscal multipliers in the COVID19 recession

Journal

Journal of International Money and Finance, 126(Quart. J. Econ. 132 3 2017)

ISSN

0261-5606

Authors

Auerbach, Alan
Gorodnichenko, Yuriy
McCrory, Peter B
et al.

Publication Date

2022-09-01

DOI

10.1016/j.jimonfin.2022.102669

Peer reviewed

Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Journal of International Money and Finance 126 (2022) 102669

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Review

Fiscal multipliers in the COVID19 recession

Alan Auerbach [a,b], Yuriy Gorodnichenko [a,b,*], Peter B. McCrory [d], Daniel Murphy [c]

[a] University of California-Berkeley, United States
[b] NBER, United States
[c] Darden School of Business, University of Virginia, United States
[d] J.P. Morgan Chase, United States

ARTICLE INFO

Article history:
Available online 18 May 2022

JEL:
E62
E32
H3

Keywords:
COVID19
Fiscal multiplier
Stimulus

ABSTRACT

In response to the record-breaking COVID19 recession, many governments have adopted unprecedented fiscal stimuli. While countercyclical fiscal policy is effective in fighting conventional recessions, little is known about the effectiveness of fiscal policy in the current environment with widespread shelter-in-place ("lockdown") policies and the associated considerable limits on economic activity. Using detailed regional variation in economic conditions, lockdown policies, and U.S. government spending, we document that the effects of government spending were stronger during the peak of the pandemic recession, but only in cities that were not subject to strong stay-at-home orders. We examine mechanisms that can account for our evidence and place our findings in the context of other recent evidence from microdata.

Contents

* Corresponding author.
E-mail address: ygorodni@econ.berkeley.edu (Y. Gorodnichenko).

https://doi.org/10.1016/j.jimonfin.2022.102669
0261-5606/

1. Introduction

In response to the record-breaking recession induced in the spring of 2020 by the COVID19 pandemic, many governments adopted unprecedented fiscal stimuli and shelter-in-place ("lockdown") policies. With the pandemic surging again in the winter months of 2020–21, further lockdowns have commenced. Little is known about the effectiveness of fiscal stimuli in such settings. Consider first the pandemic-induced massive spike in unemployment and excess capacity. Recent theoretical and empirical work suggests that fiscal stimulus is more effective in such an environment, either because supply is more responsive to demand-side stimulus in the presence of slack (the "slack" channel) or because the income decline pushes more households against their borrowing constraint, leading to high marginal propensities to consume out of additional income (the "high-MPC" channel).

Even if these channels are operative during typical periods of high unemployment and firm-level excess capacity, the pandemic episode is far from typical in that many locations implemented lockdown policies that effectively restricted subsets of spending and employment. According to recent theoretical work, such restrictions can dampen the effects of fiscal stimuli. In the case of direct government spending, contractors may be restricted from directly hiring employees, or – in the more general case of increases in households' net income – consumers may save new income for future spending, when a greater range of goods and services will be available (e.g., Auerbach et al., 2021; Guerrieri et al., 2020). Prior historical episodes lend some support to this dampening channel: the evidence in Brunet (2018) suggests that during World War II government restrictions on household spending caused lower fiscal multipliers. Similarly, the multipliers estimated by Ramey and Zubairy (2018) are lower when that war is included in the estimation sample. More generally, supply constraints, a prominent feature of the COVID19 crisis, can limit the power of fiscal policy to stimulate the depressed economy (e.g., Ghassibe and Zanetti, 2020, Jo and Zubairy, 2021).

To what extent are the slack and high-MPC channels operative during a pandemic recession, and how are they influenced by lockdown policies? We empirically address these questions by exploiting high-frequency data on government Department of Defense (DOD) spending, stay-at-home (SAH) orders, consumption, mobility, and employment during the onset of the COVID19 pandemic recession. In March and April of 2020 aggregate employment and household consumption plummeted, as did measures of worker and retail mobility. Aggregate DOD spending remained flat, potentially providing a stable source of demand during the pandemic.

Despite stable aggregate DOD spending, there was substantial variation in changes in spending across locations during the onset of SAH orders. There was also substantial independent variation in the extent of SAH orders across locations during the onset of the pandemic. As documented by (Baek et al., 2021), severe SAH orders (lockdowns) led to large declines in employment as well as worker and retail mobility, thus contributing to local economic slack but also imposing supply-side restrictions that might limit the effectiveness of demand-side stimulus. These two independent sources of cross-sectional variation permit us to compare changes in the effectiveness of DOD stimulus in unrestricted versus locked-down cities with the onset of lockdowns. We find a meaningful difference-in-differences: DOD employment multipliers are higher during the onset of the pandemic recession *only* for cities that were not subject to meaningful SAH orders. Therefore, while fiscal stimulus can remove slack, its effect on restricted economies is limited.[1]

The lack of an employment response in locked-down cities could reflect the possibility that lockdowns prevented matching between workers and employers. Alternatively, they may have prevented DOD spending from stimulating consumption, as consumers could not travel to retail or service establishments. To help disentangle these possibilities, we first examine whether DOD spending affected local consumption, as measured by Chetty et al. (2020). We find no evidence of a relationship between DOD spending and consumption, even in unrestricted locations. The fact that DOD spending increases employment but not consumption even in unrestricted locations is at first glance puzzling given that recessions are often associated with tight credit conditions, high average MPCs, and large consumption effects of fiscal stimulus (e.g., Eggertsson and Krugman, 2012). The lack of a consumption response implies that the larger (across-location average) fiscal multipliers on employment during the pandemic recession are driven by the slack channel rather than the high-MPC channel. While potentially surprising, the lack of a consumption response is consistent with recent evidence that households predominantly saved their 2020 stimulus checks (Coibion et al., 2020a). Overall, our evidence implies that fiscal stimulus is indeed more effective in recessions, but not if there are restrictions on spending or other forms of economic activity. Lockdowns effec-

[1] Our analysis uses stay-at-home orders as an indicator of the severity of lockdowns and does not attempt to disentangle the effects of stay-at-home orders *per se* from the effects of the pandemic conditions that led to the stay-at-home orders.

tively restrict the ability of the economy to absorb slack, and there are no detectable consumption responses that would contribute to stronger general equilibrium effects of government spending.

Our evidence adds to a growing literature on state-dependent fiscal effects. Much of this literature has examined whether fiscal multipliers vary across recessions and expansions (e.g., Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018), with the level of consumer debt (e.g., Demyanyk et al., 2019; Klein, 2017; Miranda-Pinto et al., 2020a), or with the sign of the change in government spending (Barnichon et al., 2020). We expand on this literature by examining how fiscal effects vary across lockdown status. In particular, we examine the difference (lockdown versus unrestricted) in differences (expansion versus recession). We add to the body of evidence in support of the slack channel in generating higher fiscal multipliers in recessions, and present new evidence that fiscal multipliers are lower in the presence of SAH orders.

2. Data

Our analysis relies on several sources of data, which we describe below. The unit of analysis is the core-based statistical area (CBSA).[2] Most of the analysis uses data at the monthly frequency but some exercises are done at other frequencies.

2.1. Government spending

Following Auerbach et al., 2020b; hereafter AGM), we use Department of Defense (DOD) spending as a measure of government spending. DOD spending accounts for approximately half of discretionary spending of the federal government. Unlike other types of spending, DOD spending is largely insensitive to business conditions and can be modelled as a demand shock from a CBSA's perspective (e.g., DOD spending does not directly enter households' utility or affect productivity/infrastructure of the local economy).

To construct a measure of spending from DOD contracts (available from the Federal Procurement Data System, or FPDS at https://www.fpds.gov), we follow AGM: for each contract, we compute average monthly spending (a contract's obligation/-contract value divided by the duration of the contract) and then sum the derived spending across contracts that are active for a given month.[3] As discussed in AGM, this measure of spending closely tracks DOD spending in the National Income and Product Accounts and U.S. Treasury payments to defense contractors reported in the Daily Treasury Statements.

Table 1 documents dramatic variation and concentration in DOD spending across CBSAs in absolute and per capita terms. For example, Dallas-Fort Worth-Arlington is a major recipient of DOD contracts ($27.1 billion in 2019), much of which is allocated to the construction of a new generation of fighter jets (see AGM for more details). Economies for smaller CBSAs can receive a very large share of their business activity from the DOD. For example, Norwich-New London, CT—the hometown of Electric Boat's nuclear submarine production—received almost $40,000 per worker from the DOD in 2019. We note that DOD spending is dispersed across many U.S. states, a feature important for studying regional variation in lockdowns during the COVID19 crisis.

Although DOD spending has a number of desirable features (e.g., it accounts to a large share of discretionary government spending -- 2 % of GDP), it is less likely to be sensitive to cyclical factors. Other types of spending—e.g., transfers to households (Economic Impact Payments, EIP) and firms (Payroll Protection Program, PPP)—accounted for the bulk for fiscal support provided during the crisis. As a result, our focus on DOD spending should be understood as a way to shed more light on the mechanisms behind government spending multipliers rather than as an evaluation of the fiscal programs implemented in response to COVID19. Furthermore, we do not engage in welfare calculations because the fiscal policy was not only concerned with supporting the aggregate demand but also—or maybe mostly—concerned with providing disaster relief to affected households and firms. In other words, the size of fiscal multipliers is only a partial measure of the government's success in achieving its social objectives.

2.2. Employment

Our main source of employment data is Local Area Unemployment Statistics (LAUS). These data are produced by the Bureau of Labor Statistics (BLS) and are based on the Current Population Survey (CPS). While BLS provides the figures for the employed and unemployed, we focus on employment data given issues with defining unemployment during the pandemic.[4] Because the CPS is a household survey, LAUS avoids double-counting workers with multiple jobs, which can happen with employer-based data, e.g., the Quarterly Census of Employment and Wages (QCEW), an administrative source for employment and payroll data utilizing state-run unemployment insurance systems. Normally, employment growth series at the CBSA-

[2] CBSA is a geographic area defined by the Office of Management and Budget that consists of one or more counties (or equivalents) anchored by an urban center of at least 10,000 people plus adjacent counties that are socioeconomically tied to the urban center by commuting.

[3] In addition to new contract obligations, the dataset also contains modifications to existing contracts, including downward revisions to contract amounts (de-obligations) that appear as negative entries. Many of these de-obligations are very large and occur subsequent to large obligations of similar magnitude. Furthermore, in many cases, de-obligations happen within days after obligations appear in the reporting system. When we observe obligations and de-obligations with magnitudes within 0.5 percent of each other, we consider both elements of the pair to be null and void as it is unlikely that any outlays were associated with these temporary obligations. This restriction removes 4.7 percent of contracts from the sample.

[4] See https://www.bls.gov/cps/effects-of-the-coronavirus-covid-19-pandemic.htm for a discussion of measurement issues.

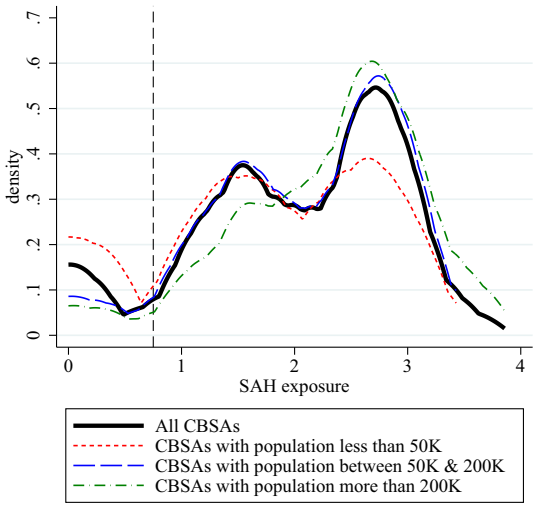

Fig. 1. Distribution of intensity for stay-at-home (SAH) orders (by number of weeks). Notes: The figure shows kernel density for the duration (in weeks) distribution of stay-at-home (SAH) as of April 11, 2020. The data are from (Baek et al., 2021). The vertical dashed line shows the cutoff used in the main analysis to classify core-based statistical areas (CBSAs) into lockdown (restricted) and no-lockdown (unrestricted) cities.

level are highly correlated between QCEW and LAUS (0.80 in April 2019), but this correlation declined sharply to 0.43 in April 2020. We conjecture that this correlation falls because QCEW and LAUS differ in the coverage of workforce; e.g., LAUS covers self-employed and gig-economy (e.g., Uber or DoorDash) workers while QCEW does not. Given that LAUS has a more comprehensive definition of employment, we use LAUS as the main source of employment data to study the COVID19 crisis. Unlike QCEW, LAUS does not have information by sector and thus we will not able to explore sectoral variation (e.g., tradables vs.

Table 1
Distribution of Department of Defense (DOD) spending and Stay-At-Home (SAH) orders.

Top 20 CBSA by absolute DOD spending, billions			Top 20 CBSA by DOD spending per worker, thousands		
CBSA	Spending	SAH	CBSA	Spending	SAH
Washington-Arlington-Alexandria, DC-VA-MD-WV	35.4	1.79	Norwich-New London, CT	39.9	2.86
Dallas-Fort Worth-Arlington, TX	27.1	2.73	Lexington Park, MD	39.2	1.86
St. Louis, MO-IL	12.3	2.58	Pascagoula, MS	39.2	1.29
Boston-Cambridge-Quincy, MA-NH	10.8	2.68	Oshkosh-Neenah, WI	28.9	2.57
Virginia Beach-Norfolk-Newport News, VA-NC	10.7	1.86	Huntsville, AL	26.6	1.14
Los Angeles-Long Beach-Santa Ana, CA	10.5	3.43	Enterprise-Ozark, AL	18.9	1.14
San Diego-Carlsbad-San Marcos, CA	9.1	3.43	Fort Polk South, LA	18.4	2.86
Baltimore-Towson, MD	9.0	1.86	Fort Leonard Wood, MO	17.6	0.86
Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	8.8	2.89	Amarillo, TX	17.1	1.84
Seattle-Tacoma-Bellevue, WA	6.8	2.86	Vicksburg, MS	16.4	1.29
Sacramento–Arden-Arcade–Roseville, CA	6.0	3.43	Sierra Vista-Douglas, AZ	15.1	1.71
Huntsville, AL	5.9	1.14	Kingsville, TX	14.3	1.43
New York-Northern New Jersey-Long Island, NY-NJ-PA	5.5	3.04	Warner Robins, GA	14.1	1.29
Louisville/Jefferson County, KY-IN	5.4	2.47	Fort Walton Beach-Crestview-Destin, FL	14.1	1.29
Norwich-New London, CT	5.3	2.86	Virginia Beach-Norfolk-Newport News, VA-NC	12.8	1.86
Orlando-Kissimmee-Sanford, FL	4.9	2.16	Fairbanks, AK	12.1	2.14
Hartford-West Hartford-East Hartford, CT	4.7	2.86	Palm Bay-Melbourne-Titusville, FL	12.0	1.29
Chicago-Joliet-Naperville, IL-IN-WI	4.6	3.11	Burlington, IA-IL	11.4	0.15
San Jose-Sunnyvale-Santa Clara, CA	4.5	3.85	Mobile, AL	10.6	1.14
Phoenix-Mesa-Glendale, AZ	4.1	1.71	Washington-Arlington-Alexandria, DC-VA-MD-WV	10.6	1.79
Average	0.3		Average	1.1	
P99	6.8		P99	16.4	
P95	1.3		P95	5.6	
P90	0.4		P90	2.8	
P50	0.0		P50	0.1	
St.Dev.	1.8		St.Dev.	3.4	

Notes: The table shows the top 20 cities (core-based statistical areas; CBSAs) by absolute or per-worker spending by the Department of Defense (DOD) in 2019. The bottom rows report moments (averages and percentiles; e.g., P99 is the 99th percentile) for all CBSAs.

Table 2
Comparison of lockdown vs. no-lockdown cities.

	Control Variables by Lockdown Status			
	Not locked down (N = 116)		Locked down (N = 824)	
	mean	median	mean	median
Population	97,500	45,141	337,234.4	81,280.5
Urban	0.60	0.64	0.62	0.62
No HS diploma	0.15	0.14	0.16	0.15
HS diploma	0.34	0.34	0.34	0.33
College Degree	0.21	0.20	0.21	0.19
Unemployment Rate	0.04	0.04	0.05	0.05
Vacancy Rate	0.12	0.11	0.13	0.11
Owner Occupancy Rate	0.71	0.71	0.70	0.71
Median Home Value	113,167	100,222	152,806	124,079
Median Household Income	43,088	43,162	45,117	43,839
DOD spending Share	0.013	0.002	0.029	0.004

Note: Summary Statistics are based on data from the 2010 Census and American Community Survey (as reported by NHGIS).

Table 3
Case study of cities with variation in Stay-At-Home (SAH) orders and Department of Defense (DOD) *spending.*

	High SAH exposure		Low SAH exposure	
	High DOD spending Dallas-Ft. Worth, TX	Low DOD spending Houston, TX	High DOD spending Omaha, NE	Low DOD spending Des Moines, IA
	(1)	(2)	(3)	(4)
Department of Defense (DOD) spending				
DOD spending, 2019, $ billion	27.09	2.58	0.67	0.06
DOD spending per worker, 2019, $ thousand	7.10	0.78	1.38	0.17
DOD spending per payroll, 2019, %	11.70	1.24	2.66	0.29
Demographics				
Employment, 2019, thousands	6371.77	5946.80	865.35	569.63
Unemployment rate, 2019, %	3.26	3.79	3.06	2.69
Share with college degree, 2010, %	30.98	28.46	32.48	32.49
Share of white, 2010, %	50.25	39.69	78.72	83.64
Dynamics during the COVID19 crisis				
Employment change, April 2020 to April 2019, %	−16.72	−18.59	−5.73	−8.93
Consumer spending, April 2020, % relative to pre-COVID19	−38.27	−36.67	−36.11	−41.37
Retail mobility, April 2020, % relative to pre-COVID19	−47.29	−45.24	−41.64	−45.27
Work mobility, April 2020, % relative to pre-COVID19	−30.46	−31.80	−25.57	−33.59

non-tradables) in multipliers during the COVID19 crisis. Following prior CBSA-level studies of fiscal stimulus, we normalize our DOD spending by the size of the local economy, as measured by 2019 payroll data from the QCEW.

2.3. Stay-at-home orders and mobility

A central element of our analysis is geographical variation in the intensity of lockdowns induced by COVID19. (Baek et al., 2021) construct a database of county-level stay-at-home (SAH) orders as of April 11, 2020, when approximately 90 percent of the U.S. population was under a lockdown order. The Baek et al. data report the number of weeks that a given CBSA has spent in a government-mandated lockdown by April 11. Fig. 1 documents considerable variation in SAH intensity, which reflects the decentralized implementation of public health measures aimed to contain the virus and to prevent local hospitals from being overwhelmed. While some CBSAs were locked down for nearly a month (e.g., the San Francisco Bay Area), some CBSAs did not have any lockdowns at all (e.g., Omaha, NE). Indeed, CBSAs with no lockdowns in April 2020 did not impose lockdowns in subsequent months and CBSAs with lockdowns introduced close to April 11, 2020 had short-lived lockdowns. To capture this group of CBSAs, we create a no-lockdown group that covers cities with SAH lasting less than 3/4 of a week as of April 11, 2020, which corresponds to a "hump" in the left-tail of the distribution.

SAH orders were not issued randomly across cities. For example, the no-lockdown group consists of relatively small CBSAs (less than 200,000 in population), see Fig. 1 and Table 2. This pattern is consistent with the notion that more densely populated areas were more vulnerable to the virus and thus were more likely to lock down. Hence, we will need to use con-

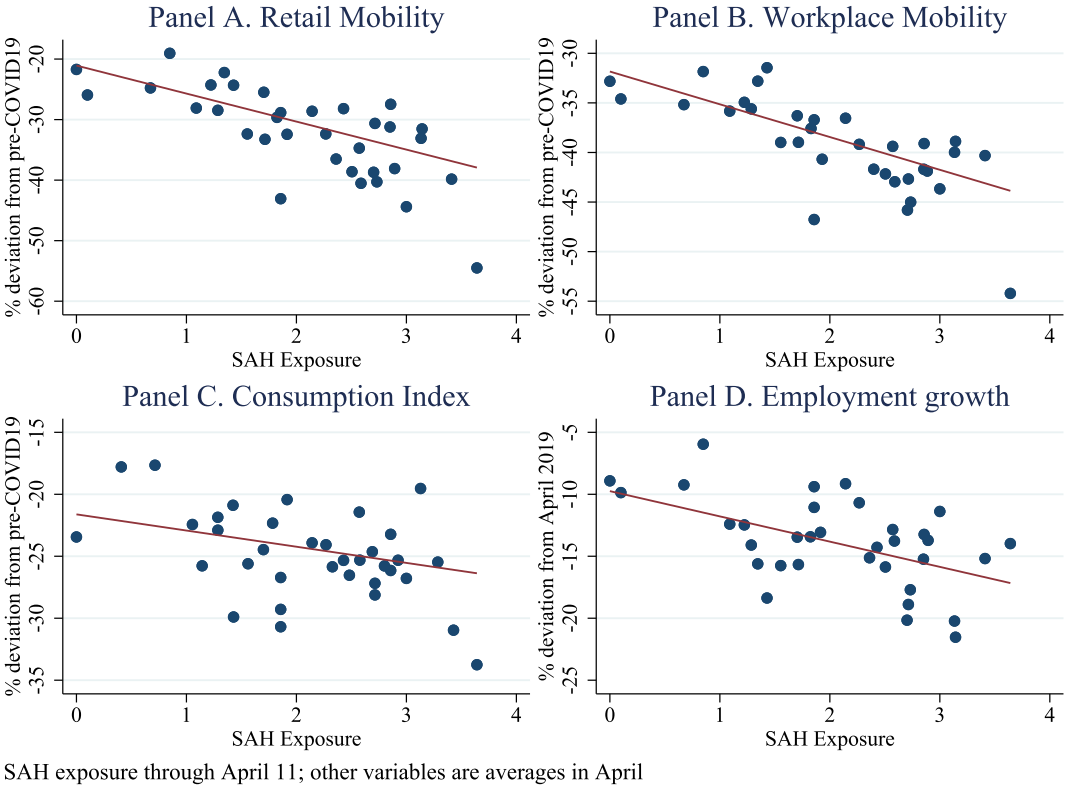

Fig. 2. Stay-At-Home (SAH) orders and economic outcomes. Notes: Each panel presents a binscatter for the exposure to stay-at-home (SAH) orders (measured in weeks) vs. an economic outcome across CBSAs.

trols to correct for possible imbalances in the lockdown vs. no-lockdown groups or match no-lockdown (unrestricted) cities to comparable lockdown (restricted) cities.

While SAH orders provide a direct measure of limits on economic activity, self-imposed restrictions (e.g., people may voluntarily limit their work/consumer/travel activities to minimize exposure to the virus) likely contributed to the contraction of economic activity. Using cellphone mobility data, Google constructs daily measures of "retail" and "workplace" mobility at the county level during the crisis, where mobility is measured as a percent of pre-COVID19 mobility.[5] While retail mobility is not a comprehensive measure of consumer spending, Baker et al. (2020) document that retail mobility is highly correlated with consumer spending. Furthermore, "retail" mobility can provide a measure of local spending in the sense that it captures the intensity with which consumers travel to local retail outlets rather than making purchases online. Weighting by employment, we aggregate these indices to the CBSA level and use them as a high-frequency indicator of mandatory and voluntary restrictions as well as economic activity.

Fig. 2 shows that SAH orders have large predictive power for the cross-sectional variation in retail/workplace mobility, consumer spending, and employment. For example, a one-week increase in the duration of SAH orders is associated with 2 percent decline in employment (April 2020 relative to April 2019) and 1.3 percent decline in consumer spending (April 2020 relative to the pre-COVID19 level[6]). Note that even CBSAs with no SAH orders demonstrate large declines for these indicators. This pattern is consistent with at least two hypotheses. First, consumers/workers can voluntarily curtail their mobility and spending in attempts to limit infections. Second, SAH orders in other locations trigger local economic contractions which via e.g. general equilibrium effects, input-output linkages or other channels can induce economic contractions in areas not directly affected by SAH orders. See (Baek et al., 2021) for further discussion for the relative merit of these two hypotheses.

2.4. Consumer spending

Consumer spending at the CBSA level is generally lacking, although some measures (e.g., new car registrations) are available. Recent work (e.g., Chetty et al., 2020) developed a measure of consumer spending (including e-commerce) based on transaction-level data provided by financial service firms. Although these data are likely incomplete (e.g., they miss cash-

[5] These data are available at https://www.google.com/covid19/mobility/.

[6] Chetty et al. (2020) use the first four complete weeks of 2020 as a measure of pre-COVID spending.

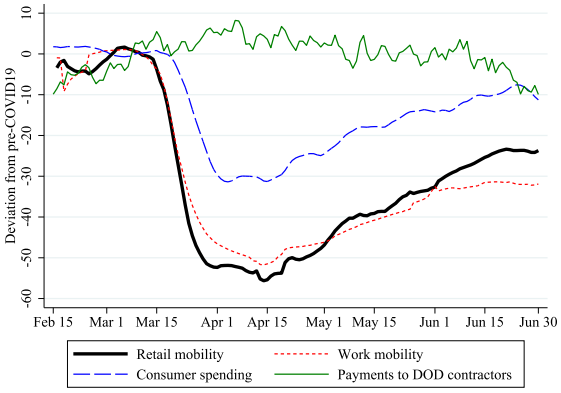

Fig. 3. Aggregate time series for Department of Defense (DOD) spending, consumer spending, and retail/work mobility. Notes: Payments to DOD contractors (daily U.S. Treasury statement) is the 30-day moving average. All other variables are 7-day moving averages. All variables are normalized to be equal to zero on average during the March 1, 2020 – March 15, 2020 period. Department of Defense (DOD) daily spending is from the U.S. Treasury. Retail and work mobility are from Google trends. Consumer spending is from Opportunity Insight. Mobility and consumer spending are computed as averages across CBSAs weighted by population.

based transactions and people with limited banking), these data give us daily series, thus allowing us to establish precise timing of events. The Chetty et al. (2020) data report consumer spending as the percent deviation from the pre-COVID19 level.

2.5. Basic trends

Fig. 3 plots the dynamics of key national-level series at the daily frequency during the COVID19 crisis. Google retail and workplace mobility indices declined by 50 percent between March 15, 2020 and April 1, 2020. Consumer spending fell by 30 percent relative to the pre-COVID19 levels. Note that this decline is somewhat delayed relative to the Google mobility indices, which likely reflects consumers' attempts to stock up food and other basic goods. After April 15, 2020, mobility and consumer spending gradually revert to pre-COVID19 levels and, by the end of June 2020, half of the decline was recovered. Workplace mobility recovered less than other indicators by this point.

These dynamics demonstrate the ferocity and depth of the economic contraction. Early estimates (Coibion et al., 2020b) and subsequent government statistics documented that approximately 20 million jobs were destroyed in April 2020. Unusually, few workers searched for new jobs—a sign of massive layoffs (perceived to be temporary at the time) and the realization that few businesses were hiring—and many workers left the labor force.

Against this background of devastation, DOD spending maintained its pace. Fig. 3 shows that there was little variation in U.S. Treasury payments to defense contractors, which is consistent with a stable flow of funding for existing DOD contracts as well as new obligations issued by the DOD. To avoid disruptions in defense contractors' operations, the DOD categorized defense contractors as essential critical infrastructure workforce on March 20, 2020 and thus exempted contractors from SAH restrictions. Furthermore, the DOD provided necessary logistical support to ensure that contractors had access to parts and materials. Thus, in contrast to private demand exhibiting a dramatic and rapid decline, DOD spending proved to be a stable, reliable source of demand in the economy. We will use this property of DOD spending to explore how government spending can affect the local economy depending on its conditions and SAH restrictions.

2.6. A case study

As we discuss above, the wide distribution of SAH orders and DOD spending provides a unique laboratory for assessing how a demand shock can stimulate a local economy when the economy exhibits slack and faces constraints for utilizing idle resources. To gather tentative insights from the data, we consider two pairs of CBSAs with similar exposure to SAH orders but different reliance on DOD contracts. Relevant statistics are in Table 3.

The first pair of cities is in Texas: Dallas-Ft. Worth and Houston. Both cities had relatively high exposure to SAH orders (almost three weeks as of April 11, 2020). The cities are roughly similar in size and other characteristics but, while Houston has a relatively small share of the economy servicing the DOD (the value of DOD spending is approximately 1 percent of Houston's payroll in 2019), Dallas-Ft. Worth has a large DOD sector (approximately 12 percent of payroll). Despite this order-of-magnitude difference in DOD spending across cities, Dallas-Ft. Worth had only slightly better outcomes for workplace mobility and employment growth. Strikingly, consumer spending and retail mobility declined more in Dallas-Ft. Worth than in Houston.

The second city pair is Omaha, NE and Des Moines, IA. These two cities had no SAH orders as of April 2011 (zero exposure). The cities are broadly similar but Omaha (2.7 percent of payroll) has roughly an order-of-magnitude more DOD spend-

ing than Des Moines (0.3 percent of payroll). In April 2020, Omaha had smaller declines in employment, consumer spending and mobility than Des Moines. This difference is consistent with the notion that DOD spending provided a reliable source of demand in Omaha which helped this city to withstand a fall in private demand.

The experience of these two pairs of cities suggests that SAH orders can contribute to the differential response of local economies to DOD demand in the dire conditions of April 2020. Specifically, high exposure to SAH orders can blur or even eliminate the difference between high-DOD-spending cities (Dallas-Ft. Worth) and low-DOD-spending cities (Houston). In contrast, low exposure to SAH orders appears to allow for a differential response to DOD spending. Obviously, this case study is only suggestive. We do a more systematic analysis of this pattern in the next section.

3. Results

Our primary objective is to examine how DOD employment multipliers differed across locations with different lockdown status during the deep recessionary month of April 2020 (relative to expansionary periods prior to April). After presenting our econometric specification, we discuss threats to identification and ways to address potential endogeneity. Then we present our baseline results for employment (the local economic outcome indicator with the highest quality) and a series of robustness checks. Finally, we examine other outcomes to better understand the mechanisms.

3.1. Empirical specification

Our empirical approach is an adaptation of the cross-sectional specification in Demyanyk et al. (2019; henceforth DLM):

$$\frac{\Delta Y_i}{Payroll_i^{Pre}} = \alpha + \beta_1 \frac{\Delta G_i}{Payroll_i^{Pre}} + \gamma State_i + \beta_2 \frac{\Delta G_i}{Payroll_i^{Pre}} \times State_i + Controls_i + \epsilon_i \quad (1)$$

where ΔY_i is the change in employment for CBSA i over a one-year period, ΔG_i is the change in DOD spending in CBSA i, $Payroll_i^{Pre}$ is average payroll between May 2018 to April 2019 (from the QCEW) measuring the size of the local economy, and $State_i$ measures the state of the local economy, which in our main specification corresponds to lockdown status in April 2020. Changes in the outcome variable and DOD spending are normalized by initial-period CBSA-level earnings such that $\beta_1 + \beta_2 State_i$ is the effect of a dollar change in DOD spending on the change in the outcome variable in CBSA i. The specification permits the effect of DOD spending to vary across CBSAs with different levels of $State_i$.

Our measure of the relevant state of the economy is an indicator for whether a CBSA was locked down in April of 2020: $State_i = 1(\text{lockdown}_i)$. Since lockdowns were temporary, we expect them to affect the economy at a high frequency and therefore we examine higher-frequency data on outcome variables. In particular, ΔY_i in our specification is the change in monthly employment between April 2020 and April 2019. For ΔG_i, we compare spending changes over the twelve months prior to and including April 2020 to changes over the year prior to and including April 2019. We examine changes in year-on-year DOD averages rather than changes over a twelve-month period due to the lack of information regarding when our measure of contract spending is associated with new production. As emphasized by AGM, the precise timing of disbursements and new production by the contractor are unobserved. High-frequency changes in DOD spending are likely to be dominated by transitory ("wealth transfer") changes rather than changes in new production.[7] The large wealth transfer component biases downward estimates of DOD multipliers relative to the true effect of the "new production component" of DOD spending.

AGM propose to further address the downward bias caused by the wealth transfer component of DOD production by using a Bartik instrument variable (IV),[8] which not only filters out the transitory ("wealth transfer") component of DOD contracts to which one might expect weaker economic responses (see AGM for further discussion), but also addresses the endogeneity concerns raised in Nakamura and Steinsson (2014) and DLM. In DLM, national spending was increasing drastically into the Great Recession on account of the Iraq War, thus contributing to the strong relevance of the IV in their study. Our empirical setting does not have the benefit of such drastic changes in national spending, which exhibited essentially no growth between 2019Q2 and 2020Q2; therefore the Bartik instrument is weak in our setting.

If a good instrument is not available, can OLS estimates uncover the existence of state dependent effects of DOD spending associated with new production, and if so how close are the estimates to the actual state-dependence? To answer this ques-

[7] To appreciate the difference between "wealth transfers" and "new production," consider the following two examples. Electric Boat Corporation builds nuclear submarines for the Department of Defense. There is usually little uncertainty about defense funding for Electric Boat, but there could be some uncertainty about when a contract is going to be signed and money transferred to Electric Boat. In this case, the econometrician can observe a lot of high-frequency variation in funds paid to Electric Boat but this variation does not materially affect employment or investment decisions of the contractor because it is assured of funding and most of the payments to the contractor are anticipated. As a result, payments to Electric Boat are "wealth transfers" that do not affect the allocation of resources at the time these payments are made. In contrast, "new production" corresponds to an unanticipated, persistent shock that affects the allocation of resources when a contract is awarded. For example, in the late 1990s, Boeing and Lockheed competed for a large contract to build a new fighter jet and neither of the firms was certain about winning the contract. When Lockheed won the contract in 2001, the uncertainty was resolved and Lockheed greatly expanded its economic activity in Dallas-Ft.Worth. Over the course of 20 years, the Department of Defense spent more than $400 billion on the program to build this jet.

[8] A Bartik instrumental variable (IV) in our context is $\frac{share_i \times \Delta G_t}{Payroll_i^{Pre}}$ where $share_i$ is the CBSA's average share of DOD contract spending over a long period and ΔG is the contemporaneous change in aggregate DOD spending.

tion, consider the following stylized setting. When the economy is functioning normally, output is determined by the following relationship.

$$Y_{it} = \alpha_0 G_{it}^{Transfer} + \beta_0 G_{it}^{NewProd} + \epsilon_{it}, \tag{2}$$

where i and t index locations and time, Y is a measure of output (employments, etc.), $G^{Transfer}$ is the anticipated component ("wealth transfer") of government spending, $G^{NewProd}$ is the unanticipated component ("new production) of government spending, and ϵ is a collection of other factors. To make algebra tractable, we assume without loss of generality that $E\left(\epsilon_{it} G_{it}^{Transfer}\right) = E\left(\epsilon_{it} G_{it}^{NewProd}\right) = E\left(G_{it}^{Transfer} G_{it}^{NewProd}\right) = 0$. Coefficients α an β are multipliers for the two types of government spending. Macroeconomic models typically predict that $\alpha_0 \approx 0$ and $\beta_0 > \alpha_0$. Usually, researchers aim to estimate β_0 but α_0 may be of interest too.

Suppose that we do not observe $G^{Transfer}$ and $G^{NewProd}$ separately. Instead, we observe the sum of these two components: $G_{it} = G_{it}^{NewProd} + G_{it}^{Transfer}$. In this case, when we estimate the following regression with OLS.

$$Y_{it} = c_0 G_{it} + u_{it}, \tag{3}$$

we obtain

$$c_0^{OLS} = \alpha_0 \frac{var\left(G_{it}^{Transfer}\right)}{var\left(G_{it}^{Transfer}\right) + var\left(G_{it}^{NewProd}\right)} + \beta_0 \frac{var\left(G_{it}^{NewProd}\right)}{var\left(G_{it}^{Transfer}\right) + var\left(G_{it}^{NewProd}\right)}.$$

Hence, the OLS estimate is a weighted average of two multipliers. If variation in the transfer component $G_{it}^{Transfer}$ is much greater than variation in the new-production component $G_{it}^{NewProd}$, the OLS estimate may be close to zero.

If one has an instrument V_{it} for $G_{it}^{NewProd}$, then one can recover β_0 and estimate a multiplier (potentially much) greater than c_0^{OLS}:

$$c_0^{IV} = \frac{cov(V_{it}, Y_{it})}{cov(V_{it}, G_{it})} = \frac{cov\left(V_{it}, \alpha_0 G_{it}^{Transfer} + \beta_0 G_{it}^{NewProd}\right)}{cov\left(V_{it}, G_{it}^{NewProd} + G_{it}^{Transfer}\right)} = \beta_0.$$

Nakamura and Steinsson (2014), AGM and others use a Bartik-type shock to construct such an instrument and, consistent with the logic above, find that IV estimates are substantially greater than OLS estimates.

Now suppose we have an economy in crisis. In this environment, output is determined by.

$$Y_{it} = \alpha_1 G_{it}^{Transfer} + \beta_1 G_{it}^{NewProd} + \epsilon_{it}$$

where $\alpha_1 > \alpha_0$ and $\beta_1 > \beta_0$. In other words, the multipliers could be state-dependent and increase in a crisis. For example, $\beta_1 > \beta_0$ could hold because there is less crowding out of private spending (supply curves are convex) and $\alpha_1 > \alpha_0$ could arise because financial frictions make cash flows affect output. It follows that, if the composition of transfer and new-production shocks is roughly constant across states, $c_1^{OLS} > c_0^{OLS}$ where c_1^{OLS} comes from estimating the following regression:

$$Y_{it} = c_1 G_{it} + u_{it}. \tag{4}$$

That is, the OLS estimate of the multiplier can increase as the economy moves from a normal state to a crisis state. However, this does not mean that the OLS recovers the state-dependent multiplier on new production (i.e., β_1). The OLS estimate continues to be biased down ($c_1^{OLS} < \beta_1$) and, in this example, represents a weighted average between α_1 and β_1. Again, β_1 can be recovered if one can filter out $G_{it}^{NewProd}$ from G_{it} with an instrumental variable or via some other way and similarly for α_1, which may be of independent interest.

Focusing on the cities that were not additionally constrained by SAH orders, we can combine regressions (3) and (4) into

$$Y_{it} = \mathbb{I}\{t = normal, not\ \ locked\ \ down\} \times b_0 G_{it} + \mathbb{I}\{t = crisis, not\ \ locked\ \ down\} \times b_1 G_{it} + error. \tag{5}$$

If the state of the economy $\mathbb{I}\{t = state, not\ \ locked\ \ down\}$ is exogenous (a likely scenario given the nature of the COVID19 pandemic) and unrelated to $G_{it}^{Transfer}$ and $G_{it}^{NewProd}$ (a plausible scenario given that government spending does not seem to change materially with COVID19 or lockdown),[9] OLS estimates of specification (5) provide us with $b_0^{OLS} \approx c_0^{OLS}$ and $b_1^{OLS} \approx c_1^{OLS}$. In particular, we are likely to get c_1^{OLS} for the early months of COVID (financial markets and the broader economy are disrupted). As the economy recovers from the initial shock, our estimate should decline toward c_0^{OLS}.

To get a sense of the ability of OLS to detect state-dependent fiscal effects in the presence of wealth transfers (and government spending that potentially responds endogenously to local economic conditions), we revisit the DLM study and report both OLS and IV estimates. Note that in this exercise, we use LAUS employment as the dependent variable, while

[9] For example, we observe (Fig. 4) little to no correlation between changes in DOD spending $\left(\frac{\Delta G_i}{Payroll_i^{Pre}}\right)$ and the intensity of SAH orders which can proxy for the relative ferocity of the local COVID19 crisis.

Table 4
Employment multipliers before and during the Great Recession.

	Dependent variable: Employment growth over the relevant period			
	IV		OLS	
	(1)	(2)	(3)	(4)
Panel A: 2006/07 to 2008/09				
$\Delta DOD_{0607\to0809}$	0.153***	−0.257	0.090***	−0.034
	(0.053)	(0.164)	(0.027)	(0.063)
$\Delta DOD_{0607\to0809} \times Debt_{06}$		0.287**		0.094**
		(0.133)		(0.044)
$Debt_{06}$		−0.019***		−0.017***
		(0.005)		(0.004)
N	827	823	827	823
R^2			0.39	0.42
Panel B: 2002/03 to 2004/05				
$\Delta DOD_{0203\to0405}$	0.206*	1.041	0.047	0.104
	(0.109)	(1.055)	(0.036)	(0.206)
$\Delta DOD_{0203\to0405} \times Debt_{02}$		−0.629		−0.044
		(0.749)		(0.152)
$Debt_{02}$		0.034***		0.030***
		(0.009)		(0.009)
N	827	823	827	823
R^2			0.20	0.24

Notes: This Table replicates the employment growth specification in Tables 3 and 8 from Demyanyk et al. (2019), with their measure of employment (based on the QCEW) replaced with ours (based on LAUS). The change in DOD spending over the indicated time periods is normalized by pre-period employee earnings. Consumer debt is based on the measure constructed by Mian et al. (2013) and includes mortgages, auto loans, credit card debt, and other forms of consumer debt.

DLM use QCEW earnings. Also note that we have a Bartik-style instrument for the period covered in DLM because aggregate DOD spending had significant variation. Table 4 reports the results from both the period around the Great Recession (Panel A) and the expansionary period in the early-to-mid-2000s (Panel B) reported in their study. Columns 1 and 2 report IV estimates, while columns 3 and 4 report OLS estimates. Notably, the estimates in columns 1 and 2 exhibit the same pattern as reported in DLM: average multipliers estimated using IV (column 1) are economically meaningful and statistically larger than zero. The employment effects are larger in high-consumer-debt cities only during the Great Recession period (when slack was high).

The OLS estimates demonstrate that both the estimate of the average employment effect (column 3) and the estimate of state-dependence (column 4) are biased downward. Nonetheless, OLS can detect a statistically significant employment effect, *but only during the Great Recession*. Comparing the estimates in column 3 across panels A and B imply that OLS can detect multipliers that vary over the business cycle (even if the magnitude is not precise), which is consistent with the view that "wealth transfers" from the government and more generally cash inflow can affect employment when financial markets and the broader economy are disrupted. Furthermore, OLS can detect state-dependence when it exists (if one interprets the IV estimate as the "true" measure of state-dependence), as the estimate of the interaction term in column 4 is positive and significant only in Panel A. These results suggests that, even if imperfect, OLS estimates can be useful to gauge the effectiveness of government spending in creating/saving jobs during crisis times.

This analysis implies that OLS can capture both a difference across periods (e.g., recessions versus expansions) and well as across states of local economies during a period of time. We will rely on this feature of OLS in interpreting our evidence. While our specification does not produce unbiased estimates of the effects of unanticipated DOD spending in normal times, we proceed under the assumption that it can provide evidence of state-dependence if such state-dependence exists. To anticipate our results, because the estimated multipliers in March 2020 are close to and indistinguishable from zero, the April 2020 multipliers can implicitly be viewed as a triple difference-in-difference. This helps to address more typical endogeneity concerns related to low-frequency changes in local government spending and economic conditions.

We take additional steps to mitigate potential endogeneity concerns. Because SAH orders were not entirely randomly assigned across cities (e.g., larger cities were more likely to be in a lockdown), we include CBSA-level controls in specification (1) to absorb potential determinants of employment changes that are correlated with DOD spending changes or vulnerability to the pandemic-induced recession. Following Mian and Sufi (2015) and DLM, these controls include: the percentage of white people in the local population, median household income, median home values, the percentage of owner-occupied housing units, the percentage with less than a high school diploma, percentage with only a high school diploma, the unemployment rate, a dummy for urban areas, the poverty rate, and the (log) of the local population. Our relevant controls are measured as of 2010 and are based on Census data when available or 3-year averages from the American Community Survey

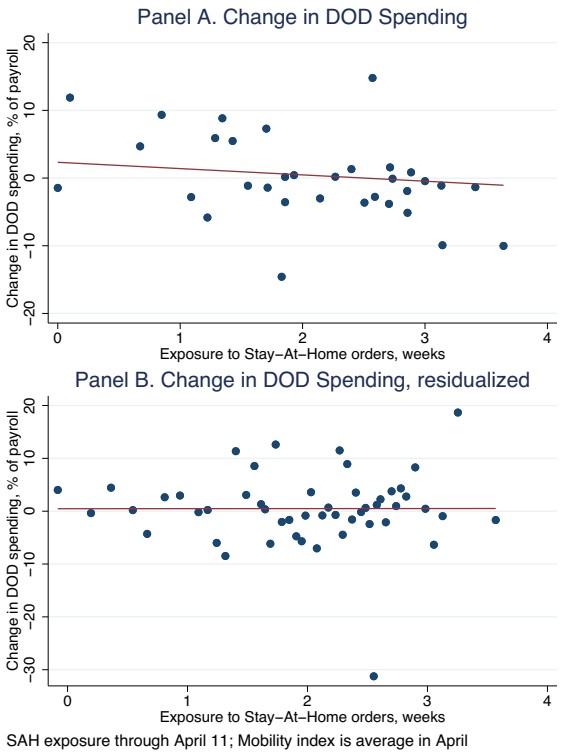

Fig. 4. Stay-At-Home (SAH) orders and change in Department of Defense (DOD) spending. Notes: Each panel presents a binscatter for the exposure to stay-at-home (SAH) orders (measured in weeks) vs. the change in Department of Defense (DOD) spending normalized by 2019 payroll across CBSAs. The top panel does not control for any CBSA characteristics. The bottom panel plots the binscatter after controlling for city size and other city characteristics. The list of controls corresponds to the list of controls used in specification (1).

(as reported by the National Historical Geographic Information System).[10] The list of controls includes the indicator for being locked down in April 2020.

A potential concern in our setting is that locked-down cities may be more or less likely to receive increases in DOD spending. Under certain conditions, controlling for lockdown status prevents any further bias in estimate of the effect of DOD spending across locations. However, if there are nonlinear effects of DOD spending, then interacting it with a correlated variable will introduce additional bias. As we discuss above, Fig. 4 suggests that changes in DOD spending (especially conditional on CBSA covariates) are independent of the severity of SAH orders and hence this concern is unlikely to be quantitatively important.[11]

3.2. Main results

Table 5 reports the results from our main specification (1) of the effect of DOD spending on employment during the pandemic month of April 2020 as well as results from various specification changes. For ease of interpretation, we measure annualized government spending in millions of dollars so that reported employment multipliers indicate the estimated number of job years created per $1 million in DOD spending in April. We report separately the estimates for unrestricted cities ($\widehat{\beta}_1$) and for locked-down cities ($\widehat{\beta}_1 + \widehat{\beta}_2$), along with standard errors. In our baseline analysis we classify cities as "high-SAH" using an intensity cutoff of 0.75 weeks of SAH orders during April.

The April employment multiplier for unrestricted cities of 22.4 reported in column (1) implies that it takes approximately $50,000 of spending to create (or save) a job-year in the month of April.[12] While it is typical to report the cost of a job per year,

[10] We explored using other/additional controls (e.g., the share of DOD spending in the local economy to address the possibility that employment changes may be related to the structural employment mix of the population rather than DOD spending changes) and found similar results.

[11] One may also be concerned about other government spending being correlated across CBSAs with the intensity of the COVID crisis and/or DOD spending, which can confound our estimates of fiscal multipliers. For example, PPP funding is negatively correlated with COVID cases (Granja et al., 2022). While we are not aware of any anti-COVID fiscal program being tied to DOD spending, we note that our focus on April 2020 should alleviate these concerns because the fiscal support programs largely reached the economy after April 2020. For example, Coibion et al. (2020a,b) report the timing of EIP checks received by households and find that most people got their checks after the relevant reference week in the CPS (that is, the week that includes the 12th of the month).

[12] Converting to a month-on-month basis is as simple as dividing the annual cost of creating a job-year by twelve. To see this note that if we did not annualize government spending then the regression coefficient would represent the number of jobs in April for each additional million dollars of spending in April, which is simply twelve times the coefficients of interest in our main specification.

Table 5
Baseline result: employment multipliers.

VARIABLES	Full sample					CBSAs by population size, April		
	April (1)	April-June (2)	April (3)	April (4)	April-June (5)	50 K or less (6)	50 K or more (7)	100 K or more (8)
$\frac{\Delta DOD}{Payroll} \times 1\{low\ SAH\}$	22.43***	13.83***	21.58*	23.32***	13.64***	29.85*	25.25***	25.59***
	(5.47)	(2.90)	(11.83)	(6.34)	(3.95)	(16.02)	(4.76)	(5.17)
$\frac{\Delta DOD}{Payroll} \times 1\{high\ SAH\}$	−1.16	−1.24	1.16	4.615	−0.201	4.70	−1.55	2.19
	(1.91)	(1.58)	(1.79)	(4.32)	(3.39)	(5.49)	(2.00)	(1.60)
Method	OLS	OLS	Huber	NN match	NN match	OLS	OLS	OLS
Observations	939	939	939	199	199	331	608	387
R^2	0.250	0.234	0.255	0.417	0.399	0.340	0.225	0.385

Notes: the table reports estimates of specification (1) for various samples and by various methods. Column (3) uses Huber-robust regression. Columns (4) and (5) use the nearest neighbor matching estimator. The top row shows employment multipliers for cities with low exposure to stay-at-home (SAH) orders, i.e., no lockdown group. The bottom row reports employment multipliers for cities with high exposure to SAH orders. Employment multipliers are measured as the number of job-years created by $1 million of Department of Defense (DOD) spending. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

one should bear in mind that the COVID crisis was unfolding at a dramatic pace. As a result, inferring a job-year from the April response may do too much extrapolation from April 2020 and it could be more appropriate to report the cost of a job on the monthly basis. In this case, our estimate implies that it takes approximately $4,000 of spending in a month to save a job for a month. Locked-down cities exhibit no detectable effect of DOD spending on employment (row 2 of column 1). These results imply the effect of DOD spending on local employment strongly depends on whether the local economy was under SAH orders (a strong state-dependent effect of DOD spending).

How does our estimate of the employment multiplier for unrestricted CBSAs compare to other employment multiplier estimates from the literature? Most cross-sectional studies of employment multipliers are conducted at the state-level and measure employment over a longer recession horizon than in our study. For example, Wilson (2012) finds that it takes approximately $100,000 of state-level American Recovery and Reinvestment Act (ARRA) spending to create a job over the twelve-month period prior to February 2010 (alternatively, $8,333 of spending in a month to create a job for a month). Chodorow-Reich et al. (2012) examine Medicaid spending associated with the ARRA and document that it takes approximately $25,000 of state-level spending to create one job-year ($2,083 to create a job for a month), which is at the upper-end of the estimates of the effects of fiscal stimulus (see Chodorow-Reich, 2019 for a review). Because of spillover effects (McCrory, 2020), multipliers tend to increase in the size of economic geography and so one might expect our estimated effects at the CBSA level to be lower than those based on state-level data. Perhaps the most comparable estimates to ours are in Dupor and McCrory (2018), who find that it takes between $67,000 and $100,000 to create a job-year at the commuting-zone level. Even stronger effects are reported by Suárez Serrato and Wingender (2016), who find that it takes only $33,000 in federal spending to create a job at the county level (e.g., $1 million creates 30 job-years).

Our estimate that it takes $4,000 of spending in a month to create a job in that month is toward the upper (lower) end of the effectiveness (spending cost) of government-induced job creation. It is potentially a downward biased estimate of the effect of "new production" DOD spending (see the discussion in Section 3.1) but also an estimate that can inform us about "wealth transfer" multipliers in crisis (α_1), an object of independent interest. One possible reason for a large effect of stimulus is that the sharp recession pushed many firms to the brink of exit. As discussed in Auerbach et al. (2021), in this environment relatively small changes in marginal firms' revenues can have large effects on their entry/exit decision and therefore large employment effects. Auerbach et al. (2020a), for example, document that DOD spending has a large effect on firm entry on average over the business cycle. A similar mechanism may operate through workers rather than firms. If the recession pushed many contractor workers' earnings up against their fixed costs of remaining in the workforce, then small changes in revenues can have large effects on employment.

To explore the persistence of the employment effect, we report (column (2)) the effect of DOD spending changes in April 2020 on the average monthly employment growth from 2019Q2 (April-June) to 2020Q2.[13] The estimated effects are smaller: converting the estimate to a monthly rate of spending, it takes about $6,000 to create (save) a job for a month in 2020Q2. This reduction in the employment multiplier is consistent with the fact that the economy improved over the course of 2020Q2, removing slack in the economy. And as documented by Auerbach and Gorodnichenko (2012), Auerbach et al. (2020b), and Demyanyk et al. (2019), fiscal multipliers tend to be increasing in the amount of slack in the economy.

3.3. Robustness

A potential threat to our interpretation of the estimates is that lockdown status is correlated with other CBSA-level determinants of fiscal effects. For example, AGM document that fiscal multipliers are larger in bigger cities, and in our setting

[13] In these specifications, we retain the same classification of locked down and not locked down cities – a classification which is based on SAH orders implemented through April 11th, 2020.

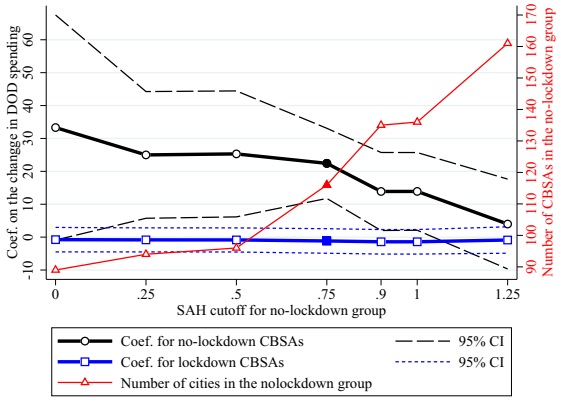

Fig. 5. Coefficient on DOD spending as a function of SAH cutoff. Notes: The figure shows the sensitivity of estimated coefficients β_1 (black line; no lockdown) and $\beta_1 + \beta_2$ (blue line; lockdown) in specification (1) to alternative cutoffs (in terms of the duration of stay-at-home (SAH) orders; SAH is measured in weeks) used to define the group of lockdown (restricted) cities. The red line shows the number of cities (CBSAs) classified as being in a lockdown. Filled markers show the baseline cutoff. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

locked-down cities tend to be larger. To address the possibility that other CBSA characteristics are driving our differential multiplier estimates, columns (4) and (5) in Table 5 reproduce the estimates from columns (1) and (2) but using a nearest neighbor estimator that averages the treatment effect of DOD spending across CBSAs that are similar in terms of their observable characteristics (the control variables). The matching estimates are similar to the baseline, which suggests that estimates are not driven by the joint distribution of SAH orders and other CBSA characteristics. To further examine how city size affects our estimates, columns (6) through (8) report estimates from specifications that restrict the sample based on population size. The estimates are similar across CBSA sizes.

Given the potential for outliers in the distribution of DOD changes (recall that Table 1 shows strong concentration of DOD spending), we also report estimates based on a Huber loss function that minimizes the influence of extreme observations (column (3)). The point estimates are nearly identical to the baseline estimates in column (1) and hence it is unlikely that our results are driven by a handful of CBSAs with unusual characteristics or changes in DOD spending.

In our baseline analysis we used a cutoff of 0.75 weeks of lockdown orders for the "high-SAH" designation, guided in part by the presence of distribution modes on either side of this level. A threshold of 0.75 weeks balances two competing objectives: *i*) to have a group of cities unaffected by SAH orders as large as possible to maximize statistical power; *ii*) to avoid contamination of this group with the cities that are affected by SAH orders. Fig. 5 reports the robustness of our baseline estimate to alternative thresholds along with the number of CBSAs classified as "unrestricted" for each threshold. As the threshold decreases below 0.75, the estimated multiplier increases (consistent with a diminished influence of SAH orders on "unrestricted" CBSAs), but so do the standard errors (consistent with a smaller sample size). As the threshold exceeds 1 week, a large mass of cities is pulled into the no-lockdown group and the estimate is attenuated toward zero.

3.4. Placebo tests: are there differential multipliers prior to the pandemic?

Our interpretation of the results from these various specifications is that the differential multiplier is driven by SAH orders, or the pandemic conditions associated with the issuance of such orders, i.e., high rates of infection, etc. (rather than other characteristics of cities that instated SAH orders). To further test this interpretation, we estimate our main specification (using the same SAH indicator) during the months prior to the pandemic. Since SAH orders had not been implemented (much less even considered a possibility), we expect to find no differential employment multiplier in these prior periods. This exercise illustrates more clearly that our design is analogous to a standard difference-in-difference approach.

Fig. 6 shows the results from cross-sectional estimates in the months before and after the onset of the pandemic. In the time prior to the pandemic there is no distinguishable difference between location types. In April, the estimated employment multiplier for unrestricted cities jumps, while it remains flat for restricted cities. This pattern mimics the pattern from the DLM OLS estimates, whereby the estimated multiplier varied across cities during the recession period but was indistinguishable from zero in the prior period. Our estimates for unrestricted cities vary little over time.

4. Exploring the channels

What are the underlying mechanisms that can account for higher multipliers in a recession only for unrestricted cities? The literature offers two mechanisms for stronger employment multipliers following a surge in unemployment. First, the slack channel implies that producers are more able to translate increases in demand into new employment and output when

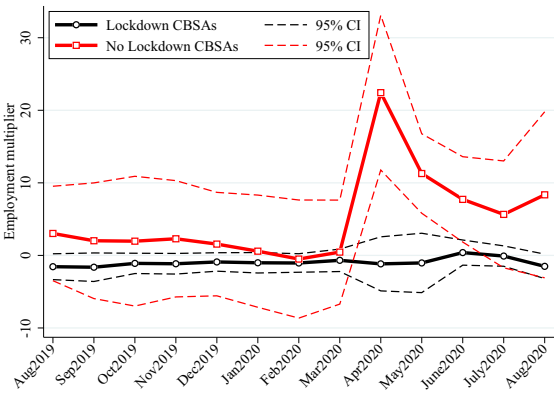

Fig. 6. Placebo Results. Notes: The figure shows the sensitivity of estimated coefficients β_1 (red line; no lockdown) and $\beta_1 + \beta_2$ (black line; lockdown) in specification (1) to alternative definitions of events. The baseline event is April 2020. Dashed lines show 95 percent confidence intervals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

local aggregate demand is initially low and so inputs are idle. The slack channel encompasses a number of potential underlying mechanisms. When demand is low, it may be easier to hire from the larger pool of unemployed workers (Michaillat, 2014). Alternatively, firms and workers may be able to accommodate more production without running into capacity constraints (e.g., Michaillat and Saez, 2015; Murphy, 2017; Auerbach et al., 2020a). A third possibility is that low aggregate demand pushes many producers' revenues down toward their fixed operating costs such that they are on the brink of exit, which underscores the importance of cash flows and potential interest in estimating transfer multiplier α_1. Additional changes in revenues (positive or negative) can have large effects on output and employment in such an environment (Auerbach et al., 2021). Broadly speaking, these slack channels predicts that employment and output are more responsive to DOD spending during a recession even if government spending does not stimulate private spending (e.g., the local supply curve is convex).

Second, households may have higher average MPCs during a recession, as tightened credit conditions and lower household income render more households credit-constrained (e.g., Eggertsson and Krugman, 2012), but the increase in MPCs may be limited in restricted locations. The high-MPC channel predicts that DOD spending increases household consumption by more during a recession with rationed credit. If enough of the new consumption is spent on local services (rather than exclusively on imported tradable goods), then the higher consumption can translate into higher local employment. However, if a subset of consumption sectors is restricted (e.g., workers cannot dine away from home), then according to recent theoretical work (Auerbach et al. (2021) and Guerrieri et al. (2020)), consumption in those locations may respond less to fiscal stimulus.

To start to disentangle these two mechanisms we turn to data on consumption. We note that a positive consumption response during the peak of the recession does not isolate the MPC channel: even the slack channel can predict strong consumption responses during a recession, as the newly hired employees spend their new income (if they have positive MPCs). And even if local consumption responds, whether it stimulates local employment depends on the extent to which local goods are included in the set of new consumption. Some pre-existing evidence suggests that this channel is small at the local level. For example, AGM document that CBSA-level local general equilibrium spillover effects (such as those arising from positive MPCs) are positive yet quantitatively small on average over the business cycle. Similarly, Demyanyk et al. (2019) document that heterogeneous (across-CBSA) fiscal multipliers are most apparent in tradable sectors that are less likely to directly benefit from local consumption.[14] In any case, a lack of a consumption response poses a substantial hurdle to the MPC mechanism.

4.1. The consumption channel

To measure the response of consumer spending, we modify specification (1) as follows:

$$\log\left(\frac{Y_i}{Y_i^{Pre}}\right) = \alpha + \beta_1 \frac{\Delta G_i}{Payroll_i^{Pre}} + \gamma State_i + \beta_2 \frac{\Delta G_i}{Payroll_i^{Pre}} \times State_i + Controls_i + \epsilon_i \tag{6}$$

where Y_i is a measure of consumer spending, expressed in percent deviations from pre-COVID19 levels so that the estimated coefficients are semi-elasticities. Table 6 reports results from the baseline specification (1) with retail mobility (columns (1)-(3)) or retail consumer spending (columns (4)-(6)) as the dependent variable. There is no clear evidence of a positive consumption response to DOD spending in areas with high or low exposure to SAH orders. Since there is no detectable differ-

[14] Demyanyk et al. (2019) also exploit household-level variation in consumer debt levels to document higher consumption responses to local defense spending. What remains unclear is whether the sum of these household-level consumption changes has a quantitatively meaningful effect on employment in the local service sector.

Table 6
Baseline result: retail mobility and spending.

VARIABLES	Retail mobility in April 2020			Retail consumer spending in April 2020		
	Full (1)	50 K or less (2)	50 K or more (3)	Full (4)	50 K or less (5)	50 K or more (6)
$\frac{\Delta DOD}{Payroll} \times 1\{low\ SAH\}$	0.54	−1.25	3.27	2.26	−12.53	1.64
	(3.00)	(9.00)	(2.95)	(5.25)	(47.40)	(6.44)
$\frac{\Delta DOD}{Payroll} \times 1\{high\ SAH\}$	−1.41	−2.32	−0.95	−0.32	−4.81**	0.47
	(1.35)	(1.96)	(0.85)	(0.63)	(1.69)	(0.32)
Method	OLS	OLS	OLS	OLS	OLS	OLS
Observations	939	330	609	729	193	536
R^2	0.60	0.52	0.57	0.14	0.16	0.17

Notes: The table reports estimates of specification (2) for various samples. The outcome variables are retail mobility or retail consumer spending in April 2020. All outcome variables are measured as percent deviations from pre-COVID19 levels. The top row shows results for cities with low exposure to stay-at-home (SAH) orders, i.e., no lockdown group. The bottom row shows results for cities with high exposure to SAH orders. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

ential consumption response to DOD spending between unrestricted and locked-down cities, we therefore cannot claim to find support for the high-MPC channel. Hence, the mechanism should rely on a reduction in MPCs for both restricted and unrestricted cities. This logic points to an aggregate factor that inhibits MPCs.

Using a survey of households participating in the Nielsen Homescan Panel, Coibion et al. (2020a) find that households reported that they would save the majority of their economic stimulus payments issued shortly after April 1, 2020. Even months into the recession, households primarily had saved rather than spent their income transfers and the MPCs in the current downturn were lower than MPCs for stimulus payments in the 2001 recession (Shapiro and Slemrod, 2003) or the Great Recession (Sahm et al., 2015). The low spending propensities out of transfer income documented in Coibion et al. (2020a) are consistent with the lack of a consumption response that we find in our data. Furthermore, Coibion et al. (2020a) find that whether a household is in a locked down city does not have a material impact on its MPC.

The apparent fall in the MPC nationwide and insensitivity to SAH exposure may arise for a variety of reasons: consumers could voluntarily avoid spending that might increase health risks (although one might expect this to be associated with SAH intensity); consumers could increase precautionary savings in response to a dramatic rise in uncertainty; consumers could perceive a nearly permanent effect of the recession on their future earnings; or consumers could be pushed to service their debt obligations before they can increase consumer spending.[15] To further understand the consumption response during this very unusual episode, it will be helpful to have additional data on consumer spending.

In any case, it appears that MPCs have been low during the pandemic even if the employment response to DOD spending was more animated in unrestricted (low SAH exposure) cities. These results suggest that while DOD may directly increase employment, it is less likely that this additional hiring generates strong second-round effects via private consumption in this exceptional environment.

4.2. The slack channel

Our evidence so far in support of the slack channel is based on differences in restrictions across cities in average fiscal effects: for unrestricted cities, employment effects of DOD spending were higher during the peak of the economy-wide recession. Here we exploit cross-sectional differences in slack to further isolate its role in driving heterogeneous fiscal effects. In particular, we compare the effects of DOD spending across locations with different levels of unemployment as of March 2020. To account for the possibility that different cities have different levels of structural unemployment, we define city-level slack as the difference between unemployment in March 2020 and its level as of 2010. Intuitively, while the labor market was apparently tight at the national level just before the pandemic struck (the unemployment rate was 3.4 and 4.4 in February and March 2020), there was a lot of regional variation in local unemployment rates across CBSAs (the standard deviation in 1.6 and 2.0 percent in February and March 2020).[16] We use the 2010 unemployment rate (from the American Community Survey which had much larger sample size than the CPS) as a benchmark for the unemployment rate in a severe recession. Cities that have above median value of the difference are classified as higher slack locations.

[15] Miranda-Pinto et al. (2020b) show that short-run MPCs near zero can be rationalized in a model with time-varying minimum consumption thresholds. These thresholds represent, for example, medical emergencies or auto repairs that are associated with large utility costs if not addressed. More generally, they represent stochastic maintenance costs for pre-committed consumption (Chetty and Szeidl, 2007). While rich households can pay for the maintenance costs out of their income or wealth, poorer households must take on debt that, if not repaid, will render them unable to afford any future adverse expenditure shocks. Adverse shocks (such as a recession) push many poor households against their minimum consumption threshold. These households use any additional income to save in the short-run before gradually increasing their consumption. The behavior predicted by their model is consistent with the lack of a short-run consumption response to DOD spending in the midst of a recession.

[16] Another option to measure the degree of slack in the early months of the COVID19 crisis is to look at the degree of wage pressure in a city. While generally intuitive, this alternative is less attractive in our case because the crisis was unfolding rapidly and wage rigidities mute the response of prices in the labor market so that we are not likely to discern variation in the slack across cities.

Table 7
Heterogeneity in the sensitivity of employment multipliers to changes in DOD spending in low-SAH cities.

	Dependent variable: Employment growth in 2020	
	April (1)	March (2)
Panel A. High/Low Slack		
$\frac{\Delta DOD}{Payroll} \times 1\{low\ \ SAH\} \times 1\{lower\ \ slack\}$	19.89***	−1.37
	(4.01)	(4.67)
$\frac{\Delta DOD}{Payroll} \times 1\{low\ \ SAH\} \times 1\{higher\ \ slack\}$	30.10**	−0.14
	(9.65)	(13.44)
$\frac{\Delta DOD}{Payroll} \times 1\{high\ \ SAH\}$	−1.17	−0.65
	(1.91)	(0.78)
Observations	939	939
R-squared	0.25	0.10
Panel B. Increases/Decreases in DOD spending		
$\frac{\Delta DOD}{Payroll} \times 1\{low\ \ SAH\} \times 1\{\Delta DOD \geq 0\}$	14.64**	−2.57
	(4.94)	(5.05)
$\frac{\Delta DOD}{Payroll} \times 1\{low\ \ SAH\} \times 1\{\Delta DOD < 0\}$	49.42*	10.95
	(26.22)	(6.59)
$\frac{\Delta DOD}{Payroll} \times 1\{high\ \ SAH\}$	−1.17	−0.68
	(1.93)	(0.81)
Observations	939	939
R^2	0.25	0.09

Notes: The table reports estimates of specification (3) for alternative event dates (April 2020 is the baseline, March 2020 is the placebo). Employment multipliers are measured as the number of job-years created by $1 million of Department of Defense (DOD) spending. Heteroskedasticity robust standard errors are reported in parentheses. ***, **, * denote statistical significance at 1, 5 and 10 percent levels.

$$\frac{\Delta Y_i}{Payroll_i^{Pre}} = \alpha + \beta_1 \frac{\Delta G_i}{Payroll_i^{Pre}} \times 1(\text{no lockdown}_i) \times 1(Lower\ \ Slack_i) + \beta_2 \frac{\Delta G_i}{Payroll_i^{Pre}} \times 1(\text{no lockdown}_i) \times 1(Higer\ \ Slack_i) + \beta_3 \frac{\Delta G_i}{Payroll_i^{Pre}} \times 1(\text{lockdown}_i) + \gamma_1 \times 1(\text{no lockdown}_i) \times 1(Lower\ \ Slack_i) + \gamma_2 \times 1(\text{no lockdown}_i) \times 1(Higer\ \ Slack_i) + \gamma_3 \times 1(\text{lockdown}_i) + Controls_i + \epsilon_i \quad (7)$$

Panel A, Column 1 of Table 7 reports employment multipliers separately for unrestricted cities with high levels of slack and for unrestricted cities with low levels of slack (as well as for restricted cities). The point estimate for high- slack cities (30.10 job-years per $millions of DOD spending) is over fifty percent larger than the estimate for lower-slack cities, although the difference is not statistically significant (p-value = 0.28). These results suggest that unrestricted cities entering the COVID19 crisis with weaker labor markets tended to have larger fiscal multipliers than unrestricted cities with stronger labor markets. This result points to slack as a potential mechanism for how a fall in aggregate private demand (and hence the corresponding increase in unemployment/slack) in unrestricted cities could be mitigated by increased government spending.

When we examine employment effects as of March 2020 (and slack as of February 2020), we find no differential effect (column 2). This result suggests that government spending could be ineffective in absorbing "initial" slack or that variation in slack (or changes in DOD spending) is too small to discern the differential fiscal effects in high- vs. low-slack areas, particularly given that our local measures of slack relative to 2010 values may incorporate some error.

As an independent test of the slack channel, we examine whether decreases in DOD spending have a larger effect than increases in spending. Theories of slack imply convex supply (AS) curves such that increases in aggregate demand push the economy against its capacity, therefore leading to relatively muted employment and output effects. Decreases in aggregate demand move the economy away from its capacity level and can have large employment effects. To test this implication, we isolate increases in DOD spending from decreases (Panel B of Table 7) for unrestricted CBSAs. The point estimate for decreases is 49.42, over three times the size of the estimate for increases. This differential effect of spending decreases during a recession is similar to the findings in Barnichon et al. (2020) using aggregate time series data. The estimates based on employment changes in March follow a similar pattern but are noisy, which could reflect the possibility that the economy was near full employment prior to the onset of the pandemic – that is, the economy was on a steep portion of the AS curve such that changes in aggregate demand had muted effects on employment. As the economy entered the recession and moved along a flatter portion of AS, there was more room for AD to affect employment (and particularly so for DOD spending decreases).

5. Conclusion

The effect of fiscal stimulus has become an increasingly relevant topic since the Great Recession. How effective are various forms of fiscal stimulus, and under what conditions? Is government spending (more) effective during a recession, and if so why? The speed and ferocity of the pandemic recession coupled with the differential restrictions imposed across cities in response to the pandemic provide a unique opportunity to understand the mechanisms underlying fiscal stimulus.

Our evidence points to the role of slack in generating higher employment multipliers during the pandemic recession for cities not subject to stay-at-home (SAH) orders. For locked-down cities, employers appear to have been prevented from hiring in response to DOD contracts, thereby shutting down the ability of the local economy to absorb labor market slack. We find no evidence that DOD spending increased local consumption (in unrestricted or locked-down cities), which casts doubt on mechanisms that rely on high-MPC households. Rather, government spending appears to be an effective means of directly increasing employment during a recession, particularly when the hiring is directed toward places that are not restricted by SAH orders. However, we cannot rule out that there were special conditions that applied during the pandemic recession that might also have limited responses through the consumption channel.

More generally, our analysis suggests that the nature of economic downturns is potentially important for the effectiveness of government spending in stimulating aggregate demand. Indeed, widespread restrictions on economic activity likely affected the supply side, thus limiting the ability of the economy to absorb government spending into new production, which contrasts with the experience in standard, demand-driven recessions. To the extent that post-COVID supply-side constraints and bottlenecks continue to be a major factor, one may anticipate that government spending multipliers may be lowered.

We focus on government spending multipliers and hope that future research can make progress in quantifying the effectiveness of counter-cyclical fiscal policies such as the Economic Impact Payment and Payroll Protection Program designed specifically to address the COVID19 crisis. Because the slack channel encompasses a number of underlying mechanisms, another fruitful avenue for future research is to explore which mechanisms are more relevant with the ultimate goal of informing the relative effectiveness of various targeted fiscal policies.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Adrien Auclert, Gabriel Chodorow-Reich, Sarah Zubairy and ASSA-2021, AEPC-2021, LACEA LAMES 2021, and BU Macro seminar participants for comments on an earlier version of the paper. This paper is not a product of the Research Department of J.P. Morgan Chase. The views expressed here reflect those of the authors only and may not be representative of others at J.P. Morgan. For disclosures related to J.P. Morgan, please see https://www.jpmorgan.com/disclosures.jsp.

References

Auerbach, A.J., Gorodnichenko, Y., 2012. Measuring the Output Responses to Fiscal Policy. Am. Econ. J. Econ. Policy 4 (2), 1–27.
Auerbach, A.J., Gorodnichenko, Y., Murphy, D., 2021. Inequality, Fiscal Policy, and COVID19 Restrictions in a Demand-Determined Economy. Eur. Econ. Rev. 137, 103810.
Auerbach, Alan J., Gorodnichenko, Yuriy, Murphy, Daniel, 2020a. Macroeconomic Frameworks: Reconciling Evidence and Model Predictions from Demand Shocks. NBER Working Paper 26365.
Auerbach, A.J., Gorodnichenko, Y., Murphy, D., 2020b. Local Fiscal Multipliers and Fiscal Spillovers in the United States. IMF Econ. Rev. 68, 195–229.
Baker, Scott, Farrokhnia, R.A., Meyer, Steffen, Pagel, Michaela, Yannelis, Constantine, 2020. How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic. Working Paper.
Baek, ChaeWon, McCrory, Peter B., Messer, Todd, Mui, Preston, 2021. Unemployment Effects of Stay-at-Home Orders: Evidence from High Frequency Claims Data. Rev. Econ. Stat. 103 (5), 979–993.
Barnichon, Regis, Debartoli, Davide, Matthes, Christian, rnichon et al., forthcoming. Understanding the Size of the Government Spending Multiplier: It's in the Sign. Rev. Econ. Stud. (forthcoming)
Brunet, Gillian, 2018. Stimulus on the Home Front: The State-Level Effects of WWII Spending. Manuscript.
Chetty, Raj, Friedman, John N., Hendren, Nathaniel, Stepner, Michael, 2020. How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. NBER Working Paper No. 27431.
Chetty, R., Szeidl, A., 2007. Consumption commitments and risk preferences. Quart. J. Econ. 122 (2), 831–877.
Chodorow-Reich, G., Feiveson, L., Liscow, Z., Woolston, W.G., 2012. Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act. Am. Econ. J. Econ. Policy 4 (3), 118–145.
Chodorow-Reich, G., 2019. Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned? Am. Econ. J. Econ. Policy 11 (2), 1–34.
Coibion, Olivier, Gorodnichenko, Yuriy, Weber, Michael. 2020a. How did U.S. Consumers Use Their Stimulus Payments? NBER Working Paper 27693.
Coibion, Olivier, Gorodnichenko, Yuriy, Weber, Michael. 2020b. Labor Markets During the COVID-19 Crisis: A Preliminary View. NBER Working Paper 27017.
Demyanyk, Y., Loutskina, E., Murphy, D., 2019. Fiscal Stimulus and Consumer Debt. Rev. Econ. Stat. 101 (4), 728–741.
Dupor, B., McCrory, P.B., 2018. A Cup Runneth Over: Fiscal Policy Spillovers from the 2009 Recovery Act. Econ. J. 128 (611), 1476–1508.
Eggertsson, G., Krugman, P., 2012. Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. Quart. J. Econ. 127 (3), 1469–11513.
Ghassibe, Mishel, Zanetti, Francesco, 2020. State Dependence of Fiscal Multipliers: the Source of Fluctuations Matters. Manuscript.
Guerrieri, Veronica, Lorenzoni, Guido, Straub, Ludwig, Werning, Iván, 2020. Macroeconomic Implications of COVID19: Can Negative Supply Shocks Cause Demand Shortages? NBER Working Paper No. 26918.

Granja, João, Makridis, Christos, Yannelis, Constantine, Zwick, Eric, 2022. Did the Paycheck Protection Program Hit the Target? Journal of Financial Economics, forthcoming.
Jo, Yoon J., Zubairy, Sarah, 2021. State Dependent Government Spending Multipliers: Downward Nominal Wage Rigidity and Sources of Business Cycle Fluctuations. Manuscript.
Klein, M., 2017. Austerity and Private Debt. J. Money Credit Bank. 49, 1555–1585.
McCrory, P.B., 2020. Tradable Spillovers of Fiscal Policy: Evidence from the 2009 Recovery Act. Mimeo.
Mian, A., Sufi, A., 2015. What Explains the 2007–2009 Drop in Employment? Econometrica 82 (6), 2197–2223.
Mian, A., Rao, K., Sufi, A., 2013. Household Balance Sheets, Consumption, and the Economic Slump. Quart. J. Econ. 128, 1687–1726.
Michaillat, P., Saez, E., 2015. Aggregate Demand, Idle Time, and Unemployment. Quart. J. Econ. 130 (2), 507–569.
Michaillat, P., 2014. A Theory of the Countercyclical Government Multiplier. Am. Econ. J. Macroecon. 6 (1), 190–217.
Miranda-Pinto, Jorge, Murphy, Daniel, Walsh, Kieran, Young, Eric, 2020a. Saving Constraints, Debt, and the Credit Market Response to Fiscal Stimulus. Mimeo.
Miranda-Pinto, Jorge, Murphy, Daniel, Walsh, Kieran, Young, Eric. 2020b. A Model of Expenditure Shocks. Mimeo.
Murphy, D.P., 2017. Excess Capacity in a Fixed-Cost Economy. Eur. Econ. Rev. 91, 245–260.
Nakamura, E., Steinsson, J., 2014. Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions. Am. Econ. Rev. 104 (3), 753–792.
Ramey, V., Zubairy, S., 2018. Government Spending Multipliers in Good times and in Bad: Evidence from U.S. Historical Data. J. Polit. Econ. 126 (2), 850–901.
Sahm, Claudia. R., Shapiro, Matthew D., Slemrod, Joel, 2015. Balance-sheet households and fiscal stimulus: Lessons from the payroll tax cut and its expiration. NBER Working Paper 21220.
Matthew, D., 2003. Shapiro, Joel Slemrod, Consumer response to tax rebates. Am. Econ. Rev. 93 (1), 381–396.
Suárez Serrato, Carlos, Juan, Wingender, Philippe, 2016. Estimating Local Fiscal Multipliers. NBER Working Paper No. 22425.

Further reading

Cox, Natalie, Farrell, Diana, Ganong, Peter, Greig, Fiona, Noel, Pascal, Vavra, Joe, Wong, Arlene, 2020. Initial Impacts of the Pandemic on Consumer Behavior: Evidence from Linked Income, Spending, and Savings Data. Brookings Papers on Economic Activity.
Guerrieri, V., Lorenzoni, G., 2017. 'Credit Crisis, Precautionary Savings and the Liquidity Trap'. Quart. J. Econ. 132 (3), 1427–1467.