

Late Adoption and Collective Action: Social Media Expansion and the Diffusion of Black Lives Matter

Annalí Casanueva Artís^{*} Vladimir Avetian^{®†}
Sulin Sardoschau^{®‡} Kritika Saxena^{®§}

December 2024

Abstract

This paper investigates the role of late adoption of social media on collective action, focusing on the Black Lives Matter (BLM) protests of 2020. Using a novel instrumental variable approach, we exploit the interaction between historical shocks to early Twitter adoption and quasi-exogenous variation in pandemic exposure, measured by super-spreading events (SSEs), to predict late Twitter adoption at the county level. We find that late adoption increases protest in support of BLM - both online and offline - but it does not facilitate right-wing or anti-BLM protest. Leveraging individual-level survey data, we find that late adopters hold more favorable views towards BLM and racial equity issues but not towards other progressive causes. Contrary to the existing literature on early adopters, our evidence suggests that late adoption drives mobilization through a change in preferences rather than (just) a reduction in coordination costs; potentially because late adopters encounter a politically consolidated and persuasive platform upon entry. Our findings highlight that the impact of social media platforms may change as they mature and their user base evolves.

Keywords: social media, collective action, protest, Twitter, BLM, late adopters

JEL classification: P16, D7

^{*}ifo Institute and Big Data Junior Research Group in Social Media annali.casanueva@gmail.com

[†]Paris Dauphine University, PSL, Chair of Governance and Regulation: vl.avetyan@gmail.com

[‡]Humboldt Universität zu Berlin : sulin.sardoschau@hu-berlin.de (corresponding author)

[§]University of Groningen: k.saxena@rug.nl The order of authors has been randomized as indicated by the superscript [®]. The authors thank Maja Adena, Hunt Allcott, Abhijit Banerjee, François Bourguignon, Eric Brousseau, Davide Cantoni, Ruben Durante, Vicky Fouka, Roberto Galbiati, Sergei Guriev, Pamina Koenig, Jonathan Lehne, Alexey Makarin, Benjamin Marx, Sultan Mehmood, Nikita Melnikov, Horacio Larreguy, Raúl Sánchez de la Sierra, Carlo Schwarz, Lore Vandewalle, Carlos Winograd, Maiting Zhuang, and Ekaterina Zhuravskaya for their useful comments and suggestions. We also thank Marlon Seror for his invaluable input. We are thankful to all the participants of the 25th Annual conference of the Society for Institutional and Organizational Economics (SIOE), the annual meeting of the European Economics Association (EEA-ESEM), the Households in Conflict (HiCN) conference, the AFSE conference, the Lisbon Meeting on Economics and Political Science, the brown bag seminar at The Graduate Institute of International and Development Studies, the Sciences Po internal seminar, the Empirical Economics Seminar at Ludwigs Maximilian University Munich, the Development Economics Network Berlin, the Applied Young Economists Webinar, the Milan Labor Lunch Seminar, the Berlin Applied Micro Seminar, and the SSES Annual Congress, participants of the ASSA San Antonio for their useful feedback.

1 Introduction

The trade-off between expanding reach and preserving unity stands at the core of many collective action problems (Ostrom, 1990; De Mesquita, 2010; Barbera and Jackson, 2019). Social media has emerged as a tool for outreach and coordination of protest (Zhuravskaya et al., 2020; Qin et al., 2024); particularly in the early stages of its roll-out (Manacorda and Tesei, 2020; Enikolopov et al., 2020). As social media platforms attract more users, the impact of late adopters on protest mobilization is ambiguous.

First, late adopters may be selected along dimensions that favor or hinder protest compared to early adopters. Second, late adopters face a different social media environment upon entry. This may curb mobilization if newcomers hold opposing views, are less engaged, or if a larger and more diverse network leads to information pollution. Conversely, late adoption may boost protest mobilization if existing networks facilitate learning about and joining a movement or alter newcomers' preferences more effectively.

We examine this question in the context of Black Lives Matter (BLM) protests in 2020, leveraging a novel instrument for the late adoption of Twitter at the county level.¹ Inspired by push-pull instruments in the migration literature, which combine historical migrant networks with economic shocks to predict the spatial distribution of new migrants (Burchardi et al., 2019), we combine the size of the historical Twitter network – established through an unrelated shock to early adoption – and quasi exogenous variation in pandemic exposure to predict the spatial distribution of new users.

In particular, we argue that the pandemic acted as a push factor by increasing local demand for social media because new users seek information, entertainment or community online. Specifically, we leverage quasi random variation in pandemic exposure through so-called super-spreading events (SSEs): the number of SSEs occurring within a 200 km radius of a county (but excluding the county itself) during the first 12 weeks of the pandemic. The pull factor draws on the concept of path dependence in technology adoption (Arthur, 1989; Bursztyn et al., 2023) which argues that larger networks offer stronger incentives for new users to adopt the platform. We leverage a historical shock to Twitter adoption spurred by attendance at the 2007 South by South-West (SXSW) festival which heavily promoted Twitter when it first launched, borrowing from Müller and Schwarz (2023). By interacting these two components, we capture the combined effect of heightened pandemic-driven demand and pre-existing network strength in predicting late adoption. Importantly, we control for both push and pull factor separately, capturing any direct effect of pandemic exposure and historical networks on protest for BLM.

We provide various pieces of evidence that support the plausibility of quasi-random exposure to the instrument. First, the interaction between SSE and SXSW is not correlated with baseline county characteristics which may influence the propensity to protest, including socioeconomic, demographic, and social capital variables. Second, the instrument is not correlated with past protest, or spatial diffusion patterns of protest measured by proximity to Minneapolis, the epicenter of the BLM movement. Third, there is no correlation with state-level lockdown policies

¹We use the terms new users and late adopters interchangeably, defined as Twitter accounts created after the outbreak of the pandemic in the U.S. and before the murder of George Floyd (January 21, 2020–May 24, 2020). We also refer to the social media platform X with its former name "Twitter" and use the terms "tweets" and "tweeting" for convenience.

or COVID-19-related racial disparities. Finally, event-study analyses confirm the temporal sequencing of the instrument’s components: Twitter adoption consistently follows, rather than precedes or coincides with SSEs, assuaging concerns that late adoption facilitates the detection of SSEs. We also verify that our instrument is not sensitive to narrowing or expanding the radius, time window, and location of SSEs (e.g., in prisons), or to different weighting schemes. We perform several additional robustness checks, which are presented after the main results below.

Since we are interested in the diffusion of BLM, we focus on the set of counties that has never protested for a BLM-related cause before and estimate a cross-sectional regression at the county level that includes state fixed effects and a battery of controls.² Our estimates show that late adoption mobilized subsequent offline and online BLM protest in counties that have never protested for a BLM-related cause before: a 1 percent increase in the number of late adopters increases the likelihood of experiencing a BLM protest in the three weeks following the murder of George Floyd by 0.4 percentage points and increases BLM-related tweets by late adopters by 0.8%, suggesting that online and offline protest are complements in our context.³

There are multiple explanations for our findings. The literature has emphasized the reduction in coordination costs as a primary driver of collective action; particularly in the early stages of its roll-out (Enikolopov et al., 2020; Manacorda and Tesei, 2020). However, this may change in the context of late adoption as these new users face a different social media environment upon entry. On the one hand, larger existing networks may still reduce coordination costs because ideological sorting and echo chambers prevent information pollution. In addition, established online communities can reduce coordination costs by offering extensive information on BLM narratives and logistics to late adopters. On the other hand, larger existing networks may favor persuasion as they provide a stronger signal of others’ preferences, particularly when posts about George Floyd are shared and liked by many people. This visibility and wealth of information may alter newcomers’ preferences more effectively.

To examine these possibilities, we check whether the surge in collective action is limited to causes in support of BLM. This allows us to assess whether the effect of late adoption is universal - even for new users that oppose the majority view on the platform.⁴ Using data from the ACLED US Crisis Monitor, we show that social media adoption intensifies anti-mask and anti-social distancing protest, which - maybe surprisingly - have been shown to appeal to the same political base as BLM protesters (Chenoweth et al., 2022). However, we do not find that protests that represent minority views on the platform at baseline surge in response to late adoption. In particular, the effects on protests organized by QAnon or "Proud Boys", as well as

²These include: the Twitter network at baseline (December 2019), COVID-19 related deaths at the county level (until May 25th 2020), as well as an array of socio-economic, political and demographic controls that may be related to late adoption and collective action. Economic controls: change in unemployment between 2019 and 2020, median household income in 2016, 3+ risk factors/community resilience, a dummy for urban counties. BLM controls: share of Black population, urban, Black poverty rate, and deadly force used by police against Black people. Political controls: Republican vote share in 2012 and 2016 and social capital (number of different types of civic organizations).

³In the context of social media use and protest in Russia, Enikolopov et al. (2020) find that a 10 percent increase in VK users (Russian equivalent to Facebook) increases the likelihood of protest by 4.6 percentage points.

⁴Newcomers to Twitter joined a platform with a well-established and prominent BLM network, where BLM was one of the leading hashtags even before 2020 (Keib et al., 2018). Twitter users, who were substantially more left-leaning than the general population conditional on age were further influenced by the fact that a minority of highly active users—disproportionately left-leaning—produced most of the tweets (Pew, 2019). This, combined with the viral circulation of protest imagery following George Floyd’s murder, created a platform environment overwhelmingly favoring a pro-BLM narrative.

other groups that support Trump (including perceived election fraud) are comparatively small and noisily estimated, indicating that new Twitter users did not mobilize right-wing protest. We also do not find an effect on counter-mobilization in the form of tweets for "All Lives Matter" and "Blue Lives Matter", which are both rallying cries in opposition to BLM.

Next, we investigate the possibility that the expansion of social media may have altered preferences, leveraging individual-level survey data. We predict individual-level late adoption with the interaction between the push-pull instrument and the age of the survey respondent. The idea behind this approach is that young individuals tend to be early adopters; older individuals tend to be late adopters, and they are the ones who will respond most to shocks to late adoption. We include county fixed effects and an array of individual level controls, effectively comparing differences in attitudes between older and younger respondents in counties that were more or less exposed to the push-pull shock.⁵ We provide evidence for the exogeneity of the individual-level instrument, showing that it does not predict other media consumption (newspaper or TV), COVID-19 salience or other progressive attitudes (abortion rights and environmental protection). Using data from the Cooperative Election Study (CES), we find evidence that social media use increases the likelihood of having attended a protest in the past year, becoming politically active on social media, voting for Biden in the 2020 presidential election and holding more progressive views towards the police and racial equity but not other progressive issues.⁶

Overall, our results show that the expansion of social media in its later stages modifies the drivers of collective action. Late adoption increases protest, but not uniformly. The mobilization potential of social media platforms is movement-specific and benefits causes with established networks. Unlike prior research on early social media adoption and protest (Enikolopov et al., 2020; Manacorda and Tesei, 2020), we find that late adoption drives mobilization by altering preferences rather than (just) reducing coordination costs. This suggests that the dominance of certain social movements on these platforms enables them to influence newcomers' attitudes more effectively than in less consolidated media environments.

We confirm the robustness of our results with an array of empirical exercises: results are robust to the spatial clustering of standard errors, alternative sample compositions (e.g., dropping coastal states and counties), model specifications (e.g., probit), controlling for changes in pandemic prevalence just before the protest trigger, expanding or narrowing the time window around protest, and controlling for ex-ante BLM protest probabilities using LASSO-selected controls. In addition, we complement our IV strategy with an event-study design, expanding the sample to all counties and collapsing the data to the county-month level. Using county and state-month fixed effects, we trace the evolution of BLM protest in places with higher levels of late adoption in the 16 months leading up to and 4 months after the unexpected protest-trigger. In line with our previous findings, we show that counties with more late adopters are on parallel trends before but go on to protest more after the murder of George Floyd.

⁵Since we can now include county fixed effects, we use the number of SSEs within the county as the push component, and control for the interaction between the push component and age as well as pull component and age. This also helps us to increase precision and achieve a stronger first stage.

⁶Using data from the PEW American Trends panel, we also find that survey respondents in counties with higher pandemic exposure consume more news about George Floyd on social media and hold more favorable views towards the BLM movement but not about the rights of undocumented migrants or about higher COVID-19 hospitalization rates among Blacks.

Our study contributes to several strands in the literature. To our knowledge, it is the first to establish a causal relationship between pandemic exposure and social media adoption. Existing research centers on supply-side constraints and initial staggered roll-outs (Manacorda and Tesei, 2020; Enikolopov et al., 2020; Müller and Schwarz, 2021; Melnikov, 2021). In contrast, we focus on isolating exogenous variation in late adoption, while controlling for the baseline Twitter network size and early adoption.

Second, we contribute to the literature on the political effects of social media (see Aridor et al. (2024) for a succinct overview of the literature). There is a rich literature examining the effect of technology adoption on protest (Campante et al., 2018; Christensen and Garfias, 2018; Guriev et al., 2021; Enikolopov et al., 2020; Manacorda and Tesei, 2020; Boyer et al., 2020), xenophobia, polarization, political preferences, social capital, and network formation (Falck et al., 2014; Boxell et al., 2017; Enikolopov et al., 2018, 2020; Guriev et al., 2021; Melnikov, 2021; Fujiwara et al., 2023; Müller and Schwarz, 2023, 2021; Campante et al., 2022; Boken et al., 2023; Enikolopov et al., 2024). Complementing this literature, we focus on newcomers to social media - the users that would have been never-takers in the absence of the pandemic - which can impact collective action in different ways. With the exception of Bursztyn et al. (2023), the literature has not yet investigated the timing of social media adoption and size of the social media network. Our findings also align with Fujiwara et al. (2023), who report that Twitter reduced Republican vote shares in the 2016 and 2020 elections but not in earlier periods, potentially because late adoption shapes behaviors and preferences more effectively than early adoption.

More generally, our analysis adds to a large literature that analyzes the determinants of social movements and protests, ranging from macro level drivers, such as local institutions or socio-economic conditions (Lipsky, 1968; Eisinger, 1973; McCarthy and Zald, 1977; Besley and Persson, 2011; Dube and Vargas, 2013; Berman et al., 2017), to micro level drivers, including individual decision making processes (Ellis and Fender, 2011; Guriev and Treisman, 2020; Sangnier and Zylberberg, 2017; Chenoweth et al., 2022; Cantoni et al., 2022) and different aspects of individual and social psychology (Guriev and Treisman, 2020; Sangnier and Zylberberg, 2017; Passarelli and Tabellini, 2017; González and Prem, 2024; Cantoni et al., 2019; Bursztyn et al., 2021), as well as the diffusion of social movements and protest across space and through networks (Becker et al., 2020; Casanueva, 2021; García-Jimeno et al., 2022). More narrowly, we add to the nascent literature on the causes and consequences of the Black Lives Matter movement and prominent police killings (Mazumder, 2018; Dave et al., 2020; Garcia and Ortega, 2024; Chenoweth et al., 2022; Ba et al., 2023; Celislami et al., 2023; Gethin and Pons, 2024).

The remainder of the paper is organized as follows. In section 2, we provide a background on the BLM movement, present motivating evidence and describe our main data sources. In section 3, we detail our empirical strategy. Section 4 describes our main results. Then, in section 5, we shed light on the role of persuasion versus coordination. Section 6 concludes.

2 Background and Data

2.1 BLM history and motivating evidence

The Black Lives Matter (BLM) movement emerged on social media after the acquittal of George Zimmerman in the deadly shooting of a Black teenager named Trayvon Martin. The movement was founded by three Black activists, Alicia Garza, Patrisse Cullors, and Opal Tometi, in July of 2013 with the aim to end systemic racism, abolish white supremacy and state-sanctioned violence (Black Lives Matter, 2020), and more generally, to “fundamentally shape whites’ attitudes toward Blacks” (Mazumder, 2019). Over the following months, an ever-increasing but small number of activists coalesced under the hashtag #BlackLivesMatter on Twitter and Facebook.

In August of 2014, after a court decision not to indict the responsible police officer in the fatal shooting of Michael Brown in Ferguson, #BLM became one of the most widely used hashtags on Twitter. The hashtag was used 1.7 million times in the three weeks following the court decision, compared to 5000 tweets in all of 2013, confirming its status as a mainstream social media phenomenon (Freelon et al., 2016; Anderson and Hitlin, 2016; Keib et al., 2018). In contrast to early adopters, new users encountered a social media environment with a large and established BLM network.

After the murder of George Floyd on May 25th, 2020, the BLM movement experienced an unprecedented expansion both geographically and demographically. Protesters took to the streets when a video of the murder of George Floyd went viral on social media, showing how police officer Derek Chauvin suffocated George Floyd using a choke-hold. The video spurred unrest in Minneapolis but the protests quickly expanded to other parts of the United States, including communities that had never engaged in BLM protests before. The number of BLM protests quadrupled in May and June of 2020, compared to previous peaks in 2016 (see Figure A1). The spike in protest activity and its coverage was not only the largest at the time but also the most extensive in the entire history of United States political protest, according to some observers (NYT, 2020, WP, 2020). Over half of the counties that recorded a BLM protest in 2020 had never mounted a protest for BLM before.

The prominence of BLM on Twitter peaked after the murder of George Floyd, when #BlackLivesMatter became the most popular hashtag on Twitter, peaking at 8.8 million mentions per day and videos on Twitter about George Floyd were watched over 1.4 billion times in the two weeks after his death (PEW, 2020). Pro-BLM narratives attracted significantly more attention on social media than anti-BLM narratives (Dunivin et al., 2022). Table A1 reports a few examples of tweets from our sample, illustrating how social media, especially Twitter, spread awareness for the Black Lives Matter movement after George Floyd’s murder. The tweets reveal a diverse set of users, including suburban moms, educators, and self-ascribed "allies" from various backgrounds. Many users emphasized the role of social media in exposing injustices and mobilizing action. These examples also illustrate how the movement gained traction among users that consider themselves new to the movement. One users writes: "I made a decision when I came on twitter to keep it strictly for work. I have other social media for expressing personal and political views. However, given the events of the last week, I feel compelled to say something - so here is my bit #BlackLivesMatter #WhitePrivilege."

Overall, the context of BLM during the pandemic offers two important features that allow us to investigate the role of late adoption in collective action. First, the BLM movement was conceived on Twitter in 2013, quickly becoming one of the most popular hashtags on the platform (McKersie, 2021). At the same time, BLM experienced an unexpected and viral protest trigger in the midst of the pandemic. Videos about the murder of George Floyd at the hands of police officer Derek Chauvin on May 25th 2020 were watched over 1.4 billion times on Twitter and the hashtag #BlackLivesMatter surged to 8.8 million mentions per day (PEW, 2020). Consequently, new users that may have joined the platform for other reasons were inevitably exposed to this viral protest trigger. Second, the pandemic was associated with record growth of new users and online activity on social media in a context where these platforms had been established and well known for over 15 years. In the first months of the pandemic, Twitter reported a surge in daily active users by 24% which “was driven by an increased engagement due to the COVID-19 pandemic” (Twitter, 2020). We think of these new users as never-takers in the absence of pandemic exposure.

2.2 Main data sources

In this section, we present the main data sources. We give a more detailed description, including exact definition, geographic unit, time frame and sources of all variables in Appendix A.

Black Lives Matter. This data comes from the crowd-sourced platform Elephrame. It provides information on the place and date of each BLM protest and estimated number of participants, as well as a link to a news article covering the protest. We extracted and geo-located all protests from August 2014 to September 2020. These protests are decidedly pro BLM. We also collected and geo-located information on street art with George Floyd-related content from the Urban Art Mapping George Floyd and Anti-Racist Street Art database. We add information on BLM and other protests from the US Crisis Monitor, a joint project between ACLED and the Bridging Divides Initiative (BDI) at Princeton University that collects real-time data on different types of political violence and protests in the US from 2020 onward.

Twitter. We collect three types of Twitter data at different points in time (before the pandemic, during the pandemic but before the murder of Floyd and in the three weeks after the murder of Floyd). We collect a total of more than 100 million tweets. First, from the Twitter API we collect a random sample of tweets, using the 100 most common words in English. Second, we collect the universe of tweets with BLM related hashtags. This includes the hashtags #BlackLivesMatter, #BlackLifeMatters, #BLM, and, separately, the #AllLivesMatter, and #BlueLivesMatter hashtags. Third, we scrape information on all followers of the official Black Lives Matter Twitter account. With the help of a geo-location algorithm, we can assign about 5 to 20% of Twitter users (depending on the sample) to counties.⁷ Importantly, we can identify the creation date of the Twitter profile. This allows us to assign "old users" (those that created their profile before January 20th 2020) and "new users" (those that created their profile between January 21st and May 25th 2020) to counties. Overall, this data allows us to measure

⁷Less than 2% of counties have no geo-located tweet.

i) online protest for and against BLM *ii*) late Twitter adoption and *iii*) information on baseline Twitter penetration.

Super spreading events (SSE). Information on SSEs are collected by a set of independent investigators and researchers from London School of Hygiene and Tropical Medicine (Leclerc et al., 2020). These are retrieved from scientific journals and news reports on SSEs, which are defined as "clusters" or "outbreaks" of COVID-19 infections with a minimum of 2 infections outside of the home.⁸ On average, each SSE is associated with 130 cases. For the whole period (January to August 2020), we identify a total of 1023 SSEs in the USA. Most commonly, events occur in nursing homes, prisons, factories, and retribution (correction facility) or medical centers. Table A4 provides descriptive statistics about each type of event. We describe the nature of these events in more detail and lay out the limitations of the SSE data set and how we address those in Appendix A.

Additional data sources. We complement these three main sources with an array of additional data sets described in more detail in Appendix A. The set of control variables comes from the American Community Survey as well as the US Census Bureau. In addition, we use data on lockdown stringency, Google searches, and mobile phone data on individual level mobility. We also exploit individual-level survey data from the Cooperative Election Study and the American Trends Panel, and election results for the 2012 and 2016 presidential elections from the MIT Election Data and Science Lab (2018).

2.3 Descriptive statistics

We report detailed summary statistics in Table A3, focusing on the sample of counties with no prior BLM protest. We use information that is available at different points in time: *i*) three weeks after George Floyd’s murder, *ii*) the day of the murder, *iii*) before the murder but after the pandemic started in January 2020, *iv*) later outcomes and *v*) baseline county characteristics before the outbreak of the pandemic.

The average likelihood of observing a BLM-related protest at the county level between May 25th and June 14th lies at about 5%. There are on average 0.06 events per county in the three weeks following George Floyd’s murder and the average number of participants is approximately 21 with a maximum of 5.5K participants. Conditional on recording a BLM protest, the average number of participants is about 440. Data from the ACLED crisis monitor reveals that, from May 25 through the end of 2020, counties recorded 1.3 pro-BLM protests on average and 3.1 protests overall. Compared to traditional protesters (not reported), counties that protest for the first time in 2020 have a lower Black population share (9% vs. 16%), are more rural (11% vs. 40% in large cities and suburbs) and have a higher vote share for Republicans in the 2016 presidential election (66% vs. 45%).

From the random sample of tweets that we collected, we can identify approximately 31 Twitter users on average before January 2020. However, this masks substantial heterogeneity in Twitter penetration across counties: some counties record over 8,200 users at baseline. A similar

⁸In our data set only 3 SSE are associated to less than 5 infections.

pattern emerges for late adopters: the number of new Twitter users in our data set ranges from zero to almost 250 of profiles created between January and May of 2020. In the three weeks following George Floyd’s murder we can identify 75 tweets per county using BLM-related hashtags. Tweets in opposition to BLM (those with #AllLivesMatter and #BlueLivesMatter) add to an average of 54 tweets per county over the same period.

3 Empirical Strategy

3.1 Estimating equation

We start by investigating the correlation between BLM protest and new Twitter users, i.e. those geo-located Twitter profiles at the county level that have been created after the outbreak of the pandemic and before the murder of George Floyd. We estimate a linear probability model of the following form:

$$BLM_c = \beta_1 \log(1 + \text{New Twitter Users})_c + \mathbf{X}_c \beta_{\mathbf{X}} + \gamma_s + \epsilon_c \quad (1)$$

Our outcomes of interest BLM_c is a dummy variable for any BLM protest in the three weeks following the murder of George Floyd. We also examine the intensive margin of offline protest and BLM protest online, specifically: *i*) the number of protests, *ii*) the log-transformed total number of participants, *iii*) the number of geo-localized tweets by new users containing BLM related hashtags or keywords and *iv*) the number of geo-localized followers of the official BLM Twitter account among new users.

All specifications include state fixed effects ζ_s which capture unobserved characteristics at the state level that could be related to both the prevalence of late adopters and our outcomes of interest. These include, for instance, lockdown stringency, differences in state laws, or policing strategies that are mandated at the state level. We include a vector of county level controls \mathbf{X}_c that account for within-state heterogeneity.

Most importantly, we control for the Twitter network at baseline (the number of Twitter profiles created before January 2020) and cumulative number of COVID-19 related deaths (just before the protest trigger). This captures the mobilization potential of existing users as well as any pandemic-related drivers of protest (such as overall agitation, opportunity cost of protesting, or salience). In addition, we include variables that are associated with participation in the BLM movement, such as a dummy for urban counties, Black population share and the poverty rate among Blacks, as well as the use of deadly force by police (i.e. number of Black people who died during an encounter with the police, from 2014 to 2019 and in 2020 up to May 25th). We also control for underlying political and attitudinal factors and socioeconomic drivers of protest, such as the vote share for Republicans in the 2012 and 2016 presidential elections, median household income, unemployment rate, community resilience (an indicator developed by the United States Census Bureau that captures the capacity of county to absorb the health impacts of pandemics), and a proxy for social capital. We cluster standard errors at the state level but show in Appendix B that our results are robust to spatial clustering.

3.2 Instrumental variable approach

3.2.1 First-stage estimating equation

The concern with a causal interpretation of β_1 is that there may be unobserved county characteristics that drive both late social media adoption and BLM protest. We mitigate concerns about reverse causality by limiting social media adoption to the pre-George Floyd period. We tackle concerns about the endogeneity of late Twitter adoption with an instrumental variable strategy, akin to push-pull instruments (or variants of shift-share or Bartik instruments) used in the labor and migration literature (see, for instance, Burchardi et al., 2019).⁹ The first stage writes as follows:

$$\begin{aligned} \log(1 + \text{New Twitter Users})_c = & \delta_1 \text{SSE}_{-c} \times \text{SXSW}_c \\ & + \delta_2 \text{SXSW}_c + \delta_3 \text{SSE}_{-c} \\ & + X_c \delta_X + \zeta_s + \varepsilon_c \end{aligned} \quad (2)$$

We exploit plausibly exogenous variation in the baseline Twitter network (SXSW_c) as a pull factor and pandemic exposure (SSE_{-c}) as a push factor to explain late Twitter adoption. The log-transformed number of new Twitter users is predicted using the interaction term $\text{SSE}_{-c} \times \text{SXSW}_c$ (excluded instrument), while controlling separately for SSE_{-c} and SXSW_c . The specification includes state fixed effects and the same set of county-level controls. Standard errors are clustered at the state level.

We argue that the pandemic acted as a large enough push factor that led to the adoption of social media in an otherwise saturated social media market. To proxy local pandemic exposure, we leverage plausibly exogenous variation in the occurrence of so-called super spreading events (SSE) in the early stages of the pandemic (we provide more detail on the rationale behind this below). SSE_{-c} is measured as the number of SSEs that occur within 200 km of the county border but not within the county until 6 weeks before the murder of George Floyd.

We combine this with a pull factor for social media adoption, drawing from the literature on path dependence in technology adoption that suggests increasing returns to joining a social media platform when the existing network is large (Arthur, 1989; Bursztyn et al., 2023). Therefore, we interact SSEs with exogenous variation in baseline Twitter networks borrowing from Müller and Schwarz (2023). The authors leverage an interactive media, film, and music festival and conference called South by South-West (SXSW) held annually in Austin, Texas. The 2007 edition heavily promoted Twitter and contributed to the diffusion of Twitter adoption: 60 percent of early Twitter adopters were connected to SXSW, and the platform’s growth accelerated disproportionately in counties with more SXSW attendees - proxied by followers of the official SXSW account who joined Twitter during the 2007 festival (Appendix A.2 provides more de-

⁹These shift-share/push-pull/Bartik-type instruments use the inflow of immigrants from a specific origin to all other counties in the US (but not the own county) or economic shocks at origin as a shifter (or push factor) and the existing migrant network as share (or pull factor) to predict county-level migration inflows. In our context, quasi random variation in both push and pull factors tackles many concerns associated these type of instruments, such as the required exogeneity of the pull component since these instruments assume a pooled exposure research design or the orthogonality of the push or shifter component of the instrument (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022)

tail). We follow this approach and measure $SXSW_c$ as the log-transformed number of SXSW followers who created their account in March 2007 in the county and neighboring counties.¹⁰

The instrument predicts a larger number of late adopters in counties with both higher pandemic exposure and a wider social media network at baseline. The presence of nearby SSEs raises local awareness and increases time spent at home, thereby elevating demand for social media. Simultaneously, the strength of the pre-existing Twitter network, rooted in initial adoption shocks like SXSW, creates the conditions necessary for this demand to translate into actual platform usage. Any direct effect of SSEs on collective action (such as fear of infection or salience of the pandemic) and any direct effect of the historical SXSW shock (such as the presence of tech-savvy early adopters) will be captured in the set of controls.

For illustration, consider two counties with contrasting instrument values: Union County, New Jersey (top decile) and Salem County, New Jersey (bottom decile). Both counties were exposed to numerous nearby super-spreading events, but only Union County has a high number of SXSW-related Twitter followers, indicating a stronger pre-existing social media network due to a higher number of attendees at the festival in 2007. Consequently, the instrument predicts higher Twitter adoption during the pandemic in Union County. In a different scenario, Suffolk County, New York (top decile) and Jefferson County, New York (bottom decile) have comparable SXSW follower counts but differ significantly in SSE exposure. Suffolk’s higher number of nearby SSEs suggests a greater predicted rise in late Twitter adoption.

3.2.2 More details on super spreading events

Super-spreaders, individuals significantly more contagious than average, are critical in the spread of infectious diseases, including COVID-19, where 70-80% of transmissions stem from just 10-20% of cases (Adam et al., 2020; Edwards et al., 2021; Endo et al., 2020; Miller et al., 2020). These super-spreading events (SSEs) involve large numbers of infections at a single gathering. Our dataset tracks outbreaks and clusters defined as "two or more test-confirmed cases of COVID-19 in a non-residential setting within 14 days" (more details in Appendix A).

On average, each SSE is associated with 130 cases, a significant number considering the average county had 164 detected cases by the end of May 2020. Most of the approximately 1,000 SSEs in our data are from the medical care sector (see Table A4). Our baseline specification controls for testing capacities with state fixed effects, and for health care facility presence using the community resilience index. In robustness checks, we also verify that results hold when excluding certain event locations (i.e. in prisons).

We focus on SSEs in the early stage of the pandemic (from January 20th to April 13, 2020) because limited public awareness about the spread of the virus and lax lockdown measures lead to significant transmission through group gatherings. During this period, we argue SSEs were primarily driven by the presence of highly infectious individuals rather than factors like risk

¹⁰Following Müller and Schwarz (2023), the specification extends the set of baseline controls X_c to include *Pre SXSW Users_c*, defined as the log of one plus the number of SXSW followers in the county and neighboring counties who created their accounts *before* March 2007 as well as its interaction with SSE_{-c} . This captures existing interest in Twitter and the SXSW festival before the 2007 edition and captures any potential compounding effect with pandemic exposure later.

preferences.¹¹ Figure A2 shows that most SSEs (solid blue line) occurred between mid-March and late April, coinciding with a rise in new COVID-19 cases (red dotted line). SSEs declined as state-imposed lockdown measures increased (green dashed line, Oxford COVID-19 Government Response Tracker).

We also leverage the geographic variation in SSE exposure, using SSEs within a 200 km (120 mile) radius of county borders, while excluding SSEs within the county itself. This approach has two advantages: unobserved factors predicting outcomes would need to predict SSEs in nearby counties, and it allows us to capture more variation in pandemic exposure, which doesn't stop at county lines. Figure 1 plots the geographic distribution of BLM protest, late adopters, SSEs and the instrument variation across U.S. counties, showing that SSEs were not spatially concentrated but spread throughout the country (see Figure 1b).

3.2.3 Relevance and first stage results

We examine the relevance of our instrument, focusing first on the role of SSEs. SSEs primarily contribute to the spread of the virus. Columns 1 and 2 of Table A5 show that SSEs significantly increase COVID-19 related deaths and cases.¹² In addition, and in line with Campante et al. (2024), SSEs may also increase the salience of the pandemic beyond its direct impact on mortality and infection rates. For instance, individuals may spend more time at home and substitute real world interactions with virtual ones, sign up to learn about the pandemic in real time, or seek to express themselves and find connection in times of crises. Using Google mobility data from 45 million mobile phones, we show in column 3 that the average time spent at home increases in response to SSEs in close proximity.

Next, we turn to the role of the existing network. Using the log-transformed number of geo-localized Twitter accounts from a random sample of Tweets in December of 2019, we check whether the combination of pandemic exposure and baseline network matters for late adoption. Columns 4 and 5 confirm that SSEs increase Twitter take-up but that this is mainly driven by counties with a larger number of Twitter users at baseline.

Lastly, we validate the pull component of the instrument, namely that early adopters at the 2007 SXSW festival still determine the number of Twitter users in December of 2019. Column 6 shows that attendees of the SXSW festival, proxied by geo-localized SXSW followers that created their profile during the festival, predicts the number of users in December of 2019. Combined, these results lend confidence to the combination of push and pull factors as key drivers of late Twitter adoption during the pandemic.

We present the corresponding first stage results in Table A6. Column 1 reports the coefficient of our baseline instrument, using SSEs within a 200 km radius and until six weeks before the protest trigger interacted with SXSW followers. The instrument is well above the conventional threshold and strongly predicts new Twitter users. Reassuringly, SSEs only predict new Twitter users through their interaction with the baseline network but not directly. Specifically, a one standard deviation increase in $SSE_{-c} \times SXSW_c$, increases Twitter take-up by 6.5%. Columns

¹¹Importantly, the characteristics that make an individual a super-spreader and make them exhale more viral particles are determined in a large part by their biological characteristics rather than by their behavior (Edwards et al., 2021)

¹²Note that for this set of exercises we exclude COVID-19 related deaths and the Twitter network at baseline from the set of controls.

2 to 6 of the same table present the first stage results for varying distances to SSEs (100km, 150km, 300km, and 400km) and time lags (4 weeks, 5 weeks, 7 weeks and 8 weeks). The results remain largely robust to those changes, albeit – expectedly – becoming weaker when we include SSEs that are more distant.

3.2.4 Plausibility of quasi random exposure

The exclusion restriction of the instrument requires that the interaction between SSEs and SXSX only impacts BLM protest through the number of new Twitter users, given the set of controls and fixed effects in our model. It is possible that the interaction between pandemic exposure and baseline Twitter use affects BLM protest beyond its impact on new Twitter users. We examine this possibility in Figure 2 (and repeat this exercise for each instrument component in Figure A3), where we predict our instrument $SSE_c \times SXSX_c$ with an array of county and state characteristics. We include the set of controls from our baseline specification and show results including and excluding state fixed effects in order to investigate potentially important correlates which are only available at the state level, such as lockdown stringency and the disproportionate death burden of COVID-19 on the Black population. Coefficients are standardized for comparability.

First, we probe the exogeneity of instrument with respect to BLM protest. We look at the geographic distance to Minneapolis, the location of George Floyd’s murder and the origin of BLM protest in 2020. If the location of SSEs and SXSX users are spatially correlated with distance to Minneapolis, we may capture geographic diffusion patterns. In addition, we look at various measures for exposure to previous BLM protest and BLM related topics, using a dummy variable for whether a county’s neighbor had a BLM protest between 2014 and 2019 or whether a county’s neighbor protested for BLM in 2020, as well as Google searches for the term BLM and Black Lives Matter in the weeks leading up to the murder of George Floyd. All coefficients are small in magnitude and insignificant with the exception of the small and marginally significant coefficient for neighboring BLM protest between 2014 and 2019, suggesting that – if anything – exposure to the push-pull instrument is driven by locations further away from centers of protest. Nonetheless, in a robustness check, we include this additional control and the results hold.

In addition, we address the concern that the instrument may predict the severity of lockdown stringency and therefore the opportunity cost of protesting. This could be because SSEs may get more attention in places with a larger social media network. While lockdown stringency would be captured in the state fixed effects, restrictions on mobility may be relevant across state lines. Moreover, if this interaction disproportionately affects the Black community, then we may capture an unobserved factor that either riles up or incapacitates prospective protesters. Reassuringly, in panel a, we do not find that lockdown stringency or the COVID19-related death burden on Blacks is associated with the instrument.

We continue this exercise with a battery of potentially relevant county characteristics, such as economic inequality and poverty all measured at baseline in 2018 from the American Community Survey (Gini, poverty rate, white and female poverty rate), education (Bachelor degree, some college education share), demographics (female above 15, population shares of age groups), economic variables (average rent, urbanization and job density) as well as social capital variables

(various religious and sports organizations etc.). We find no systematic pattern between these variables and the instrument.

In a last step, we tackle the concern that late adoption influences the detection and reporting of SSEs. Figure 3 verifies that Twitter take-up follows rather than precedes or coincides with SSEs. In an event-study design at the county and day level (with county and state-day fixed effects and standard errors clustered at the county level), we show that the timing of SSEs within a 200 km radius is associated with a subsequent increase in the number of new Twitter profiles.

3.3 Event study design

We complement the instrumental variable strategy with an event study design that allows us to account for time invariant unobserved heterogeneity at the county level. More specifically, we exploit the timing of the unexpected protest trigger - the murder of George Floyd on May 25th 2020 – and compare differences in BLM protest intensity across counties with varying levels of late Twitter adoption. We collapse the data to the county-month level¹³ and expand the time frame to 16 months before and 4 months after the protest trigger and also expand the sample to all counties.¹⁴ Our event-study specification writes as follows:

$$BLM_{ct} = \zeta_c + \delta_{st} + \sum_{\substack{k=T_0 \\ k \neq -1}}^{T_1} \beta_k \log(1 + \text{New Twitter Users})_c \times \mathbb{1}_{t=\text{May } 25\text{th}+k} + \epsilon_{it} \quad (3)$$

We include county and state-month fixed effects, which account for any county characteristic that could influence the overall propensity of a county to protest for a BLM related cause as well as any state-specific time-varying shocks that could hinder BLM protest. County fixed effects will capture any time-invariant county characteristics, including their past exposure to police violence, underlying racial demographics, or socio-economic conditions that may affect the likelihood of protests. The inclusion of state-month fixed effects allows us to control for broader state-level policies or events, such as COVID-19 restrictions or economic relief packages, that could vary over time and influence protest behavior differentially across states. The outcome *BLM* is measured as a dummy variable for any BLM protest in county *c* in month *t*. We cluster standard errors at the county level.

Our key variable of interest is the interaction term between log-transformed Twitter adoption as defined in the previous section and event time dummies. The coefficients β_k trace out the dynamic effect of Twitter adoption on BLM protest, allowing us to examine whether there are any significant pre-existing trends in protest intensity that relate to new Twitter users.¹⁵ This allows us to examine whether late adoption is associated with changes in a county’s propensity to mobilize for BLM over time.

¹³ $t = 0$ is the 25th of May; equivalently, all preceding and following months begin on the 25th.

¹⁴Note that if we only used counties with no prior BLM protest, we would achieve parallel trends mechanically.

¹⁵In a robustness check, we run a fully saturated model and interact the post dummy with the full set of county-level controls and the results hold.

4 Main Results

4.1 Late adoption and collective action

Table 1 present our main results, reporting 2SLS results in panel A and OLS results in panel B, as well as reduced form estimates and first stage results in the bottom two panels. We focus on the sample of counties with no previous BLM protest and run a linear probability model for the presence of any BLM protest in the three weeks following the murder of George Floyd. Column 1 includes state fixed effects and controls for baseline Twitter users (logged number of geo-localized Twitter accounts created before January 2020 from a random sample of Tweets) as well as pandemic exposure (cumulative COVID-19 related deaths until May 25th).¹⁶

Columns 2 to 4 successively add more controls. In column 2, we add economic county level controls, which include the average unemployment rate between May 2019 and April 2020, the median household income in 2016 from the Census Bureau Atlas, a dummy for counties that have a large urban center, and the community resilience index from the US census bureau, which captures the vulnerability of a county to the pandemic, including the share of population with health risks as well as hospital capacity and availability of medical staff. Column 3 accounts for county-level characteristics that are related to the propensity to protest for a BLM-related cause and past potential protest triggers, including the Black population share and Black poverty rate – both from the American Community Survey 2018 – as well as the use of deadly force by police against Blacks between 2014 and 2019, from *Fatal Encounters*. Column 4 also accounts for within-state political heterogeneity, including the vote share for the Republican party in the presidential elections of 2012 and 2016 and a proxy for social capital taken from Rupasingha et al. (2006), measured as the number of civic organizations per capita.

Throughout, we find a positive and significant effect of late adoption on the likelihood of observing a BLM protest. The coefficient from our preferred estimation in column 4 suggests, that a 10 percent increase in the number of late adopters increases the likelihood of a protest by 4.4 percentage points. This is a large effect considering that among the counties that never protested for a BLM related cause only 5% begin to protest in the three weeks after the murder of George Floyd.

Across all estimations, 2SLS consistently produces larger coefficients than OLS. Without exogenous variation in late Twitter adoption, counties with more new Twitter accounts may be less likely to protest, possibly due to selection bias. For example, individuals who signed up for Twitter absent the shock may already be aware and supportive of the BLM movement. Signing up for the platform may thus not change their information set and propensity to protest. The difference between OLS and IV estimates can also arise from the local average treatment effect. We estimate the effect of late adoption on BLM protests for compliers — counties where new Twitter users increased due to the push-pull shock. These new users might be more motivated to seek information and community online, making them more likely to participate in protests in response to triggering events.

¹⁶As mentioned in Section 3.2, the 2SLS, first stage and reduced form estimation additionally control for SSEs and SXSW followers separately, as well as the log-transformed number of SXSW followers that created their account before the Austin festival, as well as its interaction with SSEs, such that we capture historic interest in Twitter which may interact with SSEs and drives BLM protest. We report the coefficients for all variables in Table A7.

4.2 Offline and online BLM protest

Table 2 examines additional outcomes, including the intensive margin of offline protest and online protest. Again, we present 2SLS coefficients for log transformed new Twitter users. Column 1 reports our baseline result for reference. Column 2 examines the number of BLM protests, column 3 the aggregate number of participants across all protests in the three weeks after the murder of George Floyd's. Our results suggest that late adopters increased the frequency and scope of BLM protest in counties that have never protested before.

Next, we investigate the impact of late adoption on online protest. As suggested above, online protest requires less coordination and serves as a form of political expression that is low cost. Through this exercise, we are also able to look at the protest behavior of late adopters themselves since offline protest may include both late and early adopters as well as those that are not on the platform at all. Column 4 uses the log transformed number of geo-localized Tweets posted by late adopters in the three weeks after the murder of George Floyd that contain BLM-related hashtags and keywords. We proxy more explicit support for BLM with the log transformed number of geo-localized late adopters who follow the official BLM Twitter account in column 5. We find that a one percent increase in new Twitter users increases the number of BLM related tweets by 0.8% and the number of new followers to the official BLM Twitter account by 0.4%. Our outcome of interest in column 6 is the first principal component of the aforementioned proxies for online and offline BLM protest.¹⁷

Overall, we find that Twitter adoption increases BLM protest both online and offline. This result also stands in contrast to a strand of the social science literature that has emphasized "slacktivism" i.e. the substitution away from more effective forms of collective action towards more ineffective and symbolic activism online (Christensen, 2011; Schumann and Klein, 2015). It also suggests that late adoption can amplify mobilization for causes with a strong and established network.

4.3 Summary of robustness checks

We run a series of empirical exercises in Tables B1 to B6 which we describe briefly here and in more detail in Appendix B. We always present baseline results in column 1, which replicates the specification used in column 4 of Table 1.

We begin by showing the second stage results for variants of the instrument. We show in Table B1 that our results are not sensitive to narrowing or expanding the time window (between 4 weeks and 8 weeks before the murder of George Floyd) or geographic distance (between 100 and 400 km) of SSEs. Next, in Table B2, we change the definition of pandemic exposure as measured by SSEs. In column 2, we exclude SSEs in prisons as they may impact the public perception of exposure to the pandemic differently and may also be related to factors that drive BLM protests. In column 3, we also control for the number of SSEs in the county itself to account for correlation between neighboring and local SSEs. In column 4, we weight SSEs by their geographic distance to the county border. Throughout, our results remain robust to these

¹⁷We detail the construction of the first principal component in Table A8. We show the correlation between all measures, the eigenvalue of the first principal component and the factors loadings for each variable.

changes.

In a second set of exercises in Table B3, we examine the robustness of our results with respect to additional controls. As shown in Figure 2, the instrument is correlated (albeit the coefficient is small in magnitude and marginally significant) with the likelihood of having traditional protesters as neighboring counties. In column 2, we add this as an additional control and our results hold. Next, we examine the possibility that protest is driven by differences in the COVID-19 trajectories across counties. While our baseline specification controls for the cumulative number of COVID-19 related deaths, some counties may be more likely to protest when the salience of pandemic exposure was higher closer to the protest trigger. In columns 3 and 4, we add the number of COVID-19 related cases and deaths in the week leading up to the murder of George Floyd and the results hold.

Next, in columns 5 to 8, we define narrower control groups, comparing counties with similar ex-ante BLM protest probabilities. Using LASSO, we select the subset of relevant county characteristics that determined BLM protest in the past and create a propensity score for protesting, based on the selection of these variables (we describe this in more detail in Appendix B). In column 5, we include this variable as an additional control and confirm that our results remain the same. In columns 6 to 8, we include fixed effects for various cut-offs for the protest propensity score. We split the fixed effects along the thresholds that produce county groups of different sizes: groups of 1000 units (3 groups), 100 units (28 groups) and 10 units (277 groups). This allows us to compare counties within narrow bands of ex-ante protest probabilities. All results are similar in magnitude and remain precisely estimated throughout (except - as expected - when we include 277 fixed effects for a priori protest probabilities).

Table B4 verifies that our results are not sensitive to sample composition and spatial clustering. In columns 2 and 3, we drop coastal counties and states, assuaging concerns that these protest surges are driven by Democratic leaning states like California or New York. In columns 5 to 8, we account for spatial correlation using Conley (1999) standard error adjustments with different distance cut-offs between 100 and 300 km, as well as between neighboring counties irrespective of the distance. If anything, our results become more precisely estimated.

In Tables B5 and B6, we change the definition of our main treatment and outcome variables. Table B5 addresses concerns related to the log transformation of our treatment variable and linear probability model. In column 2, we measure the number of new users as the share of all users at baseline. In column 3, we use the inverse hyperbolic sine transformation, instead of the logged number of new users plus one. In column 4, we use the absolute number of new users. In column 5, we estimate a probit model. Table B6 shows that our results are also not sensitive to changing the definition of the outcome variable by narrowing or expanding the time window for BLM protest between 2 and 6 weeks after the murder of George Floyd. Importantly, in a placebo exercise, we also show that new Twitter users do not predict past BLM protest when including the sample of counties with prior BLM protest.

4.4 Panel evidence for all counties

Figure 4 displays the event study estimates, tracing the effect of late Twitter adoption on the probability of a BLM protest from 16 months before the trigger to 4 months after, expanding

the sample to the full set of counties. The results confirm that counties with higher levels of late adoption did not exhibit significant differences in protest activity prior to George Floyd’s murder. The coefficients for the pre-treatment period are small, close to zero, and statistically insignificant, indicating no pre-existing trends in BLM protests related to Twitter adoption. Following the murder of George Floyd, we observe a sharp and significant increase in BLM protest activity in counties with higher levels of new Twitter accounts. The largest effect occurs in the immediate aftermath (between May 25th and June 24th of 2020), where the interaction between Twitter adoption and the post dummy shows a substantial spike, suggesting that Twitter played a key role in facilitating protest. The effect persists in the months following the event, though it tapers off after the initial surge, with significant but smaller coefficients through the next four months.

We provide an array of additional exercises that confirm the robustness of the event-study results in Figure B1. We begin by taking the number of protest instead of the likelihood of protest as an outcome in Figure B1a. This assuages concerns that the intensity of protest (rather than its occurrence) differs in the pre-period across counties with different levels of late adoption. In addition, in Figure B1b, we restrict the sample to all counties that observed any protest after the murder of George Floyd. This allows us to examine whether the timing of protest differs across counties with higher levels of late adoption. As expected, we find that counties with higher levels of late adoption respond earlier to the protest trigger. Next, in Figure B1c, we present the fully saturated event study, interacting the post dummy with the set of controls used in our preferred specification. This accounts for a county’s differential response to the 2020 protest trigger depending on, for instance, the demographic composition of the county, assuaging concerns that late adoption is correlated with country characteristics that drive protest in 2020 but not before. Lastly, in Figure B1d, we expand the time frame even further, to account for differences in the propensity to protest for BLM in the 2 years before the murder of George Floyd. All of these exercises confirm our baseline result.

5 Mechanisms

Protest mobilization can be driven by two mutually non-exclusive forces: by a reduction in coordination costs and through persuasion. Leveraging early adoption, the literature has emphasized the importance of coordination costs (Manacorda and Tesei, 2020; Enikolopov et al., 2020). However, selection into early adoption may favor coordination costs over persuasion if these users are already informed, politicized and more likely to protest. In addition and in contrast to early adopters, late adopters experience social media as a politically consolidated environment upon entry - rather than a blank slate. In the following section, we examine whether persuasion or a reduction in coordination costs can explain the increase of BLM protests in response to late adoption. To this end, we examine social movements and protests that appeal to a different political base than BLM supporters and we leverage individual level survey data to examine whether late adopters changed their preferences.

5.1 Protest for other causes

Table 3 explores the role of late adoption on different forms of political expression and protest for different causes. In column 1, we examine the impact of new Twitter users on George Floyd street art and graffiti which we scrape and geo-localize from the *Urban Art Mapping George Floyd and Anti-Racist Street Art*. Graffiti has emerged as a prominent vehicle for advocacy within anti-racist movements (Mathieu, 2018; Cappelli et al., 2020). We find that the surge in Twitter use during the pandemic does not correspond with an increase in George Floyd-related street art, potentially because such these are challenging to replicate in places that have only recently engaged with BLM.

Next, in columns 2 to 5, we leverage information about the number of protests for BLM and other causes in 2020 from the ACLED US Crisis Monitor (we describe this data in more detail in Appendix A). We expand the observation period from May 25th 2020 until December of 2020 since BLM protest have crowded out other protests in May and June of 2020. In column 2, we confirm our previous findings that new Twitter users increase the number of BLM protests throughout 2020. In column 3, we look at all other protests excluding BLM and find that late adoption increases protest for other causes.

We examine specific protest causes, focusing on the two largest protest movements in 2020 beyond BLM: anti-masking protest and pro-Trump protest. First, we turn to protests that oppose social distancing rules and mask mandates. Public narratives about the political convictions of these protesters suggest that they are predominantly driven by conservative or libertarian ideologies and are distinct from the supporters of movements like BLM. However, Chenoweth et al. (2022) show that there is a large overlap between supporters of BLM and those calling for fewer public health restrictions, using individual-level survey data. In line with their findings, we detect an increase in mobilization against public health restrictions similar in magnitude to BLM protest.

In order to capture protests that are ideologically opposed to BLM, we consider pro-Trump protests as well as explicit counter-mobilization to BLM. Based on the ACLED protest descriptions, we classify protests as pro-Trump if they involve actors such as the "Proud Boys" or QAnon, protest perceived election fraud or explicitly support Trump.¹⁸ The coefficient is close to zero and imprecisely estimated, suggesting that new users do not mobilize for conspiratorial, populist or explicitly pro-Trump causes. There are many possible explanations for this finding. First, the surge in conspiratorial and "stop the steal" protests emerged in late 2020, after the presidential election and peaking in early January 2021, such that many of these are not included in the data set. Second, it is possible that new Twitter users are energized to join protests when they enter the platform but that this mobilization effect subsides over time. Third, Twitter may be particularly suited for mobilizing protest for certain causes over others, potentially because of the political make-up of its existing users.

To investigate these possibilities, we examine counter-mobilization against BLM on Twitter, focusing on the hashtags #AllLivesMatter and #BlueLivesMatter, which oppose the BLM

¹⁸The keywords we include are: Stop the Steal, Proud Boys, Qanon, election fraud, ballot, pro-Trump, Trump support, in support of Trump, WAF (Women for America First) and MAGA. The results are similar when we only include pro-Trump protest and if we exclude counter demonstrations to BLM protest (tables not reported).

movement. If salience drives the absence of right-wing populist protests, rather than the political skew of the platform when late adopters enter, we would expect to see increased online opposition to BLM. These hashtags relativize the BLM message and highlight concerns about the safety of white individuals and police officers. We scrape tweets containing these hashtags and allocate them to counties based on the locations of the users. However, and in line with the results in column 5, we do not find that new Twitter users also mobilize for causes opposing BLM online. The coefficients in columns 6 and 7 are small in magnitude (compared to pro-BLM tweets) and not significant.

Together, these findings highlight that late adoption does not universally increase protest but that it is movement-specific. It favors causes with established organizational infrastructure and broad public appeal while being less effective for more extremist or niche causes that lack widespread resonance on the platform (e.g., anti-BLM tweets, or conspiratorial and pro-Trump protests).

5.2 Attitudes on racial equity, police violence and voting

In our context, newcomers to Twitter encountered a platform with a large and established BLM network. As described in section 2, BLM was one of the most prominent hashtags on the platform - even before 2020 (Keib et al., 2018). Moreover, Twitter users were substantially more left-leaning than the overall population conditional on age.¹⁹ In fact, the majority of Tweets is produced by a minority of users, who tend to be even more left leaning than their inactive counterparts (Pew, 2019). Combined with the large and viral protest trigger, which unleashed a wave of imagery of George Floyd’s murder, late adopters were exposed to a decidedly pro-BLM narrative.

We draw from the Cooperative Election Study (CES) to assess the effect of late adoption on attitudes at the individual-level. The CES was conducted in November of 2020 and contains information on respondents’ social media consumption, attitudes towards the police and racism, vote in the 2016 and 2020 election and – importantly – their county of residence. This allows us to connect county level shocks to individual-level take-up.

5.2.1 Estimating equation

We start by investigating the correlation between social media use and a variety of outcomes Y_i , including whether the respondent has attended a protest, march or demonstration in the past year, whether the respondent takes any form of political action on social media, whether the respondent voted for Biden in the 2020 presidential elections and attitudes towards the police and about racism (see Table A9 for the exact survey questions).

$$Y_i = \beta_1 \text{Social Media Use}_i + \mathbf{X}_i \beta_{\mathbf{X}} + \gamma_c + \epsilon_i \quad (4)$$

We include county fixed effects γ_c and an array of individual level controls \mathbf{X}_i (age, gender,

¹⁹Based on a nationally representative survey conducted in 2018 that sampled 2,791 U.S. adults who have a Twitter account and agreed to allow researchers to access that account, and found that these users have a higher preference for the Democratic party.

employment status, number of children, dummies for race and for vote in the 2016 presidential election) and cluster standard errors at the county level. Our treatment is Social Media Use_{*i*} measured as a dummy variable that indicates whether the respondent has used social media over the past 24 hours. We take this as a crude indicator for social media use, combining the extensive and intensive margin.

5.2.2 Instrument for social media take-up at the individual level

In order to leverage exogenous variation in individual-level take up, we follow the same intuition as in the previous section. Twitter adoption depends on the baseline network and local pandemic exposure with the caveat that social media adoption is limited in saturated markets and will expand only where there is scope for adoption. We proxy the scope for adoption at the individual level with the age of the respondents. We conjecture that older individuals are more sensitive to our instrument since they are less likely to have already used social media at baseline.²⁰ In other words, young individuals tend to be early adopters; older individuals tend to be late adopters, and they are the ones who will respond most to shocks to late adoption.

Therefore, we instrument social media use at the individual level with the interaction of the push-pull instrument $SSE_c \times SXS W_c$ with the age of the respondent. We also add the interaction of age with a set of county-level instrument-related controls \mathbf{N}_c . Analogous to section 3 these include: SSE and SXS W separately, the logged number of baseline Twitter users, cumulative COVID deaths per 1000 inhabitants on May 24th, the logged number of followers of SXS W that joined before March 2017, as well as its interaction with SSE. We also control for county fixed effects ζ_c and include the same set of individual level controls \mathbf{X}_i . The first stage writes as follows:

$$\begin{aligned} \text{Social Media Use}_i = & \delta_1 SSE_c \times SXS W_c \times Age_i \\ & + Age_i \times \mathbf{N}_c \delta_{\mathbf{N}} + \mathbf{X}_i \delta_{\mathbf{X}} + \zeta_c + \varepsilon_i \end{aligned} \quad (5)$$

Our estimation compares differences in political preferences between older and younger respondents in high versus low exposure counties, i.e. counties with a greater shock to social media adoption. Note that in this exercise, SSE_c is measured as the number of SSEs in the county of the respondent, rather than outside the county for two reasons: first, the county fixed effects allow us to examine *within* county differences across individuals that are more or less likely to take up Twitter during the pandemic. This also allows us to assuage concerns about any unobserved differences across counties that may drive the occurrence of SSEs and of SXS W attendance independently of age. Second, by targeting variation in late adoption within counties, we gain a more precise estimate of how social media use influences individual-level political outcomes.²¹

The identifying assumption requires that the triple interaction $SSE_c \times SXS W_c \times Age_i$ only impacts political preferences through the adoption of social media. In other words, within

²⁰About 90% of respondents under the age of 33 – the bottom age quartile – have used social media in the last 24 hours compared to about 60% in the top quartile above the age of 63 (see Figure A4).

²¹We show in Table A10 that SSEs outside of the county deliver weaker first stages.

the same county older respondents will only change their political preferences in response to the push-pull instrument through a higher likelihood of signing up to social media compared to younger respondents. Any age-differential responses to baseline social media penetration or pandemic exposure will be captured in the interactions $SXSW_c \times Age_i$ and $SSE_c \times Age_i$ respectively. Any unobserved characteristics at the county level that determine social media use and political preferences independently of age will be captured in the fixed effects. Any individual-level heterogeneity related to gender, household size, employment status, race and political preferences in 2016 will be captured in the set of controls.

The exclusion restriction is violated if, for instance, baseline social media penetration combined with exposure to SSEs induced the older generation to reduce mobility and isolate more than younger respondents, such that their overall media consumption has increased - including social media. This could be because older respondents felt particularly vulnerable to the pandemic, which may have been more salient in places with a larger share of early adopters. We assess this concern by estimating a placebo first stage, looking at the consumption of different media outlets. We exploit information in the CES on the use of TV and newspapers over the past 24 hours. We report in Table A11 that our instrument predicts social media use but it does not predict the consumption of other media outlets, mitigating the concern that overall media consumption increased differentially for older respondents. We also verify whether the triple interaction predicts the salience of the pandemic. We do not find evidence for higher infection rates among older respondents and their family members or that older respondents reduced work hours as a result of the pandemic. Lastly, we check whether the push-pull instrument disproportionately affects counties where older respondents are more progressive overall, leveraging information on attitudes that are unrelated to BLM, including views on abortion or environmental protection. Throughout we do not find evidence that our instrument is correlated with other attitudes.

5.2.3 Results

We present the results on social media and political preferences in Table 4. OLS, reduced form and first stage results are presented in Table A12. Throughout, the first stage F-statistics are above the conventional threshold. In column 1, we exploit information on whether the respondent has participated in any protest, march or demonstration within the last year. Our 2SLS estimates suggest that social media use increases the likelihood of joining a protest. Similarly, column 2 suggests that respondents are also more likely to be politically active on social media. This includes posting or forwarding a story, photo, video, link or comment on social media, as well as following a political event, reading a story or watching a video about politics.

Next, we examine the voting behavior and attitudes of respondents. Specifically, we show in column 3 that respondents are more likely to vote for Biden in the 2020 election when they increase their use of social media in response to the individual level social media adoption shock. Note that we control for the respondent's vote in 2016, thereby exploiting changes in political preferences across the two election cycles.

In column 4, we build an index that captures leniency towards police officers – a topic that has been prominently discussed within the BLM community, including a sub-strand of the movement

that demands to defund the police. The index is constructed as the first principal component of seven questions, ranging from the requirement of police officers to wear body cameras, a national registry for violent police officers, a ban on the use of choke holds and more. Higher values indicate higher leniency towards the police.²² We find a negative and precisely estimated coefficient, suggesting that social media take-up increases the demand for more scrutiny of the police.

In column 5, we repeat this exercise, this time using attitudes towards racism as an outcome. The index is comprised of four question with higher values of the index indicating less awareness of racial inequalities. The set of questions includes views on the advantages of white people in society, the fact that racial problems are not isolated situations, that slavery created conditions that make it hard for Blacks to advance, and that the success of other minorities proves that Blacks can do the same. For this index, we also find consistently negative but noisily estimated effects of social media use on attitudes towards racism, indicating that respondent become more aware of racial issues.

Lastly, we gather more evidence on whether the observed differences may be driven by selection rather than persuasion. For this purpose, we investigate attitudes towards two unrelated issues: support for abortion rights and environmental protection. These topics were crowded out by racial equity, police brutality as well as health and safety debates on social media in 2020. This exercise allows us to assess whether the shift in preferences is universal or whether it is driven by issue-specific persuasion on the platform and it allows us to assuage concerns that compliers are generally more prone to changing their preferences when connecting to social media. Columns 6 and 7 show that the estimates are small in magnitude and insignificant, confirming that attitudes towards women’s reproductive health and environmental protection do not change.

Overall, we find that the age disparity in progressive attitudes decreases in response to late adoption. In other words: older respondents converge more to the preferences and behaviors of younger respondents in counties exposed to the push-pull shock.²³ They are more likely to protest, be politically active on social media and have more progressive attitudes towards the police and racial equity - but not other progressive issues. These results highlight that the influence of social media is issue-specific and that persuasion can be shaped by the existing network and issues amplified on the platform.

5.3 Attitudes towards BLM

We exploit more short-term, BLM-specific information from the PEW American Trends Panel Survey, which was conducted at the height of BLM protest between June 4th and June 10th of 2020. We describe the underlying data in more detail in Appendix Section A. Since the location of the respondent is anonymized in the survey, we cannot apply any of the identification strategies

²²See Table A9 for exact wording of these questions. We detail the construction of the first principal component in Table A13. We show the correlation between all measures, the eigenvalue of the first principal component and the factors loadings for each variable. Every variable is coded in the expected direction: they all appear with a positive coefficient, except the question about increasing police presence.

²³We show in Figure A5 that older respondents hold more conservative views on average. Looking at the outcomes considered in Table 4 by age cohort, we find that the older age cohorts are on average less likely to protest, be politically active on social media, vote for Biden or hold views that are critical of police and that are anti-racist.

discussed in the previous sections. The only available information is the severity of exposure to COVID-19 in the respondent’s county of residence in June 2020. However, the rich set of individual-level controls and placebo checks assuage concerns about omitted variable bias. In all regressions, we control for respondents’ race, whether or not they live in a metropolitan area, gender, age, education, income and whether or not they lean towards the Democratic party.

Table 5 shows the results. Columns 1 and 2 show the intensity and form of news consumption in the context of George Floyd’s murder. Higher levels of COVID-19 are positively and significantly associated with more news consumption about George Floyd and more social media news consumption about George Floyd. Then, we analyze whether this change in mode of news consumption is accompanied by a change in attitudes. In column 3 and 4, we find that respondents are also more likely to agree with the statement that the BLM protests arise because of structural racism and not as an excuse for criminal behavior.

In column 5, we verify whether pandemic exposure increased the salience of racial inequality more generally and extends to settings unrelated to George Floyd and BLM. We do not find evidence that respondents are more likely to report that higher hospitalization rates of Blacks during the pandemic are caused by circumstances beyond their control, rather than personal choices or lifestyle. To rule out that exposure to COVID-19 in the earlier stages of the pandemic is just a proxy for more progressive leaning counties, we use in column 6 an additional question that deals with an unrelated issue: legal status for undocumented immigrants. Individuals living in counties with higher exposure to COVID-19 are not more likely to prefer more rights for undocumented immigrants, alleviating some of the concern about unobserved heterogeneity.

6 Conclusion

Our findings highlight how the later stages of social media expansion reshape the dynamics of collective action. Late adopters drive protest activity not (only) by lowering coordination costs but by altering preferences, particularly in movements with well-established networks like Black Lives Matter. This movement-specific effect underscores the importance of platform consolidation, where dominant narratives can shape the attitudes of new users. Unlike early adopters, who encounter an open, less politicized environment, late adopters encounter platforms already aligned with prominent causes, amplifying their influence.

Overall, our study suggests that late adoption draws a more ideologically distant group of users onto the platform, but the persuasive power of dominant narratives can lead to greater ideological alignment with the prevailing network. This dynamic highlights the potential for social media platforms with right-leaning networks to similarly influence late adopters in the opposite direction, reinforcing existing polarization. In addition, the political capture of those platforms and the subsequent reshaping of the user base could have important consequences for political persuasion - which not need to be active content moderation but driven by changes in the platform’s network composition.

These findings have broad implications for understanding the role of social media in shaping political mobilization, electoral outcomes, and ideological divides in digital spaces. This may also explain why many of the concerns about social media are associated with their expansion rather

than initial roll-out. Future research should explore how the dynamics of social media adoption evolve as platforms mature, particularly the interplay between network consolidation, user composition, and the shaping of political preferences. Investigating how these factors differ across platforms with varying ideological leanings or cultural contexts could provide deeper insights into the broader implications of social media on collective action and political polarization.

References

- Adam, Dillon C, Peng Wu, Jessica Y Wong, Eric HY Lau, Tim K Tsang, Simon Cauchemez, Gabriel M Leung, and Benjamin J Cowling, “Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong,” *Nature Medicine*, 2020, 26 (11), 1714–1719.
- Anderson, Monica and Paul Hitlin, “The hashtag# BlackLivesMatter emerges: Social activism on Twitter,” *Pew Research Center*, 2016, 15.
- Aridor, Guy, Rafael Jiménez-Durán, Ro’ee Levy, and Lena Song, “The economics of social media,” *Journal of Economic Literature*, 2024, 62 (4), 1422–1474.
- Arthur, W Brian, “Competing technologies, increasing returns, and lock-in by historical events,” *The Economic Journal*, 1989, 99 (394), 116–131.
- Askitas, Nikolaos, Konstantinos Tatsiramos, and Bertrand Verheyden, “Estimating worldwide effects of non-pharmaceutical interventions on COVID-19 incidence and population mobility patterns using a multiple-event study,” *Scientific reports*, 2021, 11 (1), 1972.
- Ba, Bocar A, Roman Rivera, and Alexander Whitefield, “Market Response to Racial Uprisings,” Working Paper 31606, National Bureau of Economic Research August 2023.
- Barbera, Salvador and Matthew O Jackson, “A model of protests, revolution, and information,” *Revolution, and Information (October 2019)*, 2019.
- Becker, Sascha O, Yuan Hsiao, Steven Pfaff, and Jared Rubin, “Multiplex network ties and the spatial diffusion of radical innovations: Martin Luther’s leadership in the early reformation,” *American Sociological Review*, 2020, 85 (5), 857–894.
- Berman, Nicolas, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig, “This Mine Is Mine! How Minerals Fuel Conflicts in Africa,” *American Economic Review*, June 2017, 107 (6), 1564–1610.
- Besley, Timothy and Torsten Persson, “The logic of political violence,” *The Quarterly Journal of Economics*, 2011, 126 (3), 1411–1445.
- Black Lives Matter, “Home - Black Lives Matter,” <https://blacklivesmatter.com/> 2020. Accessed: 2020-10-31.
- Boken, Johannes, Draca Mirko, Nicola Mastrorocco, and Arianna Ornaghi, “The Returns to Viral Media: The Case of US Campaign Contributions,” Technical Report, CEPR DP 18337 2023.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel, “Quasi-experimental shift-share research designs,” *The Review of Economic Studies*, 2022, 89 (1), 181–213.
- Boxell, Levi, Matthew Gentzkow, and Jesse M Shapiro, “Greater Internet use is not associated with faster growth in political polarization among US demographic groups,” *Proceedings of the National Academy of Sciences*, 2017, 114 (40), 10612–10617.
- Boyer, Pierre C, Thomas Delemotte, Germain Gauthier, Vincent Rollet, and Benoît Schmutz, “Social Media and the Dynamics of Protests,” Technical Report, CESifo Working Paper Series 8326, <http://dx.doi.org/10.2139/ssrn.3618851> 2020.
- Burchardi, Konrad B, Thomas Chaney, and Tarek A Hassan, “Migrants, ancestors, and foreign investments,” *The Review of Economic Studies*, 2019, 86 (4), 1448–1486.
- Bursztyn, Leonardo, Benjamin R Handel, Rafael Jimenez, and Christopher Roth, “When Product Markets Become Collective Traps: The Case of Social Media,” Working Paper 31771, National Bureau of Economic Research October 2023.
- , Davide Cantoni, David Y Yang, Noam Yuchtman, and Y Jane Zhang, “Persistent political engagement: Social interactions and the dynamics of protest movements,” *American Economic Review: Insights*, 2021, 3 (2), 233–50.
- Campante, Filipe, Emilio Depetris-Chauvin, and Ruben Durante, “The Virus of Fear: The Political Impact of Ebola in the United States,” *American Economic Journal: Applied Economics*, 2024, 16 (1), 480–509.

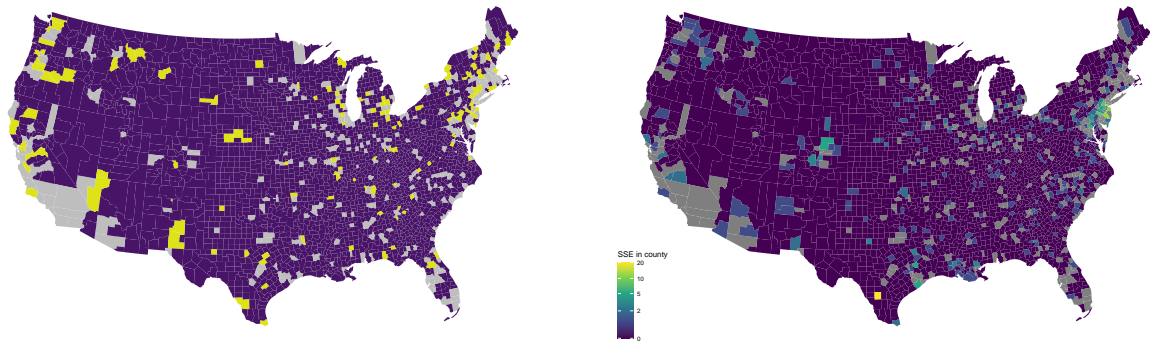
- , **Ruben Durante**, and **Andrea Tesei**, “Media and social capital,” *Annual Review of Economics*, 2022, 14, 69–91.
- , – , and **Francesco Sobbrío**, “Politics 2.0: The multifaceted effect of broadband Internet on political participation,” *Journal of the European Economic Association*, 2018, 16 (4), 1094–1136.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang**, “Protests as Strategic Games: Experimental Evidence from Hong Kong’s Anti-Authoritarian Movement,” *Quarterly Journal of Economics*, May 2019, 134 (2).
- , **Louis-Jonas Heizlsperger, David Y Yang, Noam Yuchtman, and Y Jane Zhang**, “The fundamental determinants of protest participation: Evidence from Hong Kong’s antiauthoritarian movement,” *Journal of Public Economics*, 2022, 211, 104667.
- Cappelli, Mary Louisa et al.**, “Black Lives Matter: The emotional and racial dynamics of the George Floyd protest graffiti,” *Advances in Applied Sociology*, 2020, 9 (10), 323.
- Casanueva, Annalí**, “Can chants in the street change politics’ tune? Evidence from the 15M movement in Spain,” *Available at SSRN 3831874*, 2021.
- Celislami, Elda, Stephen Kastoryano, and Giovanni Mastrobuoni**, “Strategic Bureaucratic Opacity: Evidence from Death Investigation Laws and Police Killings,” Technical Report 16609, IZA Discussion Papers 2023.
- Chenoweth, Erica, Barton H Hamilton, Hedwig Lee, Nicholas W Papageorge, Stephen P Roll, and Matthew V Zahn**, “Who Protests, What Do They Protest, and Why?,” Technical Report, National Bureau of Economic Research 2022.
- Christensen, Darin and Francisco Garfias**, “Can you hear me now? How communication technology affects protest and repression,” *Quarterly Journal of Political Science*, 2018, 13 (1), 89.
- Christensen, Henrik Serup**, “Political activities on the Internet: Slacktivism or political participation by other means?,” *First Monday*, 2011.
- Colella, Fabrizio, Rafael Lalive, Seyhun Orcan Sakalli, and Mathias Thoenig**, “Inference with Arbitrary Clustering,” IZA Discussion Papers 12584, Institute of Labor Economics (IZA) August 2019.
- Conley, Timothy G**, “GMM estimation with cross sectional dependence,” *Journal of Econometrics*, 1999, 92 (1), 1–45.
- Dave, Dhaval M, Andrew I Friedson, Kyutaro Matsuzawa, Joseph J Sabia, and Samuel Safford**, “Black Lives Matter Protests and Risk Avoidance: The Case of Civil Unrest During a Pandemic,” Working Paper 27408, National Bureau of Economic Research June 2020.
- Dube, Oeindrila and Juan F Vargas**, “Commodity price shocks and civil conflict: Evidence from Colombia,” *The Review of Economic Studies*, 2013, 80 (4), 1384–1421.
- Dunivin, Zackary Okun, Harry Yaojun Yan, Jelani Ince, and Fabio Rojas**, “Black Lives Matter protests shift public discourse,” *Proceedings of the National Academy of Sciences*, 2022, 119 (10), e2117320119.
- Edwards, David A, Dennis Ausiello, Jonathan Salzman, Tom Devlin, Robert Langer, Brandon J Beddingfield, Alyssa C Fears, Lara A Doyle-Meyers, Rachel K Redmann, Stephanie Z Killeen et al.**, “Exhaled aerosol increases with COVID-19 infection, age, and obesity,” *Proceedings of the National Academy of Sciences*, 2021, 118 (8), e2021830118.
- Eisinger, Peter K**, “The conditions of protest behavior in American cities,” *The American Political Science Review*, 1973, 67 (1), 11–28.
- Ellis, C.J. and J. Fender**, “Information cascades and revolutionary regime transitions,” *Economic Journal*, 2011, 121 (553), 763–792.
- Endo, Akira, Sam Abbott, Adam J Kucharski, Sebastian Funk et al.**, “Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China,” *Wellcome Open Research*, 2020, 5 (67), 67.

- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova**, “Social media and protest participation: Evidence from Russia,” *Econometrica*, 2020, 88 (4), 1479–1514.
- , **Maria Petrova, and Konstantin Sonin**, “Social media and corruption,” *American Economic Journal: Applied Economics*, 2018, 10 (1), 150–74.
- , – , **Gianluca Russo, and David Yanagizawa-Drott**, “Socializing Alone: How Online Homophily Has Undermined Social Cohesion in the US,” *Available at SSRN*, 2024.
- Falck, Oliver, Robert Gold, and Stephan Heblich**, “E-lelections: Voting Behavior and the Internet,” *American Economic Review*, 2014, 104 (7), 2238–65.
- Freelon, Deen, Charlton D McIlwain, and Meredith Clark**, “Beyond the hashtags: #Ferguson, #Blacklivesmatter, and the online struggle for offline justice,” *Center for Media & Social Impact, American University*, 2016.
- Freyaldenhoven, Simon, Christian Hansen, Jorge Pérez Pérez, and Jesse M. Shapiro**, “Visualization, identification, and estimation in the linear panel event-study design.,” *Advances in Economics and Econometrics: Twelfth World Congress*, forthcoming. Stata package and article, Video series, R package.
- Fujiwara, Thomas, Karsten Müller, and Carlo Schwarz**, “The effect of social media on elections: Evidence from the United States,” *Journal of the European Economic Association*, 2023, p. jvad058.
- García-Jimeno, Camilo, Angel Iglesias, and Pinar Yildirim**, “Information Networks and Collective Action: Evidence from the Women’s Temperance Crusade,” *American Economic Review*, 2022, 112 (1), 41–80.
- Garcia, Raffi E and Alberto Ortega**, “Racial Protests and Credit Access,” Working Paper 32477, National Bureau of Economic Research May 2024.
- Gethin, Amory and Vincent Pons**, “Social Movements and Public Opinion in the United States,” Working Paper 32342, National Bureau of Economic Research April 2024.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik instruments: What, when, why, and how,” *American Economic Review*, 2020, 110 (8), 2586–2624.
- González, Felipe and Mounu Prem**, “Police Violence, Student Protests, and Educational Performance,” *Review of Economics and Statistics*, 2024, 106 (3), 712–727.
- Guriev, Sergei and Daniel Treisman**, “A theory of informational autocracy,” *Journal of Public Economics*, 2020, 186, 104158.
- , **Nikita Melnikov, and Ekaterina Zhuravskaya**, “3G Internet and confidence in government,” *The Quarterly Journal of Economics*, 2021, 136 (4), 2533–2613.
- Hale, Thomas, Samuel Webster, Anna Petherick, Toby Phillips, and Beatriz Kira**, “Oxford COVID-19 government response tracker,” *Blavatnik School of Government*, 2020, 25.
- Jinjarak, Yothin, Rashad Ahmed, Sameer Nair-Desai, Weining Xin, and Joshua Aizenman**, “Accounting for global COVID-19 diffusion patterns, January–April 2020,” *Economics of Disasters and Climate Change*, 2020, 4, 515–559.
- Keib, Kate, Itai Himelboim, and Jeong-Yeob Han**, “Important tweets matter: Predicting retweets in the # BlackLivesMatter talk on twitter,” *Computers in human behavior*, 2018, 85, 106–115.
- Lancaster, Tony**, “The incidental parameter problem since 1948,” *Journal of Econometrics*, 2000, 95 (2), 391–413.
- Leclerc, Quentin J, Naomi M Fuller, Lisa E Knight, Sebastian Funk, Gwenan M Knight, cmmid covid 19 Working Group et al.**, “What settings have been linked to SARS-CoV-2 transmission clusters?,” *Wellcome Open Research*, 2020, 5.
- Lipsky, Michael**, “Protest as a political resource,” *The American Political Science Review*, 1968, 62 (4), 1144–1158.
- Manacorda, Marco and Andrea Tesei**, “Liberation technology: Mobile phones and political mobilization in Africa,” *Econometrica*, 2020, 88 (2), 533–567.

- Manski, Charles F, Daniel McFadden et al.**, *Structural analysis of discrete data with econometric applications*, MIT Press Cambridge, MA, 1981.
- Mathieu, Lilian**, “Art and Social Movements,” *The Wiley Blackwell Companion to Social Movements*, 2018, pp. 354–368.
- Mazumder, Soumyajit**, “The persistent effect of US civil rights protests on political attitudes,” *American Journal of Political Science*, 2018, *62* (4), 922–935.
- , “Black Lives Matter for Whites’ Racial Prejudice: Assessing the Role of Social Movements in Shaping Racial Attitudes in the United States,” Technical Report, Harvard University May 2019.
- McCarthy, John D and Mayer N Zald**, “Resource mobilization and social movements: A partial theory,” *American Journal of Sociology*, 1977, *82* (6), 1212–1241.
- McKersie, Robert B**, “The 1960s Civil Rights Movement and Black Lives Matter: Social protest from a negotiation perspective,” *Negotiation Journal*, 2021, *37* (3), 301–323.
- Melnikov, Nikita**, “Mobile Internet and Political Polarization,” Available at SSRN 3937760, 2021.
- Mesquita, Ethan Bueno De**, “Regime Change and Revolutionary Entrepreneurs,” *American Political Science Review*, 2010, *104* (3), 446–466.
- Miller, Danielle, Michael A Martin, Noam Harel, Omer Tirosh, Talia Kustin, Moran Meir, Nadav Sorek, Shiraz Gefen-Halevi, Sharon Amit, Olesya Vorontsov et al.**, “Full genome viral sequences inform patterns of SARS-CoV-2 spread into and within Israel,” *Nature Communications*, 2020, *11* (1), 5518.
- MIT Election Data and Science Lab**, “County Presidential Election Returns 2000-2016,” 2018.
- Müller, Karsten and Carlo Schwarz**, “Fanning the flames of hate: Social media and hate crime,” *Journal of the European Economic Association*, 2021, *19* (4), 2131–2167.
- and – , “From hashtag to hate crime: Twitter and antiminority sentiment,” *American Economic Journal: Applied Economics*, 2023, *15* (3), 270–312.
- Ostrom, Elinor**, *Governing the commons: The evolution of institutions for collective action*, Cambridge University Press, 1990.
- Passarelli, Francesco and Guido Tabellini**, “Emotions and political unrest,” *Journal of Political Economy*, 2017, *125* (3), 903–946.
- Pew**, “Twitter Conversations About News Events: Politically Polarized, Digitally Divided,” 2019. Accessed: 2024-12-04.
- Qin, Bei, David Strömberg, and Yanhui Wu**, “Social Media and Collective Action in China,” *Econometrica*, 2024, *92* (6), 1993–2026.
- Rupasingha, Anil, Stephan J Goetz, and David Freshwater**, “The production of social capital in US counties,” *The Journal of Socio-Economics*, 2006, *35* (1), 83–101.
- Sangnier, Marc and Yanos Zylberberg**, “Protests and trust in the state: Evidence from African countries,” *Journal of Public Economics*, August 2017, *152*, 55–67.
- Schumann, Sandy and Olivier Klein**, “Substitute or stepping stone? Assessing the impact of low-threshold online collective actions on offline participation,” *European Journal of Social Psychology*, 2015, *45* (3), 308–322.
- Takhteyev, Yuri, Anatoliy Gruzd, and Barry Wellman**, “Geography of Twitter networks,” *Social Networks*, 2012, *34* (1), 73–81.
- Wooldridge, Jeffrey M**, *Introductory econometrics: A modern approach*, Cengage Learning, 2015.
- Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov**, “Political effects of the Internet and social media,” *Annual Review of Economics*, 2020, *12*, 415–438.

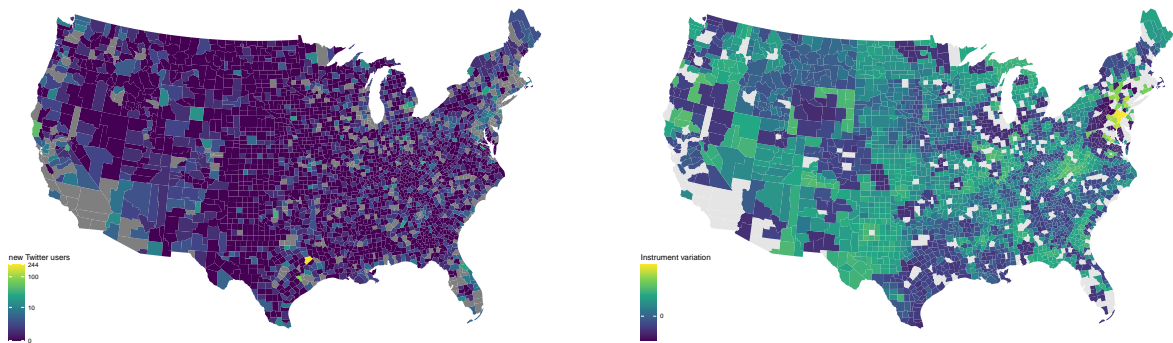
Figures and Tables

Figure 1: Geographic distribution of BLM protest, social media and super spreading events (SSEs)



(a) BLM protest after George Floyd

(b) Location of SSEs



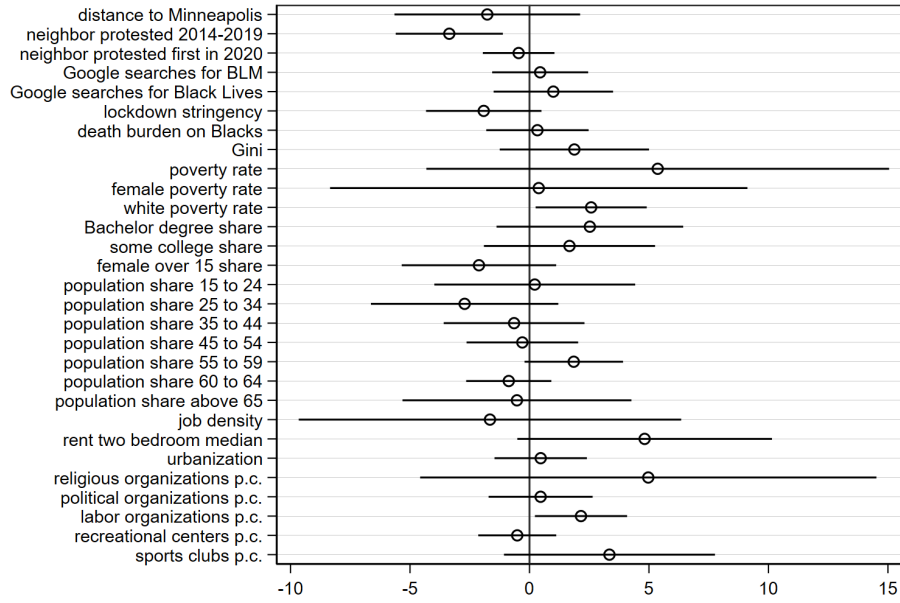
(c) Late Twitter adoption
(log-transformed number of new Twitter users)

(d) Instrument variation SSE \times SXS
(net of controls)

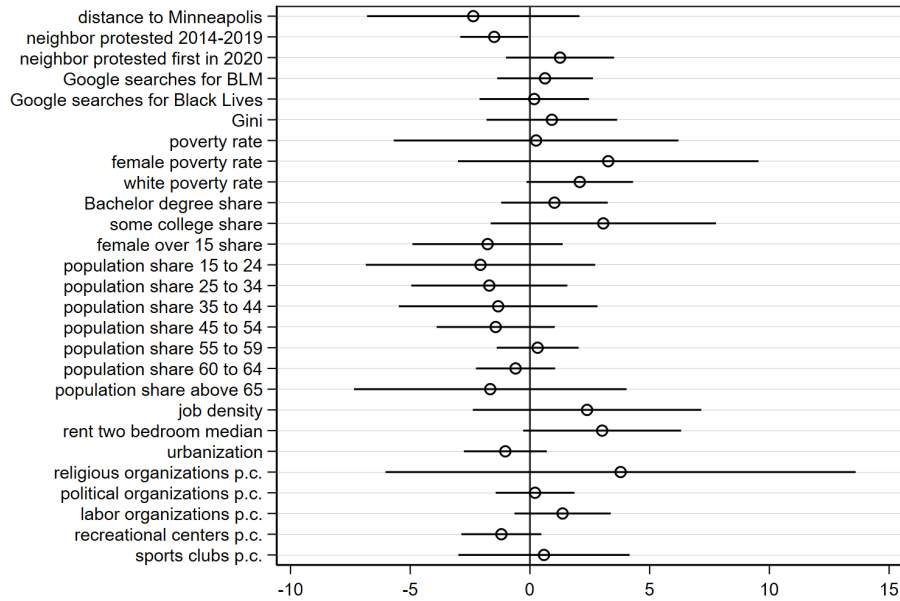
Note: Geographic distribution of outcome, treatment and instrument across counties. Sub-figure (a) yellow indicates the location of counties with at least one BLM protest as reported by *Elephrame* in the three weeks following George Floyd's murder among counties that never protested before. Gray indicates counties that has a BLM protest before May 25 2020. Dark purple color indicates counties without BLM protest in 2020. Sub-figure (b) shows the number of SSE occurring in each county, until 6 weeks before the murder of George Floyd. Sub-figure (c) shows the log of one plus the number of new Twitter profiles created between January and May 2020 observed in a random sample of tweets. Sub-figure (d) shows the variation of the instrument, i.e. the interaction SSE \times SXS net of all controls and state fixed effects.

Figure 2: **Plausibility of quasi random exposure** (dep. var.: $SSE_{-c} \times SXSW_c$)

(a) without state fixed effects

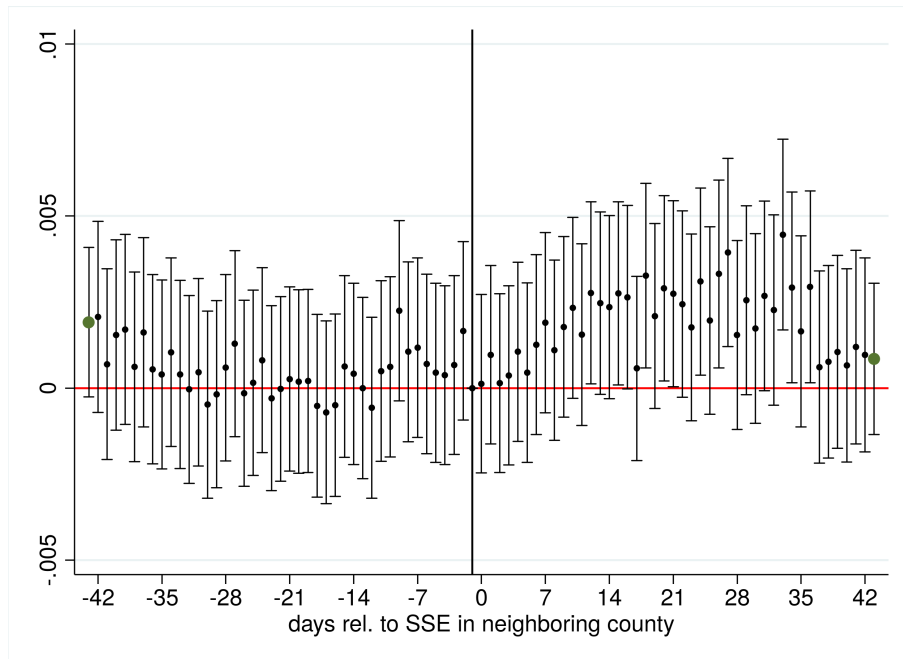


(b) with state fixed effects



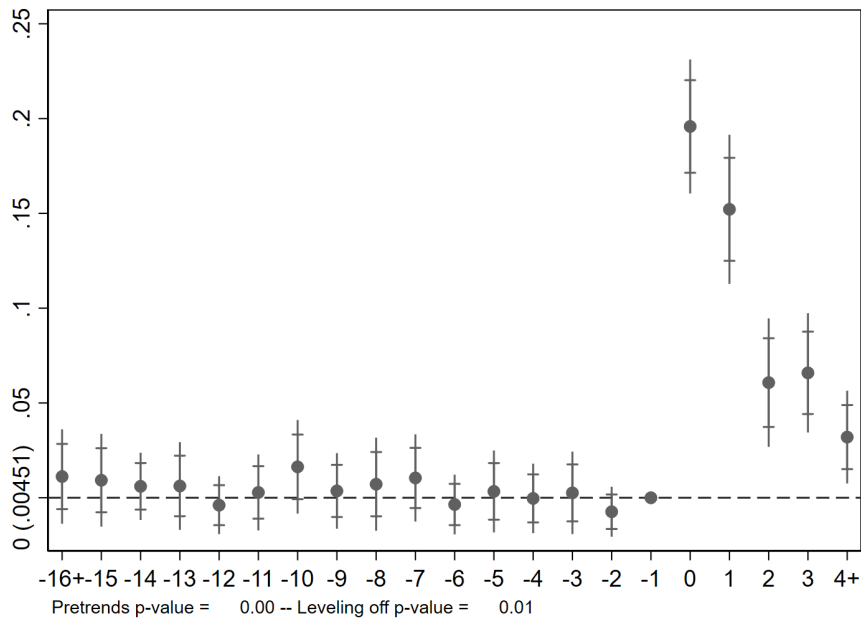
Note: Plausibility of the instrument exogeneity. We show standardized point estimates with 95% confidence intervals, clustered at the state level. We run the baseline specification but use our instruments as the outcome variable. The instrument is the interaction between the cumulative number of SSEs within a 200km radius outside of the county up to six weeks before the murder of Floyd SSE_{-c} and the baseline Twitter network $SXSW_c$, which is the logarithm of one plus the number of SXSW followers who created their account in March 2007 in the county and neighboring counties. We exclude state fixed effects in panel (a) since the death burden on Blacks and lockdown stringency is only available at the state level. We describe all variables in detail in Table A2.

Figure 3: **Timing of super spreading events (SSEs) and subsequent Twitter adoption**



Note: Event study graph according to following specification: $\text{New Twitter Users}_{ct} = \sum_{k=T_0}^{-1} \beta_k \times \text{SSE}_{-c,k} + \sum_{k=0}^{T_1} \beta_k \times \text{SSE}_{-c,k} + \zeta_c + \delta_t + \epsilon_{ct}$ with county and time fixed effects. Time is measured in days before and after a SSE within 200km of the county border.

Figure 4: **Event study at the county-month level: BLM protest after murder of George Floyd in counties with higher number of late adopters**



Note: Event study for the full samples of counties (including those with prior BLM protest) using *xtevent* by Freyaldenhoven, Hansen, Pérez and Shapiro (forthcoming) following specification 3 at the county-month level in the 16 months leading up to and 4 months following the murder of George Floyd. Regressions include county fixed effects and state-month fixed effects. We interact the post George Floyd period with the log of one plus the number of new Twitter users created after the outbreak of the pandemic but before the murder of George Floyd. Standard errors are clustered at the county level. Outcome is measured as a dummy variable for any BLM protest in that month.

Table 1: **Late adoption increases BLM protest**

| | At least one BLM protest | | | |
|------------------------------|--------------------------|----------------------------|------------------------|------------------------|
| Panel A: 2SLS | (1) | (2) | (3) | (4) |
| new Twitter users | 0.377*** (0.117) | 0.439*** (0.139) | 0.439*** (0.141) | 0.443*** (0.146) |
| Kleibergen-Paap F stat | 34.96 | 27.21 | 24.14 | 25.89 |
| Panel B: OLS | | | | |
| new Twitter users | 0.0318*** (0.00963) | 0.0303*** (0.00902) | 0.0292*** (0.00875) | 0.0297*** (0.00862) |
| Panel C: Reduced form | | | | |
| SSE \times SXS | 0.0891*** (0.0196) | 0.104*** (0.0282) | 0.104*** (0.0288) | 0.104*** (0.0297) |
| Panel D: First stage | | New Twitter users (log +1) | | |
| SSE \times SXS | 0.237*** (0.0400) | 0.237*** (0.0454) | 0.236*** (0.0480) | 0.234*** (0.0460) |
| Observations | 2767 | 2767 | 2767 | 2767 |
| Mean dep. var. | 0.0477 | 0.0477 | 0.0477 | 0.0477 |
| State FE | Yes | Yes | Yes | Yes |
| Pandemic exposure | Yes | Yes | Yes | Yes |
| Baseline Twitter | Yes | Yes | Yes | Yes |
| Economic controls | | Yes | Yes | Yes |
| BLM controls | | | Yes | Yes |
| Political controls | | | | Yes |

Note: Estimation results from specification 1. New Twitter users are measured as the log of one plus new geo-located accounts at the county level created after the beginning of the pandemic but before George Floyd's murder based on a random sample of tweets. Instrument $SSE_{-c} \times SXS_{-c}$ is the push-pull instrument for pandemic Twitter take-up, combining the cumulative number of SSEs outside of the county within a 200km radius SSE_{-c} with an instrument for baseline Twitter penetration SXS_{-c} following Müller and Schwarz (2023). All estimations include state fixed effect, the cumulative number of COVID-19 related deaths until May 25th (pandemic exposure) and the log transformed number of Twitter accounts created before January 2020 from a random sample of Tweets (baseline Twitter) as controls. 2SLS, reduced form and first stage additionally include SSE_{-c} and SXS_{-c} separately, as well as the logarithm of one plus the number of SXS Twitter followers before the Austin festival and its interaction with SSE_{-c} . The outcome is a dummy variable for any BLM protest in the three weeks following the murder of George Floyd, using information from *Elephrame*. Economic controls: average unemployment rate in the preceding year, median household income in 2016, pandemic resilience, a dummy for urban counties. BLM controls: share of Black population, Black poverty rate, and deadly force used by police against Black people in 2014-2019 and in 2020. Political controls: Republican vote share in 2012 and 2016 and social capital (number of civic organizations per capita). We report Kleibergen-Paap rkWald F statistic. Standard errors (in parentheses) are clustered at the state level. First stage and reduced form coefficients are multiplied by 100 for readability. *** p<0.01, ** p<0.05, * p<0.1

Table 2: Late adoption and scope of offline and online BLM protest

| | at least one BLM protest | number of BLM protests | log total participants | log BLM Tweets by new users | log BLM account followers among new users | 1st PC BLM protest |
|------------------------|-----------------------------|---------------------------|---------------------------|--------------------------------|--|-----------------------|
| 2SLS: | (1) | (2) | (3) | (4) | (5) | (6) |
| new Twitter users | 0.443*** (0.146) | 0.704*** (0.259) | 2.036** (0.806) | 0.882** (0.331) | 0.375*** (0.127) | 3.224*** (0.994) |
| Observations | 2767 | 2767 | 2767 | 2767 | 2767 | 2767 |
| Mean dep. var. | 0.0477 | 0.0636 | 0.234 | 0.201 | 0.200 | 0 |
| Kleibergen-Paap F stat | 25.89 | 25.89 | 25.89 | 25.89 | 25.89 | 25.89 |
| State FE | Yes | Yes | Yes | Yes | Yes | Yes |
| All controls | Yes | Yes | Yes | Yes | Yes | Yes |

Note: Estimation results from specification 1. New Twitter users are measured as the log of one plus new geo-located accounts at the county level created after the beginning of the pandemic but before George Floyd’s murder based on a random sample of Tweets. Instrument $SSE_{-c} \times SXS_{W_c}$ is the push-pull instrument for pandemic Twitter take-up, combining pandemic exposure SSE_{-c} with baseline Twitter penetration SXS_{W_c} following Müller and Schwarz (2023). Columns 1 to 3 use protest information from *Elephrame*, reporting respectively a dummy variable for any BLM-related protest in the three weeks following the murder of George Floyd, the number of these protests and the total number of participants. Column 4 reports the logarithm of one plus number of geo-located Tweets by new Twitter users that use at least one BLM-related hashtag or keyword in the three weeks following the murder. Column 6 reports the logarithm of one plus number of geo-located accounts of new Twitter users that follow the official BLM account @BlkLivesMatter. Outcome of column 6 is the first principal component of all outcomes used in columns 1 to 5. All specifications include state fixed effects. They also include SSE_{-c} and SXS_{W_c} separately, COVID per capita, old Twitter users, the (log) number of SXSW Twitter followers before the Austin festival and its interaction with SSE_{-c} . Control variables include: the share of Black population, urban, median household income, unemployment share, Black poverty rate, 3+ risk factors/community resilience, Republican vote share in 2012 and 2016, social capital (number of civic organizations per capita) and deadly force used by police against Black people. We report Kleibergen-Paap rkWald F statistic. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Late adoption and other protests: graffiti, counter mobilization, and COVID-19 protest

| 2SLS: | George Floyd Street Art (1) | All 2020 BLM protest (2) | All 2020 other protest (3) | Anti mask & distancing (4) | Populist & pro-Trump (5) | log All Lives Matter Tweets (6) | log Blue Lives Matter Tweets (7) |
|-------------------------------|-----------------------------------|--------------------------------|----------------------------------|----------------------------------|--------------------------------|---------------------------------------|--|
| new Twitter users | -0.210 (0.257) | 6.498*** (2.081) | 13.75*** (3.994) | 4.642*** (1.091) | -0.0707 (0.169) | -0.0641 (0.175) | 0.304 (0.258) |
| Observations | 2767 | 2767 | 2767 | 2767 | 2767 | 2767 | 2767 |
| Mean dep. var. Mean dep. var. | 0.00939 | 1.314 | 1.811 | 0.365 | 0.0549 | 2.581 | 1.024 |
| Kleibergen-Paap F stat | 25.89 | 25.89 | 25.89 | 25.89 | 25.89 | 25.89 | 25.89 |
| State FE | Y | Y | Y | Y | Y | Y | Y |
| All controls | Y | Y | Y | Y | Y | Y | Y |

Note: Estimation results from specification 1. New Twitter users are measured as the log of one plus new geo-located accounts at the county level created after the beginning of the pandemic but before George Floyd's murder based on a random sample of Tweets. Instrument $SSE_{-c} \times SXS_{W_c}$ is the push-pull instrument for pandemic Twitter take-up, combining pandemic exposure SSE_{-c} with baseline Twitter penetration SXS_{W_c} following Müller and Schwarz (2023). Outcome in column 1 is scraped from the Urban Art Mapping George Floyd and Anti-Racist Street Art database. Outcomes in columns 2 to 5 come from the ACLED US Crisis Monitor 2020 and count the number of protests associated with each topic after May 25, 2020. Outcomes in column 6 and 7 are the log of 1 + number of geo-localized Tweets that contain the hashtags or phrases #AllLivesMatter and #BlueLivesMatter. All specifications include state fixed effects. They also include SSE_{-c} and SXS_{W_c} separately, COVID per capita, old Twitter users, the (log) number of SXS_{W} Twitter followers before the Austin festival and its interaction with SSE_{-c} . Control variables include: the share of Black population, urban, median household income, unemployment share, Black poverty rate, 3+ risk factors/community resilience, Republican vote share in 2012 and 2016, social capital (number of different types of civic organizations) and deadly force used by police against Black people. We report Kleibergen-Paap rkWald F statistic. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table 4: Individual-level evidence on social media use and political preferences

| 2SLS | has protested | political activity | voted for | PC1 | PC1 | Placebo | |
|----------------------------------|--------------------|---------------------|-------------------|---------------------|-------------------|-------------------|-------------------|
| | for any cause | on social media | Biden in 2020 | police | racism | abortion | environment |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| Social media use | 0.472** (0.231) | 1.068*** (0.220) | 0.523* (0.310) | -2.434** (1.065) | -1.066 (1.039) | 0.0916 (0.264) | -0.274 (0.232) |
| Observations | 48,382 | 48,382 | 48,382 | 48,308 | 47,093 | 48,369 | 48,366 |
| Mean dep. var. | 0.0783 | 0.526 | 0.426 | 0 | 0 | 0.611 | 0.699 |
| Kleibergen-Paap F stat | 13.96 | 13.96 | 13.96 | 13.88 | 11.05 | 14.12 | 14.04 |
| Age \times Instrument controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Individual controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| County FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

Note: Individual level regression following specification 4. Outcomes are taken from the Cooperative Congressional Election Study conducted in November of 2020. Column 1 is dummy variable for whether respondent has attended a protest, march or demonstration in the past year. Column 2 is a dummy variable for whether the respondent has used social media in the past 24 hours to post or forward a story, video or link about politics, or to post a comment about politics, or to watch a video about politics. Column 3 is a dummy variable for whether respondent voted for Biden in the 2020 presidential election. Column 4 is the first principal component of questions that ask whether the respondent supports more oversight of police including allowing individuals and their families to sue police officers, ending the DoD program that sends military surplus to police departments, creating a national registry of police who have been investigated, banning the use of choke holds, requiring police to wear body cameras, having more or less police in the streets. Column 5 is the first principal component of anti-racist attitudes, including that white people have advantages in society, that racial problems are not isolated situations, that slavery created conditions that make it hard for Blacks to advance, that the success of other minorities does not prove that Blacks can do the same. Column 6 asks whether respondents support the statement "A woman should always be allowed to obtain an abortion as a matter of choice". Column 7 asks whether respondents support the statement "The Environmental Protection Agency should have the power to regulate Carbon". Treatment is a dummy variable indicating if the respondent used social media in the past 24 hours. The instrument is the previous push-pull instrument $SSE_c \times SXSX_c$ interacted with the age of the respondent, but using SSEs inside of the county instead of SSEs in neighboring counties. All models include county fixed effects, individual controls (age, gender, employment status, number of children, dummies for race and for vote in the 2016 presidential election), and interaction of age with instrument controls (SSE and SXSX individually, pre-SXSX followers and its interaction with age, COVID deaths per thousand and baseline social media penetration). We report Kleibergen-Paap rkWald F statistic. Standard errors are clustered at the county level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Individual-level evidence on pandemic exposure, news consumption and attitudes: PEW June 2020

| | News consumption | | Attitudes towards BLM | | Other Attitudes | |
|---------------------------------------|-----------------------------|--|---|---|---|--|
| | Follow news about GF (1) | Receive news about GF on social media (2) | BLM protest to justify criminal behavior (3) | BLM protest because of structural racism (4) | Higher Black COVID hospitaliz. not their fault (5) | Rights of undocumented migrants (6) |
| COVID-19 deaths per capita [category] | 0.0445*** (0.00952) | 0.0252* (0.0152) | -0.0195* (0.0107) | 0.0240*** (0.00886) | 0.00612 (0.00629) | -0.00496 (0.00531) |
| Observations | 9,048 | 9,048 | 9,048 | 9,048 | 9,048 | 9,048 |
| R-squared | 0.066 | 0.152 | 0.145 | 0.104 | 0.129 | 0.082 |
| Mean dep. var. | 3.336 | 2.721 | 3.341 | 3.600 | 1.605 | 1.797 |
| Black | Y | Y | Y | Y | Y | Y |
| Metropolitan area | Y | Y | Y | Y | Y | Y |
| Female | Y | Y | Y | Y | Y | Y |
| Democrat | Y | Y | Y | Y | Y | Y |
| Age [category] | Y | Y | Y | Y | Y | Y |
| Education [category] | Y | Y | Y | Y | Y | Y |
| Income [category] | Y | Y | Y | Y | Y | Y |

Note: Individual-level regressions of COVID-19 related deaths at the county level measured as categories low [1], medium [2] and high [3] on various attitudinal outcomes. Data from the June 2020 wave of the PEW American Trends Panel Survey. PEW does not provide exact location of respondent and only provides coarse info on pandemic exposure (i.e. three categories). All columns include controls for various characteristics of the respondent: a dummy for Black, living in a metropolitan area, identifying as female, as Democrat, and for levels of age (18-29, 30-49, 50-64, 65+), education (high school or less, some college, college graduate +) and income (30K or less, 30-75K, 75 or more). Outcomes are measured as dummy variables based on the following questions: column 1: "How closely have you been following news about the demonstrations around the country to protest the death of George Floyd, a black man who died while in police custody?"; column 2: How much, if any, news and information about the demonstrations to protest the death of George Floyd have you been getting on social media (such as Facebook, Twitter, or Instagram)?; column 3: How much, if at all, do you think each of the following has contributed to the demonstrations to protest the death of George Floyd? Some people taking advantage of the situation to engage in criminal behavior; column 4: How much, if at all, do you think each of the following has contributed to the demonstrations to protest the death of George Floyd? Longstanding concerns about the treatment of black people in the country; column 5: Do you think the reasons why black people in our country have been hospitalized with COVID-19 at higher rates than other racial or ethnic groups have more to do with... Circumstances beyond people's control; column 6: Which comes closer to your view about how to handle undocumented immigrants who are now living in the U.S.? There should be a way for them to stay in the country legally, if certain requirements are met. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Online Appendix

Contents

| | |
|---|-----------|
| Appendix A: Data Appendix | 38 |
| A.1 Details on SSEs: the push component | 38 |
| A.2 Details on SXSXW: the pull component | 38 |
| A.3 Description of Data Sources | 40 |
| A.3.1 Pandemic-related Data | 40 |
| COVID-19. | 40 |
| Community resilience. | 40 |
| Lockdown stringency. | 40 |
| A.3.2 Social Media and Protest Data | 40 |
| New and old Twitter users. | 40 |
| Twitter in the three weeks following George Floyd. | 40 |
| BLM account followers. | 41 |
| Geo-location of tweets and accounts. | 41 |
| Google mobility. | 41 |
| Google searches. | 41 |
| BLM protest from Elephrame. | 41 |
| BLM and other protest from ACLED US Crisis Monitor. | 42 |
| George Floyd Street Art. | 42 |
| A.3.3 Survey Data and other County Controls | 42 |
| The American Trends Panel survey by Pew Research Center. | 42 |
| The Cooperative Election Study. | 42 |
| Use of deadly force by police. | 43 |
| Additional county-level controls. | 43 |
| A.4 Additional Tables and Figures | 44 |
| Appendix B: Robustness Checks | 63 |
| Changing the radius and time window around SSEs. | 63 |
| Excluding SSEs in prisons. | 63 |
| Controlling for SSEs in the county. | 63 |
| Weighting SSEs by distance. | 63 |
| Controlling for past protest in neighboring counties. | 63 |
| Controlling for COVID-19 prevalence before the protest trigger. | 63 |
| Controlling for propensity to protest. | 64 |
| Excluding coastal counties and states. | 64 |
| Accounting for spatial correlation. | 64 |
| Treatment definition. | 64 |
| Probit estimation. | 64 |
| Plausibility of exclusion restriction for BLM. | 65 |
| Outcome definition. | 65 |
| Event study robustness. | 65 |

Appendix A: Data Appendix

A.1 Details on SSEs: the push component

Our identification strategy relies, in part, on records of SSEs in the early stages of the pandemic. In this section, we discuss the limitations of the SSE data set and how we address these in the empirical section. The data set is collected from various sources by researchers from the London School of Hygiene and Tropical Medicine and published as a free access database for researchers and the media under the *SARS-CoV-2 Superspreading Events from Around the World* Project.

A main challenge in the construction of this database is that there is no standard definition of a SSE. The database mainly refers to "outbreak" and "clusters" for which they use the UK Government Public Health Definition: "two or more test-confirmed cases of COVID-19 among individuals associated with a specific non-residential setting with illness onset dates within a 14-day period." The outbreak definition is expanded to "identified direct exposure between at least 2 of the test-confirmed cases in that setting (for example under one metre face to face, or spending more than 15 minutes within 2 metres) during the infectious period of one of the cases when there is no sustained local community transmission - absence of an alternative source of infection outside the setting for the initially identified cases". In our data set, the minimum number of cases associated with a SSE is 3 and only 3 SSEs have fewer than 5 cases.

The database draws from one main source: Leclerc et al. (2020) who performed a systematic review of available literature and media reports to find settings reported in peer reviewed articles and media with "outbreak" or "cluster" characteristics. There were various extensions to this data set, using articles of journalists, expanding that data set to second and third generation events by Swinkels (2020), and including the Western Pacific Region for a project of the World Health Organisation (under the project lead of Fatim Lakha, also from the London School of Tropical Medicine and Hygiene). We will primarily draw from Leclerc et al. (2020), as we focus on SSEs in the United States during the early stages of the pandemic. Note that the data we use was downloaded at the beginning of 2021 and may not include updates or corrections made afterwards.

There are various limitations in the measurement of SSEs. First, there exists some uncertainty about the exact date of the SSE. If, for instance, there was a COVID-19 cluster at a worker dormitory, the exact date of the transmission event is difficult to narrow down. In these cases, researchers make an approximation based on the timing of tests and overall case numbers. We address this concern by using the cumulative number of SSEs until a certain cut-off date (first week of April in the baseline version of the instrument), thereby not relying on the specific timing of the SSE. We also show that our results are not sensitive to different cut-off dates. Second, for many SSEs there is uncertainty about the number of individuals infected. The database always uses the lowest number cited in the articles about the SSE but actual numbers can be much higher. The actual detected number of cases will be related to testing capacity and potentially other unobserved factors at the county level. For this reason, we use the most simple version of the instrument, i.e. counting the number of SSEs rather than using the cases associated with the SSE. Third, the GPS coordinates of SSEs are almost always approximate. For instance, when an SSE occurred somewhere in city A, typically the database uses GPS coordinates for a random location within that city, not for the precise location. Again, we verify that our results are not sensitive to changing the radius around SSEs to account for potential measurement error.

Overall, the measurement error in SSEs would only bias our results if it is somehow related to the counties' overall propensity to join Twitter or to protest (and is not captured in the set of controls or state fixed effects). It is worth noting that we control for SSEs separately and the number of COVID-19 related deaths in the county in both the first and second stage estimation. This means that any direct effect of SSEs in neighboring counties as well as their direct effect through local pandemic exposure will be captured in the set of controls. However, if SSEs capture underlying county characteristics that interact with early adoption through the SXSX shock, we may capture heterogeneity in the effect of early adoption on protest rather than its effect through late adoption. Therefore, we run a balance test of SSEs with respect to an array of county characteristics in Figure A3. Reassuringly, we do not find a systematic correlation between county characteristics and SSEs, conditional on state fixed effects and baseline controls.

A.2 Details on SXSX: the pull component

This appendix details the construction of the instrument for the historical Twitter network, i.e. the pull factor of our instrument. South by Southwest (SXSX) is a film, interactive media, and music festival

and conference held annually in Austin, Texas. During the March 2007 edition Twitter was heavily promoted, leading to a rapid increase in the social network’s popularity. Müller and Schwarz (2023) use this event to construct an instrument for Twitter penetration in the US by exploiting the fact that, through network effects, places that had more accounts created by visitors to SXSW continued to have more accounts created later on. It is not possible to directly measure the accounts created by SXSW attendees: instead, Müller and Schwarz (2023) measure the number of followers of the official account of the festival (@SXSW) that joined Twitter during the month of the festival (March 2007).

To reproduce this instrument, we collect information of all the followers of the @SXSW account of the South by Southwest festival, the date they joined Twitter, and the location set in their profile. The dataset we end up with is not entirely identical to the one used by Müller and Schwarz (2023): some users created on or before March 2007 might have started or stopped following SXSW later. They might also have changed their location between the time Müller and Schwarz collected their dataset and when we collected ours (2019 versus November 2021). Finally, our geolocation method might be different: we automatically geocode the location given by the user using Nominatim, as described in the Data section. Müller and Schwarz (2023) do not detail their geolocation method. For comparison, we attribute 52% of users to US counties (excluding imprecise locations and locations outside the US). In comparison, Fujiwara et al. (2023) (reusing this instrument) are able to locate 58% of users that joined Twitter between 2006 and 2008 using a similar method.

For each county we compute the number of followers whose account was created in March 2007 and the number of users whose account was created before this date. With our data collection and user localization strategy, we find users that follow @SXSW and joined in March 2007 in 172 counties, only 67 of which did not have BLM events before (Müller and Schwarz (2023) find 155 affected counties). To increase the number of treated counties, and thus the power of our identification, we also consider users in neighboring counties: assuming that Twitter presence diffuses, in part, geographically,²⁴ these counties should also have a higher number of Twitter users. We find 817 such counties, 618 of which did not have a BLM protest before.

We approximate our pull factor with the number of users that joined during SXSW while controlling for the number of SXSW followers that joined before. This accounts for the overall interest in the SXSW festival and for the general interest in Twitter in the county before the 2007 edition. $SXSW_c$ is the logarithm of one plus the number of SXSW followers who created their account in March 2007 in the county and neighboring counties, and $Pre\ SXSW\ Users_{sc}$ is the logarithm of one plus the number of SXSW followers in the county and neighboring counties that created their account before March 2007.

Müller and Schwarz (2023) argue that SXSW 2007 provides a quasi-random shock to Twitter adoption, leveraging this as a unique opportunity to study the geographic spread of the platform. They emphasize that SXSW heavily promoted Twitter at the time, leading to a distinct spike in new user registrations among attendees, which helped drive Twitter adoption in attendees’ home counties. This “adoption shock” aligns with the concept of path dependence in technology diffusion, where an initial, localized user base can establish a network that fuels further uptake over time.

The authors provide evidence that counties with new SXSW-related Twitter followers experienced significant, lasting increases in Twitter activity following the event, without any observable pre-existing trends in social media usage before SXSW. This sharp uptick in adoption following the festival suggests that SXSW acted as an exogenous trigger for Twitter uptake in certain counties, rather than being influenced by prior demand for social media in those areas. The subsequent adoption pattern in these counties follows an S-shaped curve typical of technology diffusion, with a strong initial growth phase that stabilizes over time. Müller and Schwarz further demonstrate that early Twitter users were disproportionately linked to SXSW, with a high share of users directly or indirectly connected to the event. This close tie between early Twitter users and the SXSW network indicates that the festival had an outsized role in shaping Twitter’s geographic diffusion at its inception, lending credibility to SXSW’s role as a quasi-random exposure mechanism. Additionally, the authors control for the possibility of selection bias by comparing SXSW with other popular festivals like Coachella and Burning Man. While there is some overlap in demographics among followers of these festivals, only SXSW appears to have led to a substantial, sustained increase in Twitter adoption. This cross-festival comparison supports their argument that SXSW’s impact on Twitter uptake was uniquely influential and largely independent of other county characteristics, conditional on controls.

²⁴This assumption is also made by Müller and Schwarz at the level of a county. Here we just extend it to neighboring counties.

A.3 Description of Data Sources

A.3.1 Pandemic-related Data

COVID-19. Data on COVID-19 related deaths and cases in the USA at the county level comes from the New York Times. This data set provides the cumulative count of cases and deaths every day for each county in the USA, starting from January 21, 2020 when the country's first COVID-19 case was reported. A key limitation of COVID-19 cases data is that it depends on the testing facilities and the availability of test kits in the region. We therefore mainly rely on COVID-19 related deaths as a measure of exposure to the pandemic. We also obtain data on daily COVID-19 hospitalizations and deaths by race and ethnicity at the state-level from the Center for Disease Control and Prevention.

Community resilience. One of the most important COVID-19 related control variables used in our empirical analysis is the ability of counties to cope with the pandemic. This variable comes from the United States Census Bureau's Community Resilience Estimates. These estimates measure the capacity of individuals and households to absorb, endure, and recover from the health, social, and economic impacts of a disaster such as a hurricane or a pandemic. For each county the population living under each of 11 risk factors is estimated and these factors are aggregated into 3 composite risk factors: (i) population with 0 risk factors; (ii) population with 1-2 risk factors, and (iii) population with 3 or more risk factors. These risk factors are based on households' and individuals' socio-economic and health conditions. Risk factors include: low Income-to-Poverty Ratio, single or zero caregiver household, unit-level crowding defined as > 0.75 persons per room, communication barriers (defined as either limited English-speaking households or no one in the household over the age of 16 with a high school diploma), no one in the household is employed full-time, disability seriously constraining a significant life activity, no health insurance coverage, being aged 65 years or older, household without a vehicle and household without broadband Internet access. For our analysis we look at populations within each county that are classified as living under 3 or more risk factors.

Lockdown stringency. We use data from the Oxford COVID-19 Government Response Tracker (Hale et al., 2020) to measure the restrictiveness of the government's pandemic policy. Use of this data is inspired by recent work which shows that stringent policies lead to lower mortality, mobility and consequently spread of infection during the pandemic (Jinjarak et al., 2020; Askitas et al., 2021). This data provides four key indices (i) an overall government response index, (ii) a containment and health index combining containing measures and health messaging, (iii) an economic support index, and (iv) a stringency index which captures the strictness of lockdown-style policies. Each of these indices reports values between 1 and 100 and varies across states and weeks. We focus on the last index.

A.3.2 Social Media and Protest Data

New and old Twitter users. In March of 2021, we collected through the Twitter Academic Research API a random sample of 3 million English-language tweets posted between May 4 and May 24 2020, by searching for tweets containing the 100 most common words in English at random instants in the sampling period. In addition, through a similar procedure, we collected one million tweets from 765,000 users containing the word "the" during random intervals in the week of December 1-7 2019. Each observation (each tweet) contains the text and timestamp of the tweet, the name and user identifier of the profile, the creation date of the profile and in some cases the location of the Twitter profile (more on geo-location below). Old users are defined as those Twitter profiles from the 2020 random sample of tweets that have created their profile before December 31st 2019. New users are defined as those that created their profile between January 21st 2020 (after the outbreak of the pandemic) and before May 25th 2020 (before the murder of George Floyd). Pre-existing Twitter users in 2019 are defined as the users seen in the December 2019 sample. In order to replicate the Müller and Schwarz (2023) instrument, in November 2021, we also collected the locations and profile creation dates of all of 639,915 followers of the South by South-West Festival @SXSW Twitter account.

Twitter in the three weeks following George Floyd. We collected tweets related to BLM using the Twitter API. In particular, we collected the universe of tweets that contain the keywords

“BLM”, “Black Lives Matter”, “Black Life Matters” or “George Floyd”,²⁵, between May 25 and June 14. For each tweet, we extract the time and text of the tweet, the user identifier, the user’s stated location, and account creation date. From this, we construct a measure for online protest which is the logarithm of (one plus) the number of geo-localized BLM related tweets in the three weeks following the murder of George Floyd. Based on this data, we also construct the (log) new BLM Twitter users, which are the Twitter accounts that end up tweeting about BLM in the three weeks following the murder of George Floyd but have created their account between January 21st and May 24th of 2020.

BLM account followers. As an additional outcome, we use the number of all followers of the official BLM account @Blklivesmatter. We collected the followers and their geolocation in February 2022. This gap between the period of analysis and the date of data collection can lead to measurement error because we do not know the starting date of following. Accounts that followed the official BLM account may stop following it and accounts that are computed as followers may start following just a few hours before the collection. Similarly, geolocation of accounts may have changed between the period of study and the date of data collection.

Geo-location of tweets and accounts. We follow the literature in assigning the location of a tweet or a user by extracting information on their self-reported location from their Twitter profile (Enikolopov et al., 2020; Takhteyev et al., 2012; Müller and Schwarz, 2023). Not all users report a location and among those who do, not all state a valid location (e.g., “in the heart of Justin Bieber”) so we restrict the sample to the users that state a valid location that can be matched to a USA county (in particular, we exclude users whose location only mentions a state). The location is an arbitrary text field which is not meant to be machine-readable. We use the Nominatim geocoding engine (based on the Open Street Map database) to find the coordinates of the most likely match for the location. We then filter out all locations outside the US and all locations that are too vague (i.e. that map the whole country or a whole state). Finally, we map these coordinates to counties using the US Census Bureau cartographic boundary files. This approach has clear limitations as it relies only on self-reported locations and may not be representative of the whole Twitter universe. Reassuringly, counties without localizable tweets only form a tiny minority: out of the 3107 counties in our universe, only 47 (1.5%) are not attributed any tweet.

Google mobility. We use data on mobility, collected through mobile phones that use Google apps (such as Google Maps). This data collects information on the time a person spent on certain mobility tasks like the time spent in parks, being at home, doing groceries, in the transit stations and finally at their workplace (as identified by Google). This information is then aggregated at the county level to measure the aggregate daily mobility. It is constructed relative to the average mobility on the same days of the week in the January 3 - February 6 2020 period. We use this data for the period between March 1st 2020 and May 24th 2020.

Google searches. We also use the Google Trends data to analyze patterns of search activity before the death of George Floyd. Each variable is a normalized index of search activity for a given search term. The indices are specified on a Nielsen’s Designated Market Area (DMA) level. A DMA is a region of the United States that consists of counties and ZIP-codes. There are 210 DMA regions covering the US. Search activity is averaged across the period of interest: each observation is a number of the searches of the given term divided by the total searches of the geography and time range, which is then normalized between regions such that the region with the largest measure is set to 100. The important limitation of the Google Trends data is that an index of search activity is an integer from zero to one hundred with an unreported privacy threshold. We use the search terms BLM and Black Lives for the three weeks leading up to the murder of George Floyd.

BLM protest from Elephrame. Elephrame is a crowd-sourced platform that collects data on Black Lives Matter and other protests. It provides information on the place and date of each BLM protest and estimated number of participants, as well as a link to a news article covering the protest. The observation period starts with the first BLM demonstration for Eric Garner on 7/19/2014 and

²⁵These keywords are considered both in when appearing separated with space, or without spaces as a hashtag (e.g. #BlackLivesMatter)

consists of any public demonstration or public art installation focused on “communicating the value of a Black individual or Black people as a whole”. Each observation is manually collected from sources that include press, protest organizers, participants and observers. We extracted all protest records from August 2014 to September 2020 and geo-coded their location.

BLM and other protest from ACLED US Crisis Monitor. The Armed Conflict Location & Event Data Project (ACLED) dataset compiles instances of political violence and demonstrations across the United States, particularly in response to significant societal and political tensions such as those following the contentious 2020 general election, incidents of police violence, and large-scale protests like those related to the Black Lives Matter movement. The dataset’s assembly involves the aggregation and coding of events based on a diverse array of sources, including over 2800 national, regional, and local media outlets, to ensure comprehensive coverage and mitigate biases inherent in any single source. Major variables in the ACLED dataset encompass detailed codings of event types (e.g., protests, violence against civilians, strategic developments), locations (with specificity down to the city level and, in certain cases, sub-city regions in major cities like New York and Los Angeles), dates, and actors involved, including both state forces and non-state actors like militias and protest groups. The dataset distinguishes between various security forces (e.g., police, national guard) and non-state actors, including armed groups and militias, with a nuanced classification system that accounts for the complexities of political violence and protest actions within the U.S. context. ACLED reports 21,582 protest in the United States in 2020. 40% of the protests in the data set are linked to the pandemic.

George Floyd Street Art. We extract information on the location of street art representing or referring to George Floyd from the Urban Art Mapping George Floyd and Anti-Racist Street Art database. The crowd-sourced website run by researchers from the University of St. Thomas documents street art from around the world created in the aftermath of the murder of George Floyd. Their archive is a repository of images made available for research and education. The website contains geo-tagged information and images of George Floyd related street art, which we match to counties. The data does not contain time stamps and has no information on when these images were added. For this reason, we can only interpret the street art as cross-sectional snapshots at time of accessing the website in January of 2022. Overall, we record 2183 images across 70 counties. Most of the images (1467) are recorded in Minneapolis.

A.3.3 Survey Data and other County Controls

The American Trends Panel survey by Pew Research Center. To zoom in and move from county to individual level analysis, we employ the American Trends Panel survey (ATP), conducted by Pew Research Center. The panel is based on a representative sample of U.S. adults who participate via self-administrated online web-survey. Participants with no internet access were provided with tablets and wireless connection to answer the survey, which is crucial for studying the effects of social media. For our analysis we draw from the panel wave 68 that took place from June 4 to June 10, 2020. Participants were questioned on a wide range of topics, including the Black Lives Matter movement, police brutality, ideologies and politicians, race relations, social issues, the coronavirus and president Donald Trump. The survey also contains a group of demographic variables, which are included in our analysis as controls: race, age, gender, education, income, political leanings, level of urbanisation of participants’ region. It is important to note that the participants’ region of residence is anonymised, therefore the exact data on COVID-19 cases and deaths is not available. However, the panel does include a categorical version of this data: whether the prevalence of COVID-19 in the respondents’ region is low, medium or high. We can make only associative conclusions based on this limited information.

The Cooperative Election Study. The 2020 Cooperative Election Study (CES), spearheaded by a collaboration of 60 research teams, offers an expansive dataset from a survey administered to 61,000 participants across the United States. This study, facilitated by YouGov, was designed to capture a wide array of information regarding Americans’ voting behaviors, political attitudes, and electoral experiences in the context of the 2020 elections. The CES was conducted in two waves to capture both pre-election and post-election sentiments among U.S. voters. The pre-election wave of the survey took place from September 29 to November 2, 2020, encompassing the crucial final weeks leading up to the election. Following the election, the post-election wave was conducted from November 8 to December 14, 2020,

allowing for the collection of respondents' reactions to the election outcomes, their voting experiences, and their perspectives on the electoral process. The survey has information on the location of respondents at the ZIP code level which we then merge to counties in our data set.

Use of deadly force by police. We obtain this from the collaborative platform Fatal Encounters. This data is collected by a multi-disciplinary team at the University of Southern California. The results are published as part of the *National Officer-Involved Homicide Database*. The data is available from 2000 onward and contains the name, gender, race, and age of each victim and the specific address where the death occurred, among other variables.

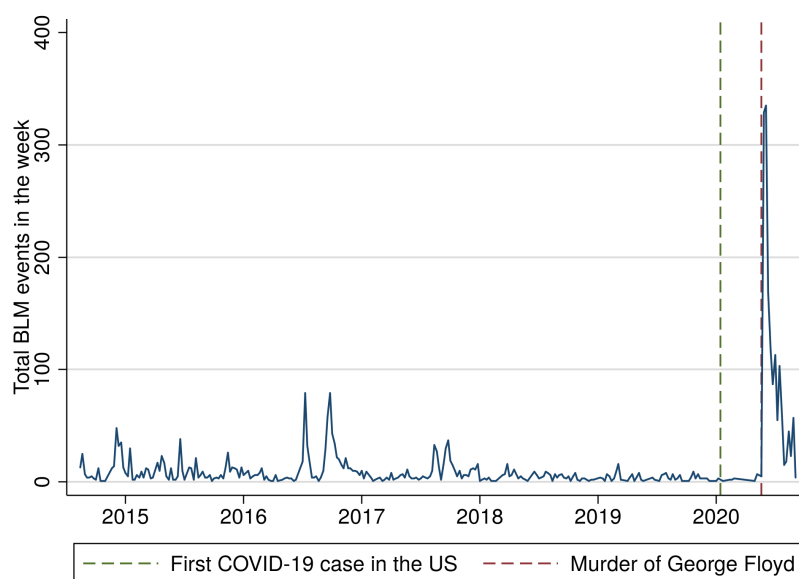
Additional county-level controls. We include unemployment data available on a monthly basis at the county level from the Local Area Unemployment Statistics of the US Bureau of Labor Statistics, which we average over the 12 months preceding the death of George Floyd (May 2019 to April 2020). We also use the total population, population by ethnicity, and income statistics such as Black poverty rate and median household income (all in 2018), as well as past Republican vote share (in 2012 and 2016) from the American Community Survey. We use a dummy for urban counties which is constructed from the Office of Management and Budget's February 2013 delineation of metropolitan and micropolitan statistical areas.²⁶ The measure of social capital that we use aggregates the information on the number of local organizations.²⁷

²⁶2013 NCHS Urban-Rural Classification Scheme for Counties, Vintage 2012 postcensal estimates of the resident U.S. population. NCHS Urbanization levels are designed to be convenient for studying the difference in health across urban and rural areas. This classification has 6 categories: large "center" metropolitan area (*inner cities*), large "fringe" metropolitan area (*suburbs*), median metropolitan area, small metropolitan area, micropolitan area and non-core (non-metropolitan counties that are not in a micropolitan area).

²⁷This includes: (a) civic organizations; (b) bowling centers; (c) golf clubs; (d) fitness centers; (e) sports organizations; (f) religious organizations; (g) political organizations; (h) labor organizations; (i) business organizations; and (j) professional organizations.

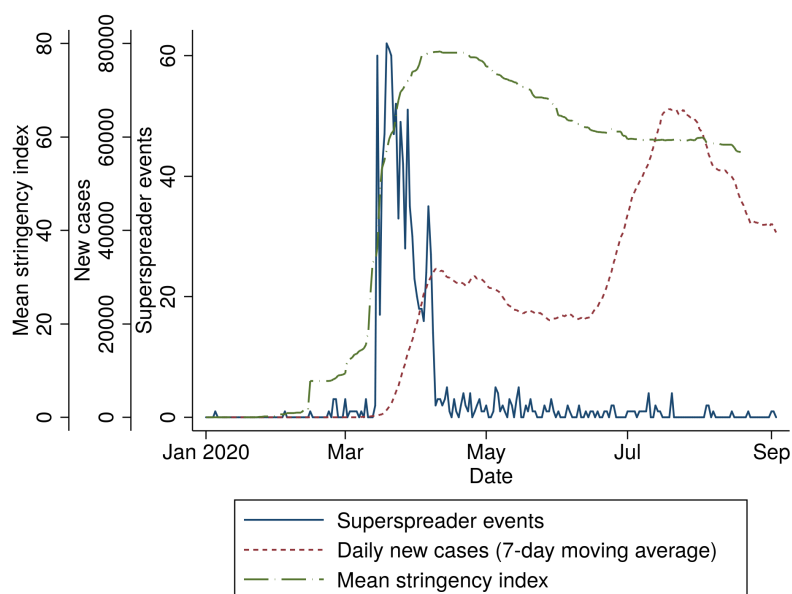
A.4 Additional Tables and Figures

Figure A1: BLM events over time



Note: Number of BLM events per week in the US from August 2014 to September 2020. The green vertical line denotes the week of the first confirmed COVID-19 case in the US (January 21, 2020), and the red vertical line denotes the week of the murder of George Floyd (May 25, 2020).

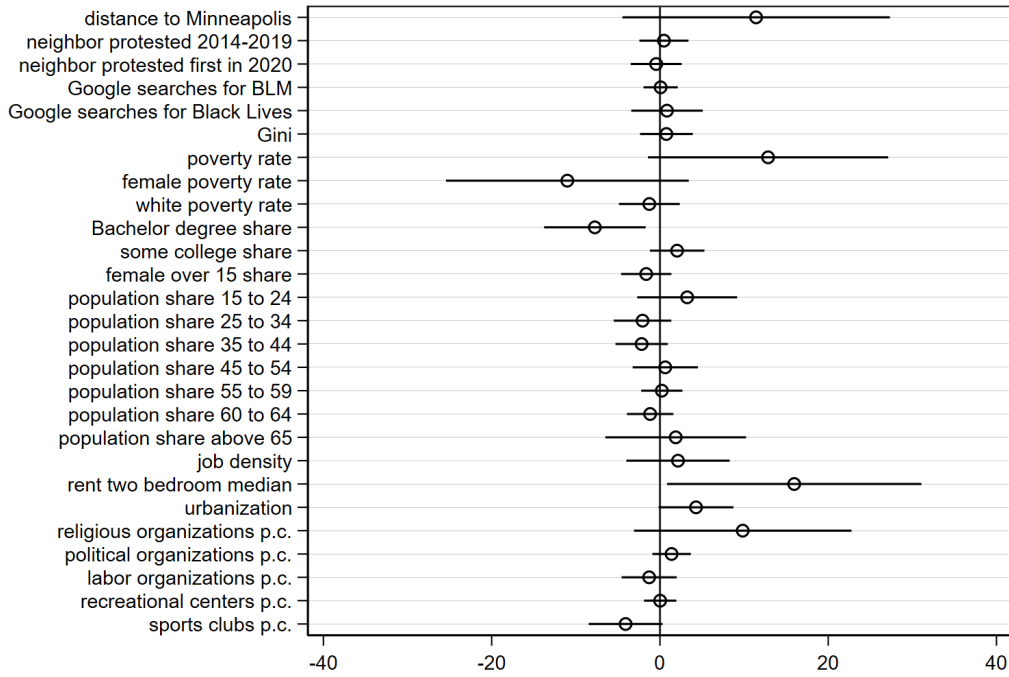
Figure A2: Super spreading events arise in early stages of the pandemic when lockdown stringency is low but COVID-19 prevalence is high enough



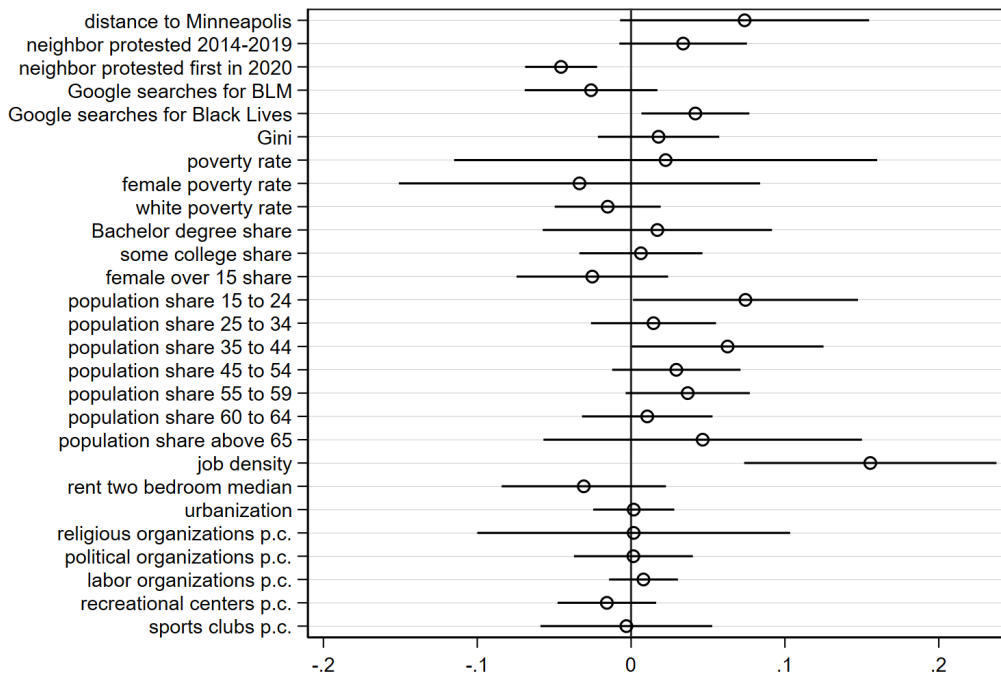
Note: Plot showing the evolution of the number of SSE events, social distancing measures and new COVID-19 cases in the US. Solid (blue) line represents the number of daily total SSEs over time (January 2020 to September 2020). Dashed (green) line shows the daily average lockdown stringency index across all US states, as measured by the Oxford COVID-19 response tracker. Dotted (red) line shows the number of daily new COVID-19 cases as recorded by the New York Times.

Figure A3: **Plausibility of quasi random exposure to instrument components**

(a) Dep. var.: SSE_c (with state fixed effects)

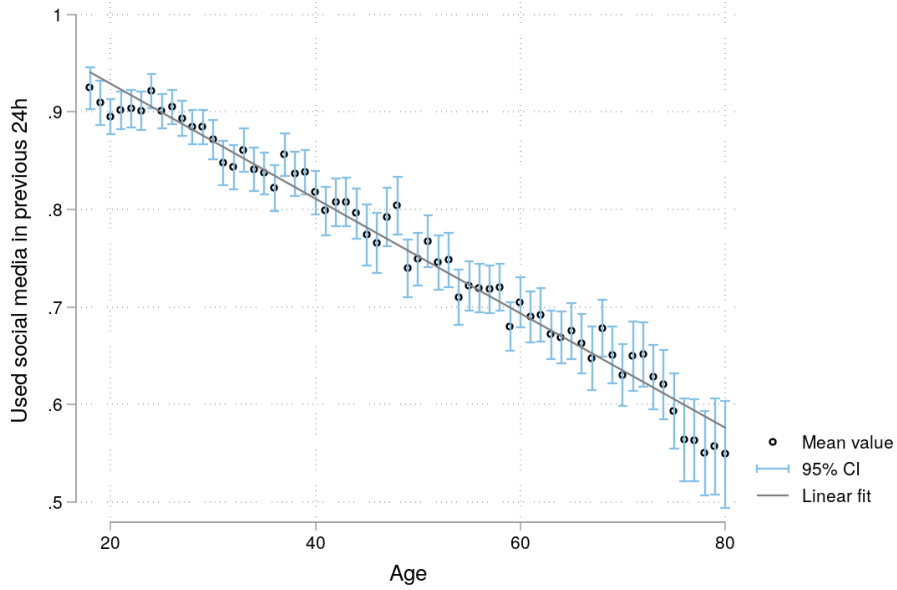


(b) Dep. var.: $SXSW_c$ (with state fixed effects)



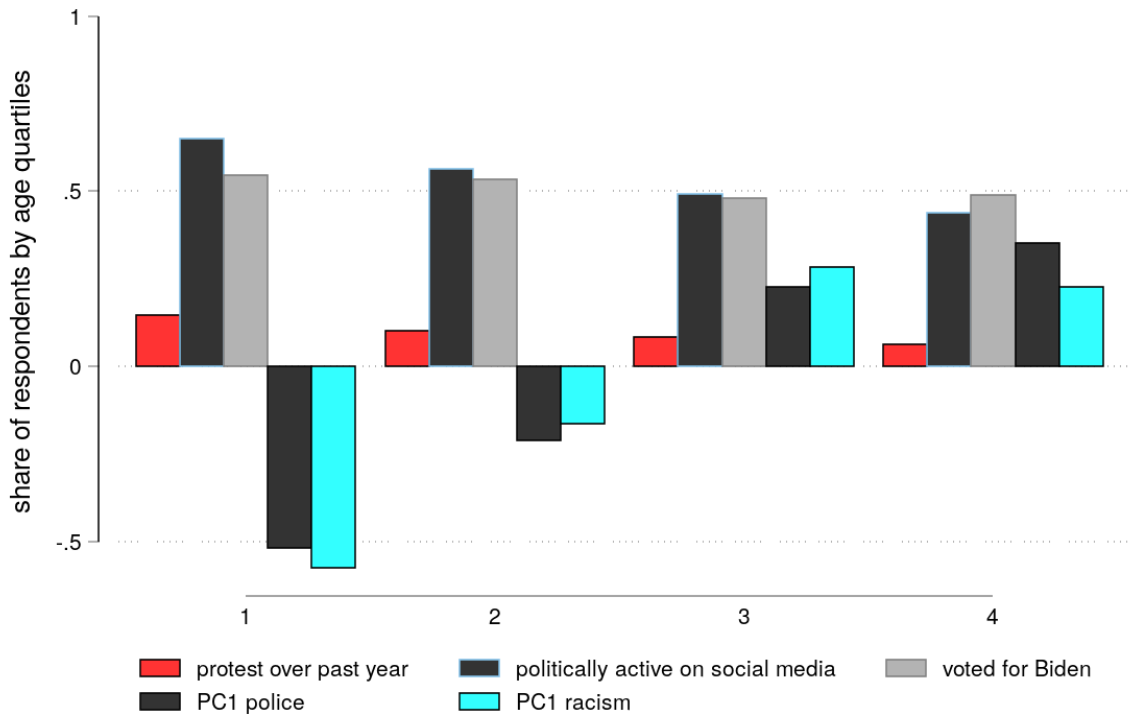
Note: Plausibility of the instrument exogeneity, for each individual component. We show point estimates with 95% confidence intervals, clustered at the state level. We run the baseline specification but use our instrument components as outcomes. Top panel shows as the outcome SSE_c - the cumulative number of SSEs within 200km radius from the county border up to six weeks before the murder of Floyd. The bottom panel shows as the outcome $SXSW_c$, which is the logarithm of one plus the number of SXSXW followers who created their account in March 2007 in the county and neighboring counties. Both models include state fixed effects. We describe all variables in detail in Table A2.

Figure A4: Social media use by age from CES



Note: Graph of the probability of having used social media in the past 24 hours by age, among the respondents of the Cooperative Election Study (CES) conducted in November 2020.

Figure A5: Political preferences by age quartile from CES



Note: Bar graph showing share of respondents by age quartiles for the non-placebo outcomes used in Table 4, using information from the Cooperative Election Study (CES) conducted in November 2020. The age thresholds are 0-33, 34-49, 50-63 and older than 63.

Table A1: Example Tweets

| Date | Text |
|--------------|--|
| May 29, 2020 | While #BlackLivesMatter is raising awareness on Twitter, it shouldn't stop there. While you're inside with your families, talk about racism and discrimination. Especially with older generations who don't use social media and don't see further than the national new's portrayal. |
| May 30, 2020 | This is called UNITY. this is what white america doesn't want. they're afraid of the non racist whites to form patternship and unity with POC bc then they will be out numbered. I stand by my brothers #BlackLivesMatter https://t.co/EPYE9HKkBN |
| May 31, 2020 | Reach out to black friends, peers, and social media connections to LISTEN to them with the understanding that I do not know what their struggles are like as a person that has lived with privilege. #BlackLivesMatter |
| Jun 2, 2020 | If it weren't for Twitter and social media the videos of George Floyd and Ahmed Arbery would have not been seen and murderers would have walked free. Fact. #BlackLivesMatter |
| Jun 4, 2020 | (3/7) We will also be sharing courses made by the Arist community designed to educate allies. The first example: https://bit.ly/anti-racism101 This 20-day text message course will teach you about systemic racism against Black people and how you can practice anti-racist allyship.#BlackLivesMatter |
| Jun 4, 2020 | I made a decision when I came on twitter to keep it strictly for work. I have other social media for expressing personal and political views. However, given the events of the last week, I feel compelled to say something - so here is my bit #BlackLivesMatter #WhitePrivilege |
| Jun 6, 2020 | #IAmASuburbanMom and Black Lives Matter to me! I just went to a rally in a suburb of Atlanta, and there are a lot of us moms who want racial justice and change! |
| Jun 7, 2020 | White privlage means you CAN walk away from #BlackLivesMatter when you get weary and you go back to your regular routine. Our black and coloured allies don't have that privilege to simply walk away. It's their life. Recognizing our white privilege means refusing to walk away. 3/3 |
| Jun 11, 2020 | 1/ I've been trying to learn more about all the complexities of everything going on lately, and how to be a better ally, better support the #blacklivesmatter movement & simply be an anti-racist. For what it's worth, here's a few things I've found to be especially helpful: |
| Jun 13, 2020 | There was a #BlackLivesMatter car parade in my VERY white, VERY red suburban San Antonio neighborhood today. I was afraid we'd be the only car. There were 50 of us!!! |

Note: Selected examples of tweets posted after the murder of George Floyd showing an increase awareness and change in attitudes about BLM and racism.

Table A2: Sources of the variables used in the analysis

| Name | Exact Definition | Geographic Unit | Time Frame | Source |
|-----------------------------------|--|-----------------|--|--|
| COVID cases and deaths | Cumulative cases and deaths attributed to COVID-19 | County | Daily, January 21 to August 20, 2020 | NYTimes, https://github.com/nytimes/covid-19-data |
| COVID deaths by race | Cumulative deaths attributed to COVID-19, by race | State | Daily, April 12 to October 14, 2020 (depending on state) | The COVID Tracking Project, The Atlantic, https://covidtracking.com/race |
| BLM protests | Number of BLM protests on each day and estimated number of participants | County | August 2014 to August 20, 2020 | Elephrame |
| Superspreader events | Timing and location of known superspreader events | County | January to August 2020 | Swinkels, K. (2020). SARS-CoV-2 Superspreading Events Around the World [Google Sheet]. Retrieved from www.superspreadingdatabase.com |
| New Twitter users | Number of distinct users whose account was created after January 21, 2020 from the county in a random sample of 3 million English language tweets posted between May 4 and May 24, 2020. Geolocation based on the location indicated in the user's profile. | County | May 4 - May 24, 2020 | Twitter API |
| Old Twitter users | Number of distinct users whose account was created before January 21, 2020 from the county in a random sample of 3 million English language tweets posted between May 4 and May 24, 2020. Geolocation based on the location indicated in the user's profile. | County | May 4 - May 24, 2020 | Twitter API |
| Preexisting Twitter users in 2019 | Number of distinct users appearing in a random sample of 1 million English language tweets posted between December 1 and December 7 2020. Geolocation based on the location indicated in the user's profile. | County | December 1 - 7, 2019 | Twitter API |
| Time spent at home | Time spent at home, compared to pre-pandemic baseline, in the week before the murder of George Floyd | County | May 18 - May 24, 2020 | Google Community Mobility Report |
| Demographic controls | female over 15 share, population share by age, Black population share, Bachelor degree share, college education share. | County | 2018 | American Community Survey 5-year estimates |
| Economic controls | Gini index, poverty rate, female poverty rate, poverty rate by race | County | 2018 | American Community Survey 5-year estimates. |
| Economic controls | median household income, job density (2013), median rent for two bedroom (2015), | County | 2016 | Census Bureau Opportunity Atlas |
| Tweets BLM by late adopters | Number of tweets mentioning keywords related to Black Lives Matter in the 3 weeks following the murder of George Floyd, posted by accounts created between January 21 and May 24, 2020. Geolocation based on the location indicated in the user's profile. | County | May 25 - June 14, 2020 | Twitter API |

Continued on next page

| Name | Exact Definition | Geographic Unit | Time Frame | Source |
|--|---|-----------------|----------------------------|---|
| Late adopters followers @BLM | Number of users geolocated in the county and following the @Blklivesmatter Twitter account that are considered late adopters (account created between January 21 and May 24, 2020). Geolocation based on the location indicated in the user's profile. | County | | Twitter API |
| Cooperative Election Study 2020 survey results | See Table A9 for exact variables | County | 2020 | CCES |
| PEW June 2020 survey results | See notes of Table 5 | County | June 2020 | Pew Research Center American Trends Panel - Wave 68 |
| George Floyd street art | Number of pieces of street art related to the murder of George Floyd. Geolocated according to the map embedded in the web page. | County | | Urban Art Mapping: George Floyd and Anti-Racist Street Art database |
| All 2020 BLM protest | Number of entries in the ACLED US Crisis Monitor data whose associated actor or notes contain "BLM", and occurring after the murder of George Floyd. | County | May 25 - December 31, 2020 | ACLED US Crisis Monitor |
| All 2020 other protest | Number of entries in the ACLED US Crisis Monitor data whose associated actor or notes does not contain "BLM", and occurring after the murder of George Floyd. | County | May 25 - December 31, 2020 | ACLED US Crisis Monitor |
| Populist and pro-Trump protests | Number of entries in the ACLED US Crisis Monitor data occurring after the murder of George Floyd, and whose associated actors or notes contain one of the following: Stop the Steal, Proud Boys, QAnon, election fraud, ballot, pro-Trump, Trump supp, in support of Trump, WAF, or MAGA. | County | May 25 - December 31, 2020 | ACLED US Crisis Monitor |
| Anti-mask protests | Number of entries in the ACLED US Crisis Monitor data occurring after the murder of George Floyd, and whose associated actors or notes contain one of the following: pandemic, mask mandate, social distancing, coronavirus, public health order. | County | May 25 - December 31, 2020 | ACLED US Crisis Monitor |
| All Lives Matter tweets | Tweets containing "All Lives Matter" or #AllLivesMatter, posted in the 3 weeks following the murder of George Floyd. Geolocation based on the location indicated in the poster's profile. | County | May 25 - June 14, 2020 | Twitter API |
| Blue Lives Matter tweets | Tweets containing "Blue Lives Matter" or #BlueLivesMatter, posted in the 3 weeks following the murder of George Floyd. Geolocation based on the location indicated in the poster's profile. | County | May 25 - June 14, 2020 | Twitter API |
| Distance to Minneapolis | Geographical distance, in kilometers, from the center of the county to the center of Minneapolis. | County | | Coordinates: Census Bureau |

Continued on next page

| Name | Exact Definition | Geographic Unit | Time Frame | Source |
|--|--|------------------------|--------------------------|--|
| Google searches for BLM/Black Lives Matter | Index indicating the importance of searches for "BLM"/"Black Lives Matter" among all Google searches, averaged in the 3 weeks prior to the murder of George Floyd. | Designated Market Area | May 25 - June 15, 2020 | Google Trends API |
| Lockdown stringency | Measure of the strength of the government response, on May 24. | State | January - August 2020 | OxCGRT government response index |
| SXSW users, Pre-SXSW users | Followers of the @SXSW Twitter account that joined the network, respectively, in March 2007, and before March 2007. Geolocation based on the location indicated in the user's profile. | County | 2007 | Twitter API |
| Police-caused deaths, 2014-2019 | Number of Black people who died during an encounter with the police between 2014 and 2019 | County | 2014-2019 | Fatal Encounters |
| Police-caused deaths, 2020 | Number of Black people who died during an encounter with the police, in 2020 before May 25th | County | January 1 - May 24, 2020 | Fatal Encounters |
| Unemployment, 2019-2020 | Mean unemployment in the 12 months from May 2019 to April 2020 | County | May 2019 - April, 2020 | Bureau of Labor Statistics: Local Area Unemployment Statistics |
| Urbanization level | From 1 = urban to 6 = rural | County | 2013 | NCHS Urban-Rural Classification Scheme for Counties |
| Pandemic resilience index | Measure of the capacity of individuals and households to absorb, endure, and recover from the health, social, and economic impacts of a disaster such as a hurricane or a pandemic. For each county the population living under each of 11 risk factors is estimated and these factors are aggregated. This variable measures the share of the population with 3 risk factors or more. | County | 2018 | Census Bureau Community Resilience Estimates |
| Republican vote in 2012 | Share of the vote in the 2012 presidential election that was a vote for Mitt Romney. | County | 2012 | MIT Election Data and Science Lab (2018) |
| Republican vote in 2016 | Share of the vote in the 2016 presidential election that was a vote for Donald Trump. | County | 2016 | MIT Election Data and Science Lab (2018) |
| Social capital | Total number of social organizations (religious, cultural, political, etc) divided by population | County | 2014 | Rupasingha et al. (2006) |
| Organizations per capita | Number of organizations of the given type per capita in the county (religious, political, labor, recreational, sports) | County | 2014 | Rupasingha et al. (2006) |
| Propensity to protest | Propensity to have a BLM event after a notable death of a Black person caused by the police, estimated by LASSO on past data (see B) | County | 2014-2020 | Own construction |

Table A3: Summary statistics - counties without prior BLM event

| From 25th of May to 14th of June 2020: | N | Mean | SD | Min | Max |
|--|------|-----------|-----------|-----------|------------|
| Presence of BLM events | 2768 | 0.048 | 0.213 | 0.000 | 1.000 |
| Number of BLM events | 2768 | 0.064 | 0.322 | 0.000 | 5.000 |
| Total participants BLM events | 2768 | 21.026 | 172.090 | 0.000 | 5500.000 |
| log(1+ total participants BLM events) | 2768 | 0.234 | 1.140 | 0.000 | 8.613 |
| Average number participants given protest | 132 | 440.913 | 662.560 | 0.000 | 5500.000 |
| Tweets about BLM | 2768 | 74.816 | 1021.834 | 0.000 | 46141.000 |
| Tweets about BLM by new users | 2768 | 2.194 | 31.502 | 0.000 | 1482.000 |
| log(1 + tweets about BLM by new users) | 2768 | 0.201 | 0.723 | 0.000 | 7.302 |
| Tweets mentioning #AllLivesMatter | 2768 | 47.488 | 326.063 | 0.000 | 15659.000 |
| log(1 + tweets mentioning #AllLivesMatter) | 2768 | 2.581 | 1.479 | 0.000 | 9.659 |
| Tweets mentioning #BlueLivesMatter | 2768 | 6.125 | 36.206 | 0.000 | 1647.000 |
| log(1 + tweets mentioning #BlueLivesMatter) | 2768 | 1.024 | 1.114 | 0.000 | 7.407 |
| On the 25th of May 2020: | | | | | |
| COVID deaths (per 1000) | 2768 | 0.099 | 0.230 | 0.000 | 2.935 |
| Superspreader events, 6+ weeks ago, in county | 2768 | 0.148 | 0.748 | 0.000 | 20.000 |
| Superspreader events, 6+ weeks ago, 200km around county | 2768 | 21.344 | 34.380 | 0.000 | 287.000 |
| Push-pull instrument $SSE_{-c} \times SXS_{W_c}$ | 2768 | 5.921 | 27.588 | 0.000 | 390.338 |
| New Twitter users | 2768 | 1.238 | 5.671 | 0.000 | 244.000 |
| log(1 + new Twitter users) | 2768 | 0.445 | 0.666 | 0.000 | 5.501 |
| Later outcomes: | | | | | |
| Followers of @BlkLivesMatter created during pandemic | 2768 | 0.441 | 2.167 | 0.000 | 78.000 |
| log(1 + followers of @BlkLivesMatter created during pandemic) | 2768 | 0.200 | 0.444 | 0.000 | 4.369 |
| George Floyd Street Art count (UAM project) | 2768 | 0.009 | 0.206 | 0.000 | 10.000 |
| All 2020 BLM protest (ACLEd) | 2768 | 1.314 | 3.147 | 0.000 | 41.000 |
| Other 2020 protest (ACLEd) | 2768 | 1.811 | 3.388 | 0.000 | 56.000 |
| Anti mask and distancing protests (ACLEd) | 2768 | 0.365 | 1.424 | 0.000 | 19.000 |
| Populist and pro-Trump protests (ACLEd) | 2768 | 0.055 | 0.316 | 0.000 | 7.000 |
| County characteristics: | | | | | |
| Baseline Twitter users | 2768 | 31.361 | 177.499 | 0.000 | 8269.000 |
| log(1 + old Twitter users) | 2768 | 2.378 | 1.322 | 0.000 | 9.020 |
| Black police-related deaths (Fatal Encounters, 2014 - 2019) | 2768 | 0.207 | 0.724 | 0.000 | 15.000 |
| Black police-related deaths (Fatal Encounters, 2020) | 2768 | 0.014 | 0.131 | 0.000 | 3.000 |
| Unemployment rate (Bureau of Labor Statistics, 2019 - 2020) | 2768 | 4.713 | 1.575 | 0.708 | 17.442 |
| Black population share (ACS, 2018) | 2768 | 0.093 | 0.146 | 0.000 | 0.875 |
| Large cities | 2768 | 0.001 | 0.027 | 0.000 | 1.000 |
| Black poverty rate A(CS, 2018) | 2768 | 0.283 | 0.236 | 0.000 | 1.000 |
| Pandemic resilience index (Census Bureau, 2018) | 2768 | 25.957 | 5.066 | 10.684 | 48.444 |
| Vote share for republicans (MIT Election Data Science Lab, 2016) | 2768 | 0.656 | 0.141 | 0.083 | 0.960 |
| Vote share for republicans (MIT Election Data Science Lab, 2012) | 2768 | 0.614 | 0.140 | 0.060 | 0.959 |
| Median household income (Opportunity Atlas, 2016) | 2768 | 47521.697 | 12362.349 | 20170.891 | 129150.344 |
| Social capital (civil society organizations per 1000, 2014) | 2768 | 1.426 | 0.726 | 0.000 | 6.887 |
| Log(SXSW followers created before March 2017) | 2768 | 0.090 | 0.228 | 0.000 | 1.427 |
| Log(SXSW followers created during March 2017) | 2768 | 0.157 | 0.312 | 0.000 | 1.658 |

Note: Descriptive statistics of all variables used in the main results for the counties that did not host a BLM protest before the murder of George Floyd. Different panels correspond to different periods where data is measured.

Table A4: Summary statistics for super spreading events by their type

| Location of SSE event | Total events | Total Events 6 weeks before GF's murder | Total Cases |
|--------------------------|--------------|--|-------------|
| Community | 11 | 8 | 504 |
| Development Center | 12 | 12 | 1612 |
| Event/group gathering | 12 | 12 | 450 |
| Industry | 113 | 83 | 17660 |
| Medical | 137 | 130 | 13684 |
| Nursing Home | 269 | 255 | 26527 |
| Prison | 186 | 182 | 44414 |
| Rehabilitation / Medical | 258 | 246 | 26776 |
| Restaurant/Bar | 5 | 4 | 1164 |
| Retail | 3 | 0 | 46 |
| School | 2 | 2 | 173 |
| Other | 15 | 13 | 1337 |

Note: Descriptive statistics on super spreading events (SSE) by the place where they took place. First column describe the location, second column the total number of events of this category, third column the total number of event in each location during our period of consideration (up until 6 week before the murder of George Floyd), and column 4 the total number of COVID-19 cases registered.

Table A5: **SSEs predict COVID-19 deaths, cases, residential stay and new Twitter users; SXSW predicts baseline network**

| | COVID-19 cases/1000 (1) | COVID-19 deaths/1000 (2) | Residential stay (3) | Log new Twitter users (4) | Log new Twitter users (5) | Log baseline Twitter users (6) |
|--|-------------------------------|--------------------------------|----------------------------|---------------------------------|---------------------------------|--------------------------------------|
| SSE | 0.0189*** (0.00612) | 0.000901*** (0.000272) | 0.00563** (0.00231) | 0.00197*** (0.000685) | -0.00275*** (0.000822) | |
| SSE × Log baseline Twitter users (Dec. 2019) | | | | | 0.00121*** (0.000250) | |
| SXSW | | | | | | 0.373*** (0.103) |
| Observations | 2,767 | 2,767 | 1,022 | 2,767 | 2,767 | 2,767 |
| R-squared | 0.136 | 0.282 | 0.781 | 0.238 | 0.519 | 0.401 |
| Mean of dependent variable | 2.590 | 0.0990 | 10.01 | 0.445 | 0.445 | 1.738 |
| Pandemic exposure | No | No | No | No | No | No |
| Baseline Twitter | No | No | No | No | No | No |
| Economic controls | Yes | Yes | Yes | Yes | Yes | Yes |
| BLM controls | Yes | Yes | Yes | Yes | Yes | Yes |
| Political controls | Yes | Yes | Yes | Yes | Yes | Yes |
| State FE | Yes | Yes | Yes | Yes | Yes | Yes |

Note: Estimation of the effects of SSE_c (or, in column 6, $SXSW_c$) on COVID and Twitter outcomes. Outcomes in column 1 and 2 are cumulative COVID-19 related cases and deaths per 1000 inhabitants before May 25th 2020. Outcome in column 3 is based on the Google mobility index in the week of 18-24 May 2020, which measures how much time individuals spend at home relative to January 3 - February 6 2020, rather than at work, at places of commerce or in parks and recreational facilities. Outcomes in columns 4 and 5 are measured as the log of one plus new Twitter users that have created their account between January and May 2020 based on a random sample of geolocalized English-language tweets collected between May 4 and May 24 2020. Outcome in column 6 is measured as the log of one plus the number of Twitter users created before January 2020. SSE is measured as the cumulative number of SSEs within 200 km of the county border but not within the county until six weeks before the murder of George Floyd, i.e. until mid-April 2020. SXSW is measured as the log of one plus the number of followers of the official SXSW account who created their profile in March of 2007. All models include state fixed effects, as well as the full set of non-COVID, non-Twitter related county controls: the share of Black population, urban, median household income, unemployment share, Black poverty rate, 3+ risk factors/community resilience, Republican vote share in 2012 and 2016, social capital (number of civic organizations per capita) and deadly force used by police against Black people. Standard errors (in parentheses) are clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A6: **First Stage: Predicting late adoption with super spreader events (SSE_{-c})**
× South-by-South West ($SXSW_c$)

| | Log(1+ New Twitter Accounts) | | | | |
|------------------------|------------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | baseline | | | | |
| | (1) | (2) | (3) | (4) | (5) |
| Radius SSE | 200 km | 100 km | 150 km | 300 km | 400 km |
| Time lag SSE: 42 days | | | | | |
| SSE × SXSW | 0.00234*** (0.000460) | 0.00386*** (0.00114) | 0.00252*** (0.000777) | 0.00157*** (0.000409) | 0.00103** (0.000491) |
| SSE | 0.000122 (0.000575) | -0.000650 (0.00127) | 0.000292 (0.000933) | -0.000160 (0.000426) | -0.000383 (0.000338) |
| SXSW | 0.00315 (0.0343) | 0.0327 (0.0324) | 0.0220 (0.0347) | -0.0115 (0.0322) | -0.0134 (0.0404) |
| Kleibergen-Paap F stat | 25.89 | 11.51 | 10.54 | 14.81 | 4.401 |
| Time lag SSE | 42 days | 28 days | 35 days | 49 days | 56 days |
| Radius SSE: 200km | | | | | |
| SSE × SXSW | 0.00234*** (0.000460) | 0.00232*** (0.000442) | 0.00233*** (0.000452) | 0.00238*** (0.000480) | 0.00283*** (0.000562) |
| SSE | 0.000122 (0.000575) | 8.71e-05 (0.000572) | 0.000108 (0.000573) | 0.000200 (0.000598) | 0.000191 (0.000697) |
| SXSW | 0.00315 (0.0343) | 0.00201 (0.0346) | 0.00251 (0.0345) | 0.00730 (0.0337) | 0.00815 (0.0335) |
| Kleibergen-Paap F stat | 25.89 | 27.55 | 26.57 | 24.62 | 25.40 |
| Observations | 2767 | 2767 | 2767 | 2767 | 2767 |
| State FE | Yes | Yes | Yes | Yes | Yes |
| Pandemic exposure | Yes | Yes | Yes | Yes | Yes |
| Baseline Twitter | Yes | Yes | Yes | Yes | Yes |
| Economic controls | Yes | Yes | Yes | Yes | Yes |
| BLM controls | Yes | Yes | Yes | Yes | Yes |
| Political controls | Yes | Yes | Yes | Yes | Yes |

Note: This table shows the first stage regression for the log of 1 + the number of geolocalized Twitter profiles that were created between January 20th 2020 and May 25th 2020, based on a random sample of tweets. SXSW measures the log of 1 + the number of followers of the SXSW official account that joined Twitter during the festival in the county and in its neighboring counties. SSE are measured as the number of super spreading events within a certain radius (Panel A) and until a certain number of days before the murder of George Floyd on May 25th 2020 (Panel B). The baseline instrument is reported in column 1 and uses a radius of 200 km and 42 day lag. We report Kleibergen-Paap rkWald F statistics for each instrument variation at the bottom of each panel. All regressions include controls for COVID per capita, old Twitter users, the (log) number of SXSW Twitter followers before the Austin festival and its interaction with SSE_{-c} , as well as the share of Black population, urban, median household income, unemployment share, Black poverty rate, 3+ risk factors/community resilience, Republican vote share in 2012 and 2016, social capital (number of civic organizations per capita) and deadly force used by police against Black people. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table A7: Main 2SLS results (all coefficients reported)

| 2SLS: | At least one BLM protest | | | |
|---|--------------------------|-------------------------|-------------------------|-------------------------|
| | (1) | (2) | (3) | (4) |
| new Twitter users | 0.377*** (0.117) | 0.439*** (0.139) | 0.439*** (0.141) | 0.443*** (0.146) |
| old Twitter users | -0.0955** (0.0413) | -0.120** (0.0483) | -0.119** (0.0482) | -0.124** (0.0502) |
| COVID deaths before May 25th (per 1000) | 0.0414* (0.0235) | 0.0542** (0.0256) | 0.0549** (0.0270) | 0.0545* (0.0273) |
| SSE | -3.20e-05 (0.000493) | -0.000311 (0.000512) | -0.000311 (0.000514) | -0.000352 (0.000525) |
| SXSW | 0.00140 (0.0211) | -0.00845 (0.0207) | -0.00590 (0.0200) | -0.0104 (0.0210) |
| SXSW (before 03-2007) | 0.0215 (0.0483) | 0.00952 (0.0483) | 0.00948 (0.0484) | 0.00552 (0.0494) |
| SSE × SXSW (before 03-2007) | 0.000117 (0.00109) | 0.000316 (0.00105) | 0.000333 (0.00107) | 0.000370 (0.00107) |
| Unemployment rate (2019-2020) | | -0.0112** (0.00435) | -0.0115** (0.00442) | -0.0122*** (0.00451) |
| Median household income | | 6.01e-07 (9.07e-07) | 7.26e-07 (9.32e-07) | 6.11e-08 (1.03e-06) |
| Large cities | | -0.541* (0.279) | -0.545** (0.266) | -0.575** (0.258) |
| Pandemic resilience index | | -0.00197 (0.00195) | -0.00202 (0.00192) | -0.00253 (0.00207) |
| Police violence against Blacks (2014-2019) | | | -0.00529 (0.0125) | -0.00586 (0.0127) |
| Police violence against Blacks (2020) | | | -0.0298 (0.0325) | -0.0303 (0.0327) |
| Black population share | | | 0.00156 (0.0531) | -0.101 (0.0809) |
| Black poverty rate | | | 0.0300 (0.0183) | 0.0282 (0.0186) |
| Vote Republican (2016) | | | | -0.380* (0.208) |
| Vote Republican (2012) | | | | 0.244 (0.193) |
| Social capital | | | | 0.00305 (0.00745) |

Note: Estimation results from specification 1, with all coefficients reported. New Twitter users are measured as the log of one plus new geo-located accounts at the county level created after the beginning of the pandemic but before George Floyd's murder based on a random sample of tweets. Instrument $SSE_{-c} \times SXSW_c$ is the push-pull instrument for pandemic Twitter take-up, combining the cumulative number of SSEs outside of the county within a 200km radius SSE_{-c} with an instrument for baseline Twitter penetration $SXSW_c$ following Müller and Schwarz (2023). All estimations include state fixed effect, the cumulative number of COVID-19 related deaths until May 25th (pandemic exposure), the log transformed number of Twitter accounts created before January 2020 from a random sample of Tweets (baseline Twitter) as controls. 2SLS, reduced form and first stage additionally include SSE_{-c} and $SXSW_c$ separately. They also include number of SXSW Twitter followers before the Austin festival and its interaction with SSE_{-c} . The outcome is a dummy variable for any BLM protest in the three weeks following the murder of George Floyd, using information from *Elephrame*. Economic controls: average unemployment rate in the preceding year, median household income in 2016, pandemic resilience, a dummy for urban counties. BLM controls: share of Black population, Black poverty rate, and deadly force used by police against Black people in 2014-2019 and in 2020. Political controls: Republican vote share in 2012 and 2016 and social capital (number of civic organizations per capita). Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table A8: Principal component analysis of BLM protest

| | at least one BLM protest | number of BLM protests | log total participants | log tweets BLM (new users) | log followers @BLM (new users) |
|--------------------------------|-----------------------------|---------------------------|---------------------------|-------------------------------|-----------------------------------|
| at least one BLM protest | 1 | | | | |
| number of BLM protests | 0.541 [0.000] | 1 | | | |
| log total participants | 0.148 [0.000] | 0.631 [0.000] | 1 | | |
| log tweets BLM (new users) | 0.555 [0.000] | 0.557 [0.000] | 0.237 [0.000] | 1 | |
| log followers @BLM (new users) | 0.532 [0.000] | 0.577 [0.000] | 0.271 [0.000] | 0.809 [0.000] | 1 |
| PC1 coef. | .5460801 | .5475816 | .4346101 | .3314093 | .3212987 |

p-values in brackets. PC1 eigenvalue: 2.61 (52% of variance, PC2 eigenvalue: 1.29)

Note: The table reports the correlation among the BLM protest measures that compose the first principal component used in column 6 of Table 2, and the factor loadings of the first principal component.

Table A9: Cooperative Election Study 2020: Selected Survey Questions

| Category | ID | Question | Scale (low/high) |
|-----------|-------------|---|-------------------------------------|
| Police | CC20_334b | Do you support or oppose requiring police officers to wear body cameras that record all of their activities while on duty? | Support/Oppose |
| | CC20_334c | Do you support or oppose increasing the number of police on the street by 10 percent even if it means fewer funds for other public services? | Support/Oppose |
| | CC20_334d | Do you support or oppose decreasing the number of police on the street by 10 percent and increasing funding for other public services? | Support/Oppose |
| | CC20_334e | Do you support or oppose banning the use of choke holds by police? | Support/Oppose |
| | CC20_334f | Do you support or oppose creating a national registry of police who have been investigated for or disciplined for misconduct? | Support/Oppose |
| | CC20_334g | Do you support or oppose ending the Department of Defense program that sends surplus military weapons and equipment to police departments? | Support/Oppose |
| | CC20_334h | Do you support or oppose allowing individuals or their families to sue a police officer for damages if the officer is found to have “recklessly disregarded” the individual’s rights? | Support/Oppose |
| Racism | CC20_440a | White people in the U.S. have certain advantages because of the color of their skin. | Strongly agree to Strongly disagree |
| | CC20_440b | Racial problems in the U.S. are rare isolated situations. | Strongly agree to Strongly disagree |
| | CC20_441a | Irish, Italians, Jewish, and many other minorities overcame prejudice and worked their way up. Blacks should do the same without any special favors. | Strongly agree to Strongly disagree |
| | CC20_441b | Generations of slavery and discrimination have created conditions that make it difficult for blacks to work their way out of the lower class. | Strongly agree to Strongly disagree |
| Media Use | CC20_300_1 | In the past 24 hours have you [1] Used social media (such as Facebook or Youtube) [2] Watched TV news [3] Read a newspaper in print or online [4] None of these | No/Yes |
| | CC20_300d_3 | In the past 24 hours, did you do any of the following on social media (such as Facebook, Youtube or Twitter)? [1] Posted a story, photo, video or link about politics [2] Posted a comment about politics [3] Read a story or watched a video about politics [4] Followed a political event [5] Forwarded a story, photo, video or link about politics to friends [6] None of the above | No/Yes |
| Protest | CC20_430a_4 | During the past year did you attend a political protest march or demonstration? | No/Yes |
| COVID | CC20_309a | Have you or someone you know been diagnosed with the novel coronavirus (COVID-19) during the past year? (select all that apply) [1] Yes, I have [2] Yes, a family member | No/Yes |

Table A9: Cooperative Election Study 2020: Selected Survey Questions

| Category | ID | Question | Scale (low/high) |
|-----------------|-------------|--|-------------------------|
| Abortion | CC20_332a_1 | On the topic of abortion, do you support or oppose each of the following proposals? Always allow a woman to obtain an abortion as a matter of choice | Oppose/Support |
| Environment | CC20_333a_1 | Do you support or oppose each of the following proposals? Give the Environmental Protection Agency power to regulate Carbon Dioxide emissions | Oppose/Support |

Note: Framing of the questions used in our analysis from the Cooperative Election Study 2020. The values are given from the value with the lowest numerical coding to the value with highest numerical coding.

Table A10: **First stage of the individual-level regression, varying the radius of super spreader events.**

| Radius SSE | Social media use | | | | |
|-------------------------------------|------------------------|-----------------------|------------------------|-------------------------|-------------------------|
| | (1) county | (2) 50 km | (3) 100 km | (4) 200 km | (5) 300 km |
| $SSE_c \times SXS W_c \times Age_i$ | 0.0139*** (0.00371) | 0.00290* (0.00170) | 0.000635 (0.000780) | -0.000125 (0.000392) | -0.000202 (0.000288) |
| F stat | 13.96 | 2.913 | 0.663 | 0.102 | 0.492 |
| Mean dep. var | 0.745 | 0.745 | 0.745 | 0.745 | 0.745 |
| Observations | 48382 | 48382 | 48382 | 48382 | 48382 |
| $SSE_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes |
| $SXS W_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes |
| $SXS W Pre_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes |
| COVID-19 deaths $\times Age_i$ | Yes | Yes | Yes | Yes | Yes |
| Baseline Twitter $\times Age_i$ | Yes | Yes | Yes | Yes | Yes |
| Individual controls | Yes | Yes | Yes | Yes | Yes |
| County fixed effects | Yes | Yes | Yes | Yes | Yes |

Note: First stage regression with varying radius of superspreader events. Individual level regression following specification 4. The treatment variable is the push-pull instrument $SSE \times SXS W$ interacted with the age of the respondent. The outcome is a dummy variable equal to 1 if the respondent used social media in the last 24 hours, taken from the Cooperative Congressional Election Study conducted in November of 2020. All models include county fixed effects, individual controls (age, gender, employment status, number of children, dummies for race and for vote in the 2016 presidential election), and interaction of age with instrument controls (SSE and $SXS W$ individually, pre- $SXS W$ followers and its interaction with age, COVID deaths per thousand and baseline social media penetration). Column 1 uses SSEs inside the county of the respondent. Column 2 uses the SSEs outside the county of the respondent, but less than 50 km from the border of the county. Similarly, columns 3, 4 and 5 use, respectively, SSEs within 100, 200, and 300 km from the county of the respondent. The treatment variable is the push-pull instrument $SSE_c \times SXS W_c$ interacted with the age of the respondent. Standard errors (in parentheses) are clustered at the county level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A11: Placebo first stage: other media consumption, COVID exposure and other attitudes

| | social media (1) | TV (2) | newspaper (3) | COVID self (4) | COVID family (5) | COVID work reduction (6) | abortion rights (7) | environmental protection (8) |
|--|------------------------|-----------------------|----------------------|----------------------|------------------------|--------------------------------|---------------------------|------------------------------------|
| $SSE_c \times SXS\bar{W}_c \times Age_i$ | 0.0139*** (0.00371) | -0.00546 (0.00477) | 0.00368 (0.00304) | 0.00180 (0.00213) | -0.00675 (0.00556) | -0.00168 (0.00246) | 0.00128 (0.00364) | -0.00382 (0.00295) |
| F stat | 13.96 | 1.307 | 1.471 | 0.718 | 1.473 | 0.468 | 0.123 | 1.679 |
| Mean dep. var | 0.745 | 0.630 | 0.396 | 0.0427 | 0.207 | 0.0989 | 0.601 | 0.696 |
| Observations | 57472 | 57472 | 57472 | 48382 | 48382 | 48382 | 48369 | 48366 |
| $SSE_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| $SXS\bar{W}_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| $SXS\bar{W} Pre_c \times Age_i$ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| COVID-19 deaths $\times Age_i$ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Baseline Twitter $\times Age_i$ | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Individual controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| County fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

Note: Placebo first stage regression. Individual level regression following specification 4. The treatment variable is the push-pull instrument $SSE \times SXS\bar{W}$ interacted with the age of the respondent. Outcomes are taken from the Cooperative Congressional Election Study conducted in November of 2020. Column 1 is dummy variable for whether respondent has used social media in the past 24 hours. Column 2 and 3 use the same question for TV and newspapers. Column 4 to 6 ask whether the respondent has received a COVID-19 diagnosis (column 4), whether a family member has received a COVID-19 diagnosis (column 5), or if the respondent experienced reduced work hours as a result of COVID-19 (column 6). Column 7 asks whether respondents support the statement "A woman should always be allowed to obtain an abortion as a matter of choice". Column 8 asks whether respondents support the statement "The Environmental Protection Agency should have the power to regulate Carbon". The treatment variable is the push-pull instrument $SSE_c \times SXS\bar{W}_c$ interacted with the age of the respondent, with SSE_c counting SSEs inside the county. All models include county fixed effects, individual controls (age, gender, employment status, number of children, dummies for race and for vote in the 2016 presidential election), and interaction of age with instrument controls (SSE and $SXS\bar{W}$ individually, pre- $SXS\bar{W}$ followers and its interaction with age, COVID deaths per thousand and baseline social media penetration). Standard errors (in parentheses) are clustered at the county level. *** p<0.01, ** p<0.05, * p<0.1

Table A12: Individual-level evidence on social media use and political preferences

| | has protested for any cause | political activity on social media | voted for Biden in 2020 | PC1 police | PC1 racism | Placebo | |
|------------------------------|--------------------------------|---------------------------------------|----------------------------|------------------------|------------------------|-------------------------|------------------------|
| Panel A: 2SLS | (1) | (2) | (3) | (4) | (5) | abortion (6) | environment (7) |
| Social media use | 0.472** (0.231) | 1.068*** (0.220) | 0.523* (0.310) | -2.434** (1.065) | -1.066 (1.039) | 0.0916 (0.264) | -0.274 (0.232) |
| Panel B: OLS | | | | | | | |
| Social media use | 0.0416*** (0.00270) | 0.697*** (0.00339) | 0.0126*** (0.00364) | -0.177*** (0.0153) | -0.168*** (0.0141) | -0.0127*** (0.00452) | 0.0122*** (0.00465) |
| Panel C: Reduced form | | | | | | | |
| SSE × SXSXW × Age | 0.00654** (0.00259) | 0.0148*** (0.00359) | 0.00724** (0.00365) | -0.0335*** (0.0128) | -0.0138 (0.0134) | 0.00128 (0.00364) | -0.00382 (0.00295) |
| Panel D: First stage | | | | | | | |
| | Social media use | | | | | | |
| SSE × SXSXW × Age | 0.0139*** (0.00371) | 0.0139*** (0.00371) | 0.0139*** (0.00371) | 0.0138*** (0.00369) | 0.0129*** (0.00388) | 0.0139*** (0.00371) | 0.0139*** (0.00371) |
| Observations | 48,382 | 48,382 | 48,382 | 48,308 | 47,093 | 48,369 | 48,366 |
| Mean dep. var. | 0.0783 | 0.526 | 0.426 | 0 | 0 | 0.611 | 0.699 |
| Kleibergen-Paap F stat | 13.96 | 13.96 | 13.96 | 13.88 | 11.05 | 14.12 | 14.04 |
| Age × Instrument controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Individual controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| County fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

Note: Individual level regression following specification 4. Outcomes are taken from the Cooperative Congressional Election Study conducted in November of 2020. Column 1 is dummy variable for whether respondent has attended a protest, march or demonstration in the past year. Column 2 is a dummy variable for whether the respondent has used social media in the past 24 hours to post or forward a story, video or link about politics, or to post a comment about politics, or to watch a video about politics. Column 3 is a dummy variable for whether respondent voted for Biden in the 2020 presidential election. Column 4 is the first principal component of questions that ask whether the respondent supports more oversight of police including allowing individuals and their families to sue police officers, ending the DoD program that sends military surplus to police departments, creating a national registry of police who have been investigated, banning the use of choke holds, requiring police to wear body cameras, having more or less police in the streets. Column 5 is the first principal component of anti-racist attitudes, including that white people have advantages in society, that racial problems are not isolated situations, that slavery created conditions that make it hard for Blacks to advance, that the success of other minorities does not prove that Blacks can do the same. Column 6 asks whether respondents support the statement "A woman should always be allowed to obtain an abortion as a matter of choice". Column 7 asks whether respondents support the statement "The Environmental Protection Agency should have the power to regulate Carbon". Treatment is a dummy variable indicating if the respondent used social media in the past 24 hours. The instrument is the previous push-pull instrument $SSE_c \times SXSXW_c$ interacted with the age of the respondent, but using SSEs inside of the county instead of SSEs in neighboring counties. All models include county fixed effects, individual controls (age, gender, employment status, number of children, dummies for race and for vote in the 2016 presidential election), and non-OLS models include the interaction of age with instrument controls (SSE and SXSXW individually, pre-SXSXW followers and its interaction with age, COVID deaths per thousand and baseline social media penetration). We report Kleibergen-Paap rkWald F statistic. Standard errors are clustered at the county level. *** p<0.01, ** p<0.05, * p<0.1

Table A13: Principal component analysis of Cooperative Election Study 2020 data

(a) Racism

| | CC20_440a | CC20_441b | CC20_440b | CC20_441a |
|------------------|-------------------|-------------------|------------------|-----------|
| CC20_440a | 1 | | | |
| CC20_441b | 0.760 [0.000] | 1 | | |
| CC20_440b | -0.648 [0.000] | -0.608 [0.000] | 1 | |
| CC20_441a | -0.695 [0.000] | -0.746 [0.000] | 0.632 [0.000] | 1 |
| PC1 coef. | .5107706 | .5133582 | -.4689424 | -.505638 |

p-values in brackets. PC1 eigenvalue: 3.05 (76% of variance, PC2 eigenvalue: 0.43)

(b) Attitudes towards police

| | CC20_334b | CC20_334c | CC20_334d | CC20_334e | CC20_334f | CC20_334g | CC20_334h |
|------------------|--------------------|-------------------|------------------|------------------|------------------|------------------|-----------|
| CC20_334b | 1 | | | | | | |
| CC20_334c | -0.0849 [0.000] | 1 | | | | | |
| CC20_334d | 0.134 [0.000] | -0.547 [0.000] | 1 | | | | |
| CC20_334e | 0.264 [0.000] | -0.262 [0.000] | 0.318 [0.000] | 1 | | | |
| CC20_334f | 0.328 [0.000] | -0.221 [0.000] | 0.279 [0.000] | 0.378 [0.000] | 1 | | |
| CC20_334g | 0.140 [0.000] | -0.393 [0.000] | 0.505 [0.000] | 0.348 [0.000] | 0.281 [0.000] | 1 | |
| CC20_334h | 0.284 [0.000] | -0.294 [0.000] | 0.332 [0.000] | 0.363 [0.000] | 0.420 [0.000] | 0.344 [0.000] | 1 |
| PC1 coef. | .2546515 | -.376133 | .4236488 | .3838915 | .3749601 | .4073791 | .4003408 |

p-values in brackets. PC1 eigenvalue: 2.90 (41% of variance, PC2 eigenvalue: 1.18)

Note: Each table reports the correlation among the racism and police attitude measures, and the factor loadings of the first principal components used in columns 4 and 5 of Table 4. The detailed wording of the questions is given in Table A9.

Appendix B: Robustness Checks

We run a series of empirical exercises in Tables B1 to B6 which we describe in more detail here. We always present baseline results and reduced form coefficients in column 1, which replicates the specification used in column 4 of Table 1.

Changing the radius and time window around SSEs. In the baseline specification, we choose the 200km threshold as a distance of the SSE to the county border. To make sure that this choice is not driving our results, we change the radius of influence to 100 km, 150 km, 300 km and 400 km (columns 2 to 5 of Table B1). We show 2SLS, reduced form and first stage results. Throughout we find that new Twitter users increase BLM protest. However, the first stage becomes weak when we expand the radius of SSEs to 400km, potentially because the salience of pandemic exposure and thus the relevance of the push factor is lower for SSEs that are further away. We repeat this exercise, this time expanding or narrowing the time lag between SSEs and the protest trigger. Our baseline specification uses the cumulative number of SSEs until six weeks before the murder of George Floyd, i.e., 12 weeks into the pandemic. This choice is based on the idea that SSEs in the early stages of the pandemic reflect quasi-random variation in the presence of a super-spreader rather than risk preferences of the local population. SSEs further into the pandemic might be more endogenous as knowledge about the spread of the virus increased over time. We show in the bottom panel of the table that our results hold when we include SSEs until 4 weeks, 5 weeks, 7 weeks or 8 weeks before the protest trigger.

Excluding SSEs in prisons. About 15% of SSEs occurred inside prisons (see Table A4). We exclude SSEs in prisons in a robustness check in column 2 of Table B2 for two reasons. First, it is likely that by the nature of prisons, the geographical spread of cases stemming from an SSE in a prison is quite limited and less relevant for the overall population and the protesting population. Second, SSEs in prisons may affect BLM protests (and, to a lesser extent, social media usage) in ways other than through overall social media adoption, for instance, by increasing awareness of the disproportionate incarceration of Black people. While the salience of racial inequality in prisons may be a possible mechanism, with this exercise we investigate whether our results are indeed solely driven by this subsample of SSEs. We exclude SSEs in prisons in column 2 and find that our results are not sensitive to this choice.

Controlling for SSEs in the county. Our pandemic exposure variable measures the effect of having an SSE outside the county within 200 km of the county border and excluding the effect of SSEs that take place within its border. Therefore, in our analysis a county is "not affected" by an SSE if its border is either further than 200 km from the SSE, or the SSE happened within its boundaries. To assuage the concern that correlation of SSEs across counties is driving the variation in SSE exposure, we add as a control the number of SSEs that occurred within the county itself. Estimates are presented in column 3 of Table B2 and show that the results of the baseline specification are robust to the addition of this control.

Weighting SSEs by distance. In our baseline specification, we count any SSE that occurred in a 200 km radius outside the border of a county as an additional SSE affecting the county in an uniform way. However, an SSE 1 km away from the border is likely to have a different level of influence from a SSE 49 km away. To ensure that this simplification is not driving the results, we refine the level of influence by weighting the SSEs by a linear function decreasing with distance (column 4 of Table B2), giving less weight to events that are more distant. The results are robust to this distance weighting procedure. The first stage is above the conventional threshold but smaller than for our baseline instrument, potentially because we discount SSEs that still act as a push factor for social media adoption.

Controlling for past protest in neighboring counties. We observed in the balance test (Figure 2) that SSEs may be correlated with BLM protests in 2014-2019 in neighboring counties. To make sure that this does not affect the results, we control for the presence of such protests in column 2 of Table B3. This does not affect the results.

Controlling for COVID-19 prevalence before the protest trigger. The intensity of COVID-19 around the protest could directly affect the willingness and opportunity to protest (e.g. due to health concerns or local restrictions). That is, counties with a higher pandemic exposure early in

the pandemic may respond differentially to the protest trigger compared to those that have been exposed later. In column 3 and 4 of Table B3, we show that our results are not sensitive to adding a control for the number of COVID-19 cases or deaths in the 7 days before May 24th.

Controlling for propensity to protest. We focus on the sample of counties without prior BLM protest. However, there is a large heterogeneity among these counties - even conditional on the set of controls and state fixed effects. Some counties may be at the margin, while others are very unlikely to mount a pro-BLM protest even in the presence of the push-pull shock. Therefore, we attempt to account for ex-ante protest propensity in a data-driven way.

Using LASSO regression, we construct a variable measuring propensity to protest in the 3 weeks following a notable death of a Black person in an encounter with the police. We consider a death notable if it received national media coverage. We construct a dataset of 31 notable deaths between 2014 and 2019, and we select our predictors among a set of county characteristics, including a large set of socio-demographic and economic characteristics extracted from the American Community Survey (such as population, population density, race distribution, age groups, poverty rates, among others), indicators for different levels of urbanization, geographical indications (latitude, longitude, and state dummies). This allows us to assign ex-ante protest propensities (continuous measure between 0 and 1) even to the sub-set of counties with no prior BLM protest. We first use it directly as a control (column 5 of Table B3).

In addition, we include fixed effects for different levels of the propensity to protest. We group observations by groups of 1000 units (3 groups), 100 units (28 groups) and 10 units (277 groups) with similar propensity to protest and add fixed effects for each group. Results are shown in columns 6 to 8 of Table B3. This is essentially a matching strategy, where the fixed effects ensure that observations with similar propensity are compared. Results are robust to the inclusion of fixed effects, except in column 8 where - as expected - results become imprecisely estimated for groups of size 10.

Excluding coastal counties and states. Coastal states and counties might behave differently with regard to our instrument, to social media adoption, and to protest. Coastal regions are generally denser, which increases the chance of having an SSE (Figure 1 shows the density of SSEs). Coastal counties also differ in the construction of the push component of the instrument. Coastal counties naturally have fewer neighbors, which decreases their chances of being treated by the push component of the instrument. In addition, these counties have also had higher BLM protest activity in the past. We show that our results are robust to excluding coastal counties (column 2 of Table B4), as well as coastal states (column 3).

Accounting for spatial correlation. Observations are likely to be spatially correlated for several reasons. For instance, there could be spatially-correlated unobserved factors influencing the decision to protest (such as weather conditions or available TV and radio stations). Clustering by state does not entirely remove these errors because correlation across state borders remains (Colella et al., 2019). To overcome this problem, we use Conley standard errors that allow for spatial correlation within a certain distance. Column 4 to 7 of Table B4 show the estimates when allowing spatial correlation between observations in a 100, 200 and 300 km radius. Column 7 shows the estimates when allowing spatial correlation with all neighboring counties. Reassuringly, our results remain robust and the first stage becomes even stronger when accounting for spatial correlation.

Treatment definition. In Table B5, we probe the robustness with respect to changes in the definition of the treatment. Instead of the log of one plus the number of new users, we take the number of new users as a share of total users (column 2), the inverse hyperbolic sine transformation of new users (column 3), as well as the absolute number of new users (column 4). Throughout we find a positive and precisely estimated effect of new users on BLM protest.

Probit estimation. In our baseline specification the effect of social media is additive. It might be the case that the effect would be multiplicative of some characteristics of the counties. Using a Probit model accounts for this possibility. Non-linear models with many covariates (typically when using fixed effects) suffer from the incidental parameter problem resulting in bias of the estimates (Manski et al., 1981; Lancaster, 2000; Wooldridge, 2015). To reduce the extent of this problem we omit the state fixed effects, which significantly reduces the number of covariates. We use an OLS in the first stage, but

estimate the second stage with a Probit model. Results are presented in column 5 of Table B5. The Probit model delivers a precisely estimated and positive coefficient (although magnitudes are not directly comparable).

Plausibility of exclusion restriction for BLM. If our push-pull instrument $SSE_c \times SXSW_c$ were to pick up any underlying factors correlated with the overall likelihood of protesting for a BLM-related cause, then this would challenge a causal interpretation of our estimates. To probe the plausibility of the exclusion restriction, we estimate the effect of instrumented new Twitter users on the likelihood of observing past BLM protests, using the full sample of counties. In column 2 of Table B6, we show that instrumented new Twitter users do not predict the presence of BLM events between 2014 and 2019. We take this as additional evidence for the plausibility of our identifying assumption.

Outcome definition. In our baseline specification, we choose the three week window following Floyd’s murder since it captures the vast majority of BLM-related protests, while being close enough to the protest trigger. The expansion of the time frame may lead to time-varying differences across states that drive BLM protest. Narrowing the time window will not allow us to capture the time it takes to mount a protest. Nonetheless, we show that our main results are robust to reducing this time window to 2 weeks and expanding this time window to 4, 5 and 6 weeks (columns 3 to 6 of Table B6 respectively).

Event study robustness. Figure B1 shows robustness of our event study results. In panel a) we investigate the intensive margin of protest, counting the number of BLM protest in a given county and month. Panel b) focuses on the sample of counties with any BLM protest after the murder of George Floyd. This allows us to investigate whether counties with higher levels of late adoption protest earlier in response to the protest trigger. Panel c) fully saturates the model and interacts the post dummy with the full set of baseline controls. This assuages concerns that factors that are associated with late adoption and BLM protest may have impacted protest propensity in 2020 but not before. Lastly, in panel d) we expand the time frame even further into the past to account for differences in protest propensity in the early days of the movement. Throughout our estimates are consistent with our baseline findings.

Table B1: Robustness - main results with instrument variation

| | At least one BLM protest | | | | |
|---|--------------------------|--------------------------|--------------------------|---------------------------|---------------------------|
| | baseline | (2) | (3) | (4) | (5) |
| Radius SSE | (1) | (2) | (3) | (4) | (5) |
| Time lag SSE: 42 days | 200 km | 100 km | 150 km | 300 km | 400 km |
| 2SLS: new Twitter users | 0.443*** (0.146) | 0.393* (0.212) | 0.415** (0.168) | 0.494** (0.201) | 0.531** (0.261) |
| Reduced form: SSE \times SXSW | 0.00104*** (0.000297) | 0.00152** (0.000612) | 0.00105*** (0.000339) | 0.000776*** (0.000154) | 0.000546*** (0.000140) |
| First stage: SSE \times SXSW | 0.00234*** (0.000460) | 0.00386*** (0.00114) | 0.00252*** (0.000777) | 0.00157*** (0.000409) | 0.00103** (0.000491) |
| Kleibergen-Paap F stat | 25.89 | 11.51 | 10.54 | 14.81 | 4.401 |
| Time lag SSE | 42 days | 28 days | 35 days | 49 days | 56 days |
| Radius SSE: 200km | | | | | |
| 2SLS: new Twitter users | 0.443*** (0.146) | 0.438*** (0.145) | 0.441*** (0.146) | 0.452*** (0.149) | 0.444*** (0.142) |
| Reduced form: SSE \times SXSW | 0.00104*** (0.000297) | 0.00102*** (0.000286) | 0.00103*** (0.000292) | 0.00108*** (0.000319) | 0.00126*** (0.000376) |
| First stage: SSE \times SXSW | 0.00234*** (0.000460) | 0.00232*** (0.000442) | 0.00233*** (0.000452) | 0.00238*** (0.000480) | 0.00283*** (0.000562) |
| Kleibergen-Paap F stat | 25.89 | 27.55 | 26.57 | 24.62 | 25.40 |
| Observations | 2767 | 2767 | 2767 | 2767 | 2767 |
| State FE | Yes | Yes | Yes | Yes | Yes |
| Pandemic exposure | Yes | Yes | Yes | Yes | Yes |
| Baseline Twitter | Yes | Yes | Yes | Yes | Yes |
| Economic controls | Yes | Yes | Yes | Yes | Yes |
| BLM controls | Yes | Yes | Yes | Yes | Yes |
| Political controls | Yes | Yes | Yes | Yes | Yes |

Note: This table shows the main regression of presence of BLM protests on new Twitter users, varying the radius and the time lag at which super spreader events (used in the instrument) are considered. SSE are measured as the number of super spreading events within a certain radius (Panel A) and until a certain number of days before the murder of George Floyd on May 25th 2020 (Panel B). The baseline instrument is reported in column 1 and uses a radius of 200km and 42 day lag. We report Kleibergen-Paap rkWald F statistics for each instrument variation at the bottom of each panel, as well as the reduced form and first stage coefficients. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table B2: Robustness - SSE definition and weighting

| | At least one BLM protest | | | |
|--|--------------------------|--------------------------|---------------------------|--------------------------|
| | baseline (1) | (2) | (3) | (4) |
| 2SLS: new Twitter users | 0.443*** (0.146) | 0.444*** (0.143) | 0.409*** (0.136) | 0.444** (0.196) |
| Reduced form: SSE \times SXSX | 0.00104*** (0.000297) | 0.00113*** (0.000333) | 0.000978*** (0.000303) | 0.00184*** (0.000642) |
| Observations | 2,767 | 2,767 | 2,767 | 2,767 |
| Mean dep. var. | 0.0477 | 0.0477 | 0.0477 | 0.0477 |
| First stage coef. SSE \times SXSX | 0.00234 | 0.00256 | 0.00239 | 0.00415 |
| First stage s.e. SSE \times SXSX | (0.000460) | (0.000519) | (0.000491) | (0.00123) |
| Kleibergen-Paap F stat | 25.89 | 24.27 | 23.73 | 11.45 |
| Excluding SSEs in prisons | | Y | | |
| Control SSE in county | | | Y | |
| SSE distance weighting | | | | Y |
| All controls | Y | Y | Y | Y |
| State fixed effects | Y | Y | Y | Y |

Note: Results on robustness of the results to variations in the SSE measure, using different time and distance selection. The results correspond to the 2SLS results in Column 4 of Table 1 (ie. the effect of the logarithm of one plus the number of new Twitter users in the sample, instrumented by SSE \times SXSX). Column 1 reports our baseline specification, corresponding to column 4 of Table 1. Column 2 excludes SSEs that took place in prisons. In column 3, a control is added for the number of SSEs within the county 6 weeks before the murder of George Floyd. Columns 4 weighs the effect of SSEs by distance with smaller weights given to more distant SSEs. Weights are applied linearly. All specifications include the whole set of controls and state fixed effects. For 2SLS estimates, we report the Kleibergen-Paap rkWald F statistic, the first stage coefficient, the reduced form estimate, and corresponding standard errors. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table B3: Robustness - additional controls

| | At least one BLM protest | | | | | | | |
|--|--------------------------|--------------------------|--------------------------|--------------------------|---------------------------|--------------------------|---------------------------|--------------------------|
| | baseline | | | | | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| 2SLS: new Twitter users | 0.443*** (0.146) | 0.438*** (0.150) | 0.452*** (0.151) | 0.444*** (0.147) | 0.531** (0.226) | 0.467*** (0.160) | 0.478** (0.196) | 0.440 (0.263) |
| Reduced form: SSE \times SXSX | 0.00104*** (0.000297) | 0.00100*** (0.000300) | 0.00104*** (0.000298) | 0.00104*** (0.000297) | 0.000993*** (0.000340) | 0.00102*** (0.000301) | 0.000872*** (0.000293) | 0.000660** (0.000326) |
| Observations | 2,767 | 2,767 | 2,767 | 2,767 | 1,560 | 2,767 | 2,767 | 2,767 |
| Mean dep. var. | 0.0477 | 0.0477 | 0.0477 | 0.0477 | 0.0744 | 0.0477 | 0.0477 | 0.0477 |
| First stage coef. SSE \times SXSX | 0.00234 | 0.00229 | 0.00230 | 0.00233 | 0.00187 | 0.00219 | 0.00182 | 0.00150 |
| First stage s.e. SSE \times SXSX | (0.000460) | (0.000462) | (0.000452) | (0.000459) | (0.000503) | (0.000449) | (0.000439) | (0.000484) |
| Kleibergen-Paap F stat | 25.89 | 24.55 | 25.85 | 25.75 | 13.82 | 23.77 | 17.27 | 9.627 |
| Neighbor had BLM protest before 2020 | | Y | | | | | | |
| COVID cases in past 7 days | | | Y | | | | | |
| COVID deaths in past 7 days | | | | Y | | | | |
| Propensity to protest | | | | | Y | | | |
| Propensity to protest group: size | | | | | | 1000 | 100 | 10 |
| All controls | Y | Y | Y | Y | Y | Y | Y | Y |
| State fixed effects | Y | Y | Y | Y | Y | Y | Y | Y |

Note: Results on robustness of the results to variations in controls. The results correspond to the 2SLS results in Column 4 of Table 1 (ie. the effect of the logarithm of one plus the number of new Twitter users in the sample, instrumented by SSE \times SXSX). Column 1 reports the baseline specification. Column 2 controls by whether a neighboring county had a BLM protest prior to 2020. Column 3 and 4 control respectively by the COVID cases and deaths in the week preceding the murder of George Floyd. Column 5 adds a control for the propensity to protest derived from our LASSO selection model. Columns 6 to 8 add fixed effects for propensity to protest for groups of size 1000, 100 and 10 respectively. All specifications include the whole set of controls and state fixed effects. For 2SLS estimates, we report the Kleibergen-Paap rkWald F statistic, the first stage coefficient, the reduced form estimate, and corresponding standard errors. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table B4: Robustness - sample composition and spatial clustering

| | At least one BLM protest | | | | | | |
|---|--------------------------|---------------------------|-------------------------|----------------------|--------------------------|--------------------------|--------------------------|
| | baseline | | | | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| 2SLS: new Twitter users | 0.443*** (0.146) | 0.418*** (0.146) | 0.259* (0.133) | 0.443*** (0.0332) | 0.443*** (0.0202) | 0.443*** (0.114) | 0.443** (0.180) |
| Reduced form: SSE \times SXSXW | 0.00104*** (0.000297) | 0.000945*** (0.000302) | 0.000530* (0.000284) | 0.00104 | 0.00104*** (0.000173) | 0.00104*** (0.000280) | 0.00104*** (0.000303) |
| Observations | 2,767 | 2,616 | 1,697 | 2,768 | 2,768 | 2,768 | 2,768 |
| Mean dep. var. | 0.0477 | 0.0428 | 0.0371 | 0.0477 | 0.0477 | 0.0477 | 0.0477 |
| First stage coef. SSE \times SXSXW | 0.00234 | 0.00226 | 0.00204 | 0.00234 | 0.00234 | 0.00234 | 0.00234 |
| First stage s.e. SSE \times SXSXW | (0.000460) | (0.000413) | (0.000372) | (0.000419) | (0.000252) | | (0.000507) |
| Kleibergen-Paap F stat | 25.89 | 29.89 | 30.24 | 30.39 | 84.02 | | 20.82 |
| Excluding coastal | | counties | states | | | | |
| Spatial clustering | | | | 100 km | 200 km | 300 km | neighbors |
| All controls | Y | Y | Y | Y | Y | Y | Y |
| State fixed effects | Y | Y | Y | Y | Y | Y | Y |

Note: Results on robustness of the results to variations in sample composition and clustering. The base sample consists of counties with no BLM protest before George Floyd's murder, except in column 2. The results correspond to the 2SLS results in Column 4 of Table 1 (ie. the effect of the logarithm of one plus the number of new Twitter users in the sample, instrumented by SSE \times SXSXW). Column 1 corresponds to the baseline specification. Column 2 and 3 exclude, respectively, coastal counties and coastal states. Column 4 to 7 use Conley standard errors to account for spatial correlation: column 4, 5 and 6 allow a spatial correlation at distances 100 km, 200 km and 300 km respectively. Column 7 allows correlation between direct neighbors. All specifications include the whole set of controls and state fixed effects. For 2SLS estimates, we report the Kleibergen-Paap rkWald F statistic, the first stage coefficient, the reduced form estimate, and corresponding standard errors. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table B5: Robustness - treatment definition, non-linear estimation

| | At least one BLM protest | | | | |
|--|--------------------------|--------------------------|--------------------------|--------------------------|----------------------|
| | baseline (1) | (2) | (3) | (4) | (5) |
| 2SLS: new Twitter users | 0.443*** (0.146) | 8.154* (4.662) | 0.363*** (0.124) | 0.0926*** (0.0282) | 1.299** (0.538) |
| Reduced form: SSE \times SXSX | 0.00104*** (0.000297) | 0.00104*** (0.000297) | 0.00104*** (0.000297) | 0.00104*** (0.000297) | 0.00378 (0.00251) |
| Observations | 2,767 | 2,618 | 2,767 | 2,767 | 2,766 |
| Mean dep. var. var. | 0.0477 | 0.0504 | 0.0477 | 0.0477 | 0.444 |
| First stage coef. SSE \times SXSX | 0.00234 | 0.000125 | 0.00286 | 0.0112 | 0.00244 |
| First stage s.e. SSE \times SXSX | (0.000460) | (6.44e-05) | (0.000616) | (0.00382) | (0.000432) |
| Kleibergen-Paap F stat | 25.89 | 3.750 | 21.49 | 8.601 | |
| New users transformation | log(1+new users) | new users / total | IHS | untransformed | log(1+new users) |
| Probit second stage | | | | | Y |
| All controls | Y | Y | Y | Y | Y |
| State fixed effects | Y | Y | Y | Y | |

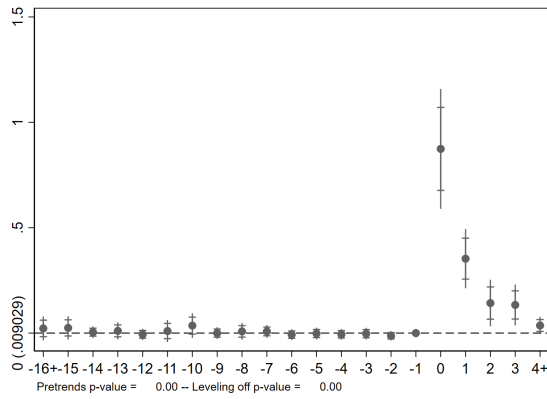
Note: Results on robustness of the results to variations in treatment definition, and non-linear estimation. The sample consists of counties with no BLM protest before George Floyd's murder. The results correspond to the 2SLS results in Column 4 of Table 1 (ie. the effect of new Twitter users in the sample, instrumented by SSE \times SXSX). Column 1 corresponds to the baseline specification, where new users are expressed as the logarithm of one plus the count of new users. Column 2 uses the ratio of new users over total users. Column 3 transforms the number of new users with the inverse hyperbolic sine function. Column 4 takes the absolute number of new users. Column 5 uses the usual logarithm of one plus new Twitter users, but estimates the second stage using a probit regression (excluding state fixed effects). All specifications include the whole set of controls and state fixed effects. For 2SLS estimates, we report the Kleibergen-Paap rkWald F statistic, the first stage coefficient, the reduced form estimate, and corresponding standard errors. Standard errors (in parentheses) are clustered at the state level. *** p<0.01, ** p<0.05, * p<0.1

Table B6: Robustness - outcome definition

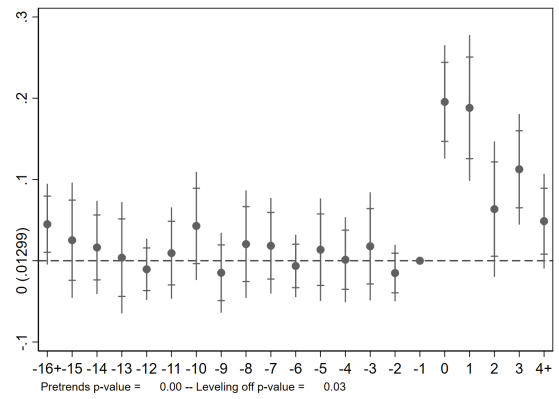
| | Presence of BLM events | | | | | |
|------------------------------------|------------------------|-------------------|-------------------|---------------------|--------------------|--------------------|
| | baseline | | | | | |
| | 3 weeks | Past events | 2 weeks | 4 weeks | 5 weeks | 6 weeks |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| 2SLS: new Twitter users | 0.443*** (0.146) | -0.249 (0.327) | 0.265* (0.142) | 0.412*** (0.143) | 0.272** (0.121) | 0.344** (0.146) |
| Observations | 2,767 | 3,106 | 2,767 | 2,767 | 2,767 | 2,767 |
| Mean dep. var. | 0.0477 | 0.108 | 0.0354 | 0.0542 | 0.0868 | 0.0665 |
| First stage coef. SSE \times SXS | 0.00234 | 0.00172 | 0.00234 | 0.00234 | 0.00234 | 0.00234 |
| First stage s.e. SSE \times SXS | (0.000460) | (0.000369) | (0.000460) | (0.000460) | (0.000460) | (0.000460) |
| Kleibergen-Paap F stat | 25.89 | 21.60 | 25.89 | 25.89 | 25.89 | 25.89 |
| All controls | Y | Y | Y | Y | Y | Y |
| State fixed effects | Y | Y | Y | Y | Y | Y |

Note: Robustness of the results to changes in the window in which BLM events are considered. Results correspond to the 2SLS results in Column 2 of Table 2 (ie. the effect of the logarithm of one plus the number of new Twitter users in the sample, instrumented by SSE \times SXS). The sample consists of counties with no BLM protest before George Floyd's murder, except in column 2 where the full sample is considered. Column 1 corresponds to the baseline specification. Column 2 predicts past BLM events (likelihood of observing a BLM event between 2014 and 2019) and uses all counties instead of only the counties with no BLM protest before George Floyd's murder. Columns 3, 4, 5 and 6 present different time windows for BLM protests: 2, 4, 5 and 6 weeks. All specifications include the whole set of controls and state fixed effects. We report the Kleibergen-Paap rkWald F statistic, the first stage coefficient and the mean of dependent variable. Standard errors (in parentheses) are clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

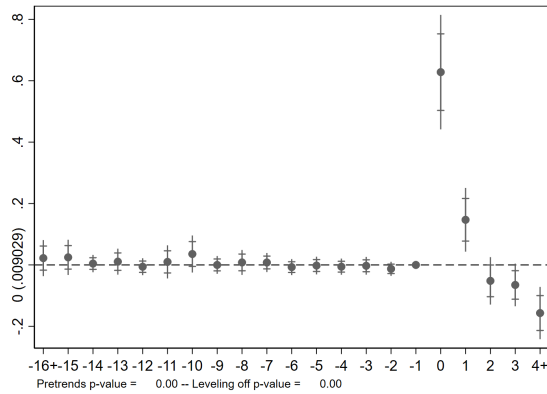
Figure B1: **Robustness: Event study for BLM protest after murder of George Floyd in counties with more late adopters**



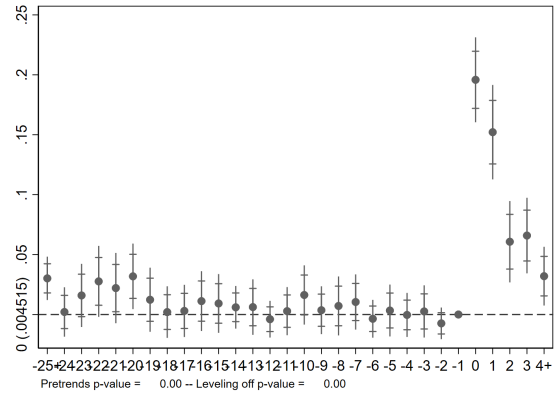
(a) Change outcome: number of protests in each month



(b) Change sample: counties with any BLM protest in 2020



(c) Change controls: fully saturated model with Post \times Controls X_c



(d) Change time frame: expanded to 24 months before event time

Note: Event study for the full samples of counties (including those with prior BLM protest) using *xtevent* by Freyaldenhoven, Hansen, Pérez and Shapiro (forthcoming) following specification 3 at the county-month level in the 16 months leading up to and 4 months following the murder of George Floyd. Regressions include county fixed effects and state-month fixed effects. We interact the post George Floyd period with the log of one plus the number of new Twitter users created after the outbreak of the pandemic but before the murder of George Floyd. Standard errors are clustered at the county level. Outcome is measured as a dummy variable for any BLM protest in that month.