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

Poverty and economic dislocation reduce compliance with COVID-19 shelter-in-place protocols☆

Austin L. Wright[a], Konstantin Sonin[a,d,∗], Jesse Driscoll[b], Jarnickae Wilson[c]

[a] Harris School of Public Policy, University of Chicago, United States
[b] School of Global Policy and Strategy, University of California, San Diego, United States
[c] J.P. Morgan, United States
[d] HSE University, Moscow, Russia

ARTICLE INFO

ABSTRACT

Shelter-in-place ordinances were the first wide-spread policy measures aimed to mitigate the spread of COVID-19. Compliance with shelter-in-place directives is individually costly and requires behavioral changes across diverse sub-populations. Leveraging county-day measures on population movement derived from cellphone location data and the staggered introduction of local mandates, we find that economic factors have played an important role in determining the level of compliance with local shelter-in-place ordinances in the US. Specifically, residents of low income areas complied with shelter-in-place ordinances less than their counterparts in areas with stronger economic endowments, even after accounting for potential confounding factors including partisanship, population density, exposure to recent trade disputes, unemployment, and other factors. Novel results on the local impact of the 2020 CARES Act suggest stimulus transfers that addressed economic dislocation caused by the COVID-19 pandemic significantly increased social distancing.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Shelter-in-place ordinances were the first major policy response to the COVID-19 pandemic. The effectiveness of these local mandates in the US relied heavily on compliance of the population at large without widespread legal enforcement. Efficient implementation of shelter-in-place policies has been found to reduce social contact and risks of interpersonal COVID-19 transmission (Hsiang et al., 2020; Viner et al., 2020; Anderson et al., 2020; Bai et al., 2020; Matrajt and Leung, 2020; Allcott et al., 2020) and to improve health outcomes during the pandemic (Fowler et al., 2020; Kapoor et al., 2020; Chudik et al., 2020; Dave et al., 2020b). Though the economic consequences of these policies are substantial (Goolsbee and Syverson, 2020; Stapleton, 2020; Baker et al., 2020; Gormsen and Koijen, 2020), local non-compliance creates public health risks and may contribute to regional spread (Lewnard and Lo, 2020; Chen et al., 2020). Clarifying how economic conditions influence compliance provides actionable insights for policy makers and public health officials responding to the COVID-19 pandemic. More generally, the setting provides a natural experiment for engaging a classic issue in political economy: what factors influence compliance with government policies (Gaventa, 1982; Scott, 1985; Tyler, 1990; Levi, 1997)?

In this paper, we study compliance with shelter-in-place policies during the critical early months of the 2020 pandemic using novel data on population movement, local policy changes, and variation in economic endowments. Our research design leverages the staggered introduction of local-level shelter-in-place mandates in a difference-in-differences design (Goodman-Bacon, 2018). The central assumption of our econometric technique is that treated and control units exhibit common trends (Bertrand et al., 2004; Angrist and Pischke, 2009). The use of event study designs enables us to assess the plausibility of the common trends assumptions using microlevel changes in cellphone movement (Goodman-Bacon and Marcus, 2020). Our findings reveal robust evidence of common trends prior to the introduction of local policy interventions to stop the spread of COVID-19. These results give confidence that our findings can be interpreted as causal evidence.

In addition, our research design flexibly accounts for a number of potential confounding factors, including government restrictions on social and business activities, local COVID severity, exposure to recent trade disputes, political partisanship, exposure to slanted media, unemployment, industry dynamics, as well as population density. This approach implements recent advances in econometrics regarding staggered difference-in-difference designs as well as dynamic heterogeneous effects (Goodman-Bacon, 2018; Goodman-Bacon and Marcus, 2020).

We find compelling evidence that shelter-in-place policies reduced population movement and these reductions in movement increased with local income. Event study designs suggest much of the nation-wide reduction in population movement is driven by counties above the median level of income. These core findings survive a number of alternative model specifications.

We supplement our main results with novel estimates of how stimulus transfers impact social distancing. This exercise serves as both a robustness check for our main findings and a result of independent policy interest. The CARES Act stimulus transfers were intended to offset negative economic shocks caused by public policy responses to COVID-19 (e.g., business closures, reduced travel). We use unseasonable federal bank deposits as a measure of electronic stimulus transfers under the CARES Act. The recipient-level roll out of the stimulus program was haphazard with some beneficiaries receiving financial transfers weeks ahead of others. We take advantage of these cumulative transfers at the county level. Consistent with our main results about the economic determinants of compliance, we find that local stimulus shocks significantly increase social distancing overall. These results suggest that measures taken to address economic dislocation caused by the COVID-19 pandemic helped address an important source of viral spread: social interaction.

Our empirical results are consistent with the theoretical model of compliance (see Section 2). The model considers an environment, in which the "shelter-in-place" ordinance is issued, and agents decide whether to comply with these orders or not. Agents have heterogeneous economic endowments (income or wealth), which capture the assets available to individuals as well as their exposure to economic shocks (e.g., from trade conflict). This maps directly to the empirical research design (see Section 3).[1] Then, we extend the setup to cover an environment in which agents receive additional information about the probability of becoming infected if they do not comply provided by a source (e.g., a political leader or a media) that might have a bias over the importance of compliance.

In equilibrium, there is a positive relationship between person's endowment and compliance: wealthier agents are more likely to comply with shelter-in-place protocols (Proposition 1). This is intuitive: for a low-income person, the same gain due to non-compliance results in a higher marginal utility. Aggregating at the county level, we demonstrate that richer counties feature higher compliance, which is fully supported by our empirical analysis. Empirically, we find a significant decline in population movement after the local shelter-in-place policies were enacted. As predicted by the theoretical model, counties with above median income comply with shelter-in-place policies by reducing movement by an additional 60% relative the baseline policy impact.

Though our primary focus is the economic determinants of compliance, we check robustness by accounting for political factors such as partisanship and slanted media exposure that may confound this relationship. Consistent with contemporary empirical work (Gadarian et al., 2020; Painter and Qiu, 2020; Barrios and Hochberg, 2020; Brzezinski et al., 2020), our model suggests that partisan media viewership will reduce compliance with shelter-in-place policies (Proposition 2). President Trump's supporters are more skeptical of mainstream media and federal government policies in general (Iyengar et al., 2019) and early statements by administration officials and some media dismissed or deflected the COVID-19 threat (Painter and Qiu, 2020; Bursztyn et al., 2020).[2] Our empirical results show that partisanship and exposure to slanted media may play a complementary yet distinct role in shaping public compliance with shelter-in-place orders.

A number of studies have identified other factors that influence compliance. Chan et al. (2020) show that European regions with higher confidence in the health care system are more likely to comply. In Barrios et al. (2020) the level of civic capital is an important factor in explaining voluntary compliance in Europe and the US. In addition, they show that, as states began loosening restrictions, social distancing compliance was more likely to remain steady in high civic capital counties, even when it was not mandated by the law. Brodeur et al. (2020) measure the effect of stay-at-home orders on compliance. This study employs U.S. cell phone data from Unacast, and social capital and trust data from the General Social

---

[1] In the empirical design, assets available immediately prior to the onset of COVID-19 are unknown. Instead, we rely on a baseline measure of income by county (2016). We account for various other factors that might offset these initial endowments (e.g., negative economic shocks due to the US trade war).

[2] In a theoretical paper, Gitmez et al. (2020) demonstrated that preferences for media consumption for those who are unlikely to comply with the shelter-in-place ordinances (e.g., because they live in a rural area where the risk of transmission is relatively low or just because they are poor) are skewed towards media sources that de-emphasize the threat.

Survey (GSS). Their main finding is that counties with more social trust decreased their mobility significantly more once a lockdown policy is implemented.

In this study, we have not attempted to account for economic *consequences* of shelter-in-place ordinance. A number of studies document and quantify negative consequences of shelter-in-place policies in terms of customer activity (Goolsbee and Syverson, 2020; Carvalho et al., 2020; Chetty et al., 2020; Coibion et al., 2020) and job losses (Friedson et al., 2020; Chudik et al., 2020; Gupta et al., 2020; Beland et al., 2020). Our results complement these papers and should be taken into account when formulating an optimal policy response to a major public health emergency.

Finally, this study advances a long standing literature on the relationship between income and strategic decision-making, including household financial choices as well as social compliance. Economic endowments, income uncertainty, and liquidity constraints impact risk preferences and self control (Jalan and Ravallion, 2001; Banerjee and Mullainathan, 2008; Guiso and Paiella, 2008; Tanaka et al., 2010; Dupas and Robinson, 2013; Bernheim et al., 2015; Carvalho et al., 2016). The evidence we present suggests that economic endowments influence substantive behavioral choices which pose a first-order challenge to public health interventions and broader government measures to stop the interpersonal transmission of COVID-19 and other similar risks.

The rest of the paper is organized as follows. Section 2 provides a basic theoretical model of compliance. Section 3 describes our data and introduces the research design. Section 4 reports the results, and Section 5 concludes. The Appendix contains data overview, and some additional results and robustness checks for the empirical analysis.

2. Theory

Our theoretical model considers an environment, in which the "shelter-in-place" ordinance is issued, and agents decide whether to comply with these orders or not. Agents have heterogeneous economic endowments. To simplify our analysis, we make no distinction between wealth and income. This should not affect the qualitative implications as agents with more assets naturally have better access to credit.

The other economic parameter that affects individuals' decisions on whether or not to comply with shelter-in-place is the income that an agent receives if she does not comply with ordinance. This parameter encompasses both the foregone income of compliance and possible fines for non-compliance. In addition, the decision to comply in the model is affected by the health risks agents encounter and the probability of getting the disease (e.g., depending on whether they live in a rural or urban community). At the same time, our model abstracts away from other reasons for non-compliance; they can be incorporated into our framework at the expense of cumbersome algebra. Finally, citizens decisions depend on the information that they receive. We assume that the information is provided by a media that might have a certain slant.

Setup. Consider a continuum of agents who have income of w_i; for each agent $i \in [0, 1]$, w_i is obtained from distribution $F(\cdot)$ over $[\underline{w}, \overline{w}]$, $\overline{w} > \underline{w} > 0$. The only decision that each agent makes is whether or not to comply with the "shelter in place" ordinance, $a_i \in \{c, n\}$, for "comply" and "not comply", respectively. If agent i complies, she consumes her income w_i. If she does not comply, she gets an additional income r, in which case her consumption is $w_i + r$, yet risks getting coronavirus. We assume that r incorporates, in addition to benefits of noncompliance, the expected costs, e.g., fines. The utility function of agent i is quasilinear: $u_i = \ln \left(w_i + r \mathbb{I}_{\{a_i = n\}} \right) - pH$, where p is the probability to get the disease, and $H > 0$ is the expected health damage to the infected person.

Agents' optimal decisions about compliance depends on their assessment of risk probabilities. There are two possible states of the world, $s \in \{C, N\}$, for "coronavirus threat" and "no threat" ; the *ex ante* probabilities are $P(C) = \theta_0 \leq \frac{1}{2}$ and $P(N) = 1 - \theta_0$. Agent i's utility depends on both the action and the state of the world as follows.

$$\begin{array}{l} \ln w_i \text{ if } a_i = c, \\ \ln (w_i + r) - H \text{ if } a_i = n, \ s = C, \\ \ln (w_i + r) \text{ if } a_i = n, \ s = N. \end{array}$$

Before making the decision, agents might get additional information from a source (e.g., a media outlet), which operates as follows. The source receives information about the true state of the world $s \in \{C, N\}$, and is committed to report a recommendation $\widehat{s} \in \{\widehat{c}, \widehat{n}\}$ for "comply" and "go".[3] If $s = N$, then the report is $\widehat{s} = \widehat{n}$. If $s = C$, then the report is $\widehat{s} = \widehat{n}$ with probability m and $\widehat{s} = \widehat{c}$ with probability $1 - m$. Thus, a source with low m is putting more emphasis on dangers of the virus, while high m corresponds to the desire to downplay the danger.

Analysis. We start by analyzing the determinants of compliance conditional on information that agents have. For now, let θ be an exogenous probability that the coronavirus threat is serious, i.e., the estimate of probability p that enters agents' utility function. Afterwards, we will assume that θ is a function of the *ex ante* beliefs θ_0 and a media signal; this will allow us to do comparative statics with respect to media slant.

---

[3] See Gentzkow and Shapiro (2008) for the basics of this approach to modeling media.

The expected utility of compliance is

$$Eu_i(c) = P(s = C) \ln w_i + P(s = N) \ln w_i$$
$$= \ln w_i.$$

The expected utility of non-compliance is

$$Eu_i(n) = \theta (\ln(w_i + r) - H) + (1 - \theta) \ln (w_i + r)$$
$$= \ln(w_i + r) - \theta H.$$

Now let $w^*(\theta)$ be a solution of the following equation.

$$\ln w = \ln(w + r) - \theta H. \tag{1}$$

Since $\ln (w + r) - \ln w$ is a decreasing function of w, there exists at most one such solution; the appropriate choice of $\underline{w}$ guarantees the existence of a solution.

The solution to Eq. (1), $w^* = w^*(\theta, H, r)$, is the critical threshold: agents with less income than w^* do not comply, while those with the income exceeding w^* shelter in place.

The comparative statics is straightforward. When the expected damage to the infected person, H, goes up, the threshold w^* goes down, i.e., more people comply with the ordinances. When the additional income r goes up (or, equivalently, fines for non-compliance go down), the threshold w^* goes up.

In addition to results about individual decisions as a function of income, one can do comparative statics with respect to the aggregate income distribution. For example, suppose that county 1 is wealthier than county 2; e.g., county 2 is more affected by a trade war. Mathematically, this corresponds to the distribution of income with cumulative density function $F_1(\cdot)$ first-order stochastically dominating the distribution with c.d.f. $F_2(\cdot)$: for any w, $F_1(w) \leq F_2(w)$ (Mas-Colell et al., 1995). As the threshold w^* is independent of the income distribution, there is a larger share of agents sheltering in place in the wealthier county 1.

Proposition 1 summarizes the above discussion.

Proposition 1. *(i) There exists a threshold $w^* = w^*(\theta, H, r)$ such that any agent i with $w_i > w^*$ shelters in place, $a_i = c$, while every agent i with $w_i < w^*$ does not comply, $a_i = n$.*

(ii) The lower is the agent's belief that there is no danger (the lower is θ), the higher is the threshold w^*, i.e., less people shelter in place. The higher is the cost of having the disease, H, or the lower is additional income that an agent can get by non-complying, r, the lower is the threshold w^*, i.e., more agents comply.

(iii) If county 1 is wealthier than county 2, i.e., distribution of income $F_1(\cdot)$ first-order stochastically dominates distribution $F_2(\cdot)$, then there is a larger share of sheltering in place in county 1.

Informational impact. Suppose that citizen i watched the news with the media slant m. (We use the media terminology, yet our analysis applies to *any* source of information, e.g., a national or local authority, a politician, etc.) If the media report is $\widehat{s} = \widehat{c}$, then she infers that the probability that the threat is serious is

$$\theta = \theta(\widehat{c}) = P(s = C | \widehat{s} = \widehat{c}) = 1.$$

The utility of compliance with orders when the signal is $\widehat{s} = \widehat{c}$:

$$Eu_i(c | \widehat{s} = \widehat{c}) = P(s = C | \widehat{s} = \widehat{c}) \ln w_i + P(s = N | \widehat{s} = \widehat{c}) \ln w_i = \ln w_i.$$

The expected utility of non-compliance is

$$Eu_i(n | \widehat{s} = \widehat{c}) = P(s = C | \widehat{s} = \widehat{c})(\ln(w_i + r) - H) + P(s = N | \widehat{s} = \widehat{c}) \ln (w_i + r)$$
$$= \ln(w_i + r) - H.$$

Now let $w^* = w^*(\widehat{c})$ be a unique solution of the following equation.

$$\ln w = \ln (w + r) - H.$$

Similarly, if the media report is $\widehat{s} = \widehat{n}$, there is a threshold $w^*(\widehat{n})$ that separates compliers and non-compliers when the media report is $\widehat{n}$.

In both cases, agents with income below $w^*(x)$, $x \in x\{\widehat{c}, \widehat{n}\}$ do not shelter in place, and those with income above the threshold comply. The comparative statics with respect to H, the cost of getting the disease, and r, the additional income, is similar to the previous case. Also, an increase in the *ex ante* probability of danger, θ_0, leads to a lower threshold $w^*(x)$, $x \in x\{\widehat{c}, \widehat{n}\}$.

With the media slant m increasing, the threshold $w^*(\widehat{n})$ goes down as well: with a media slanted against emphasizing the danger of the coronavirus, the signal $\widehat{s} = \widehat{n}$ is less persuasive. It is a straightforward exercise to show that $w^*(\widehat{n}) > w^*(\widehat{c})$, i.e., more people shelter in place when the report recommends to shelter $(\widehat{c})$, then when the report recommends not to $(\widehat{n})$.

Finally, how does the slant, m, affect the people behavior? Agents receive report $\widehat{s} = \widehat{c}$ with probability $\theta_0(1 - m)$ and report $\widehat{s} = \widehat{n}$ with probability $m\theta + 1 - \theta$. If the media report is $\widehat{s} = \widehat{c}$, then $\int_{w^*(\widehat{c})}^{\overline{w}} dF$ people shelter in place; if $\widehat{s} = \widehat{n}$, $\int_{w^*(\widehat{n})}^{\overline{w}} dF$

people do. To make the calculus tractable, let us assume that F is a uniform distribution over $[\underline{w}, \overline{w}]$. Then the expected number of non-compliers is as follows.

$$En(m) = \theta_0 (1 - m) \frac{1}{\overline{w} - \underline{w}} (w^*(\widehat{c}) - \underline{w}) + (m\theta_0 + 1 - \theta_0) \frac{1}{\overline{w} - \underline{w}} (w^*(\widehat{n}) - \underline{w}).$$

There are two effects of media slant on the expected number of non-compliers $En(m)$. First, when m is higher, the probability that the report is $\widehat{s} = \widehat{c}$ is lower, and $w^*(\widehat{c})$ does not change with m. As $w^*(\widehat{n}) > w^*$ for any m, with a higher probability of report $\widehat{s} = \widehat{n}$, the number of non-compliers increases with slant m. At the same time, conditional on report $\widehat{s} = \widehat{n}$, a higher level of slant m leads to more sheltering in place, as citizens take the slant into account. In general, the outcome of an increase in m is ambiguous. However, as $En(m)$ is a differentiable function, and $\frac{c}{cm} En(m) > 0$ in the neighborhood of $m = 0$, there is a range of parameters for which more biased reports result in a higher noncompliance.

We summarize the discussion about the impact of information in Proposition 2.

Proposition 2. *Suppose that agents receive their information through media with slant m.*

(i) There exist two thresholds, $w^*(\widehat{c})$ and $w^*(\widehat{n})$, such that any agent i with $w_i > w^*(x)$ shelters in place, $a_i = c$, upon receiving signal $x \in \{\widehat{c}, \widehat{n}\}$ while every agent i with $w_i < w^*(x)$ does not comply ($a_i = n$). More people comply when the media report is $\widehat{c}$, i.e., $w^*(\widehat{c}) < w^*(\widehat{n})$.

(ii) Suppose, additionally, that the distribution $F(\cdot)$ is uniform. There exists a threshold $\overline{m}$, $0 < \overline{m} \le 1$, such that the expected number of non-compliers increases in media bias towards minimizing the threat, m, over $[0, \overline{m}]$.

It is a straightforward exercise to derive comparative statics results similar to those in Proposition 1 in the setup with media. The only difference is that we would need to state all the results for each of the two thresholds corresponding to different media signals. For example, the lower is the agent's *ex ante* belief that there is no danger (the lower is θ_0), the higher are both thresholds $w^*(x)$, $x \in \{\widehat{c}, \widehat{n}\}$, i.e., less people shelter in place. The higher is the cost of having the disease, H, or the lower is additional income that an agent can get by non-complying, r, the lower are both thresholds $w^*(x)$, $x \in \{\widehat{c}, \widehat{n}\}$, i.e., more agents comply.

Discussion. Our model does not differentiate agents by their ability to generate income by non-complying. This might look unrealistic as there is a high correlation between wealth and earning capacity. However, what matters for the compliance decision is the relative marginal income due to non-compliance which need not to be correlated with wealth. It is well-documented that the relative earning losses for the high-earners who worked at home were lower than that of the low-earners (Cajner et al., 2020). This makes relative marginal income of the high-earners lower than that of the low-earners.

In our model, the media is not strategic. However, it is straightforward to extend the model to the case in which the slant m is determined endogenously (e.g., Gehlbach and Sonin, 2014). The same model can be used to incorporate strategic censorship by the government. Gitmez et al. (2020) models individuals' demand for information slant.

Perhaps the most significant simplifying assumption of our basic model is that the probability of the environment being dangerous to an individual, θ, is independent of the share of people who decide, in equilibrium, to maintain social distancing and shelter-in-place. It could be argued that this assumption exaggerates the county-by-county level of unpredictability, the danger caused by out-of-locality travelers ("superspreaders"), and the like.[4] At the cost of more algebra, our model could be extended to incorporate an endogenous probability of becoming infected, but those extensions are not necessary to motivate our core empirical results.[5]

3. Data and research design

We study changes in population movement ("social distancing") after the onset of local shelter-in-place ordinances in the United States. Cellphone location data enables us to quantify population movement within and across origin counties by day. Patterns of cellphone use vary across counties. To standardize measures of population movement, all data is compared to a local day-of-week base rate calculated using data prior to the onset of COVID-19 in the United States (March 8). The outcome of interest is county-day variation in standardized population movement. This is illustrated in Fig. 1 Panel A. We provide a detailed overview of the data used in our analysis in the Appendix. Where individuals comply with local shelter-in-place laws, we observe a reduction in intersecting sets (less distance traveled on average). Although population movement declined by approximately 25% (national average) between March 8 and May 1, there is high variance in movement (Fig. 1 Panel B). This suggests variation in compliance with government ordinances. The secondary quantity of interest is how compliance varies across counties with differing economic endowments.

---

[4] See, for example, Dave et al. (2020a).

[5] That would require a new definition of equilibrium that will include the share of people who do not comply, then probability θ being a function of the expected share of non-compliers. By the Brouwer's fixed point theorem, there will be such θ that will be both a function of the equilibrium share of non-compliers and result in every agent satisfying her incentive compatibility constraints. Gitmez et al. (2020) demonstrate that our model can be extended to incorporate the endogenous probability of getting the infection.
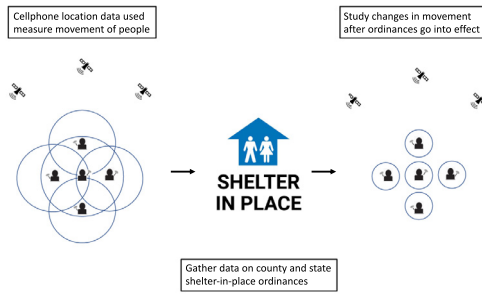
(a) Measurement Approach.
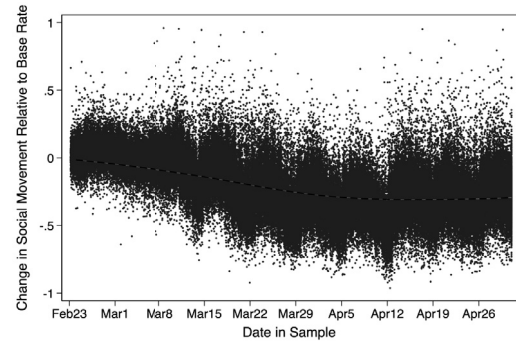
(b) Variation in Social Distancing.

Fig. 1. Design for quantifying and studying localized population movement using cellphone data. *Notes*: Panel (A): Cellphone location data is gathered passively and used to measure population movement by origin county and day. Intersecting circles indicate social proximity. Data is gathered on timing of shelter-in-place ordinances. Study assesses change in travel patterns which indicate social distancing (reduced social contact). Panel (B): Variation in cellphone-derived movement data over study period (February 23 to May 1, 2020). Base rate for daily movement after March 8 (onset of COVID-19 in United States) is history of day-of-week movement data collected in prior periods. Local polynomial regression indicates 25% decline (national average) in movement during study period.
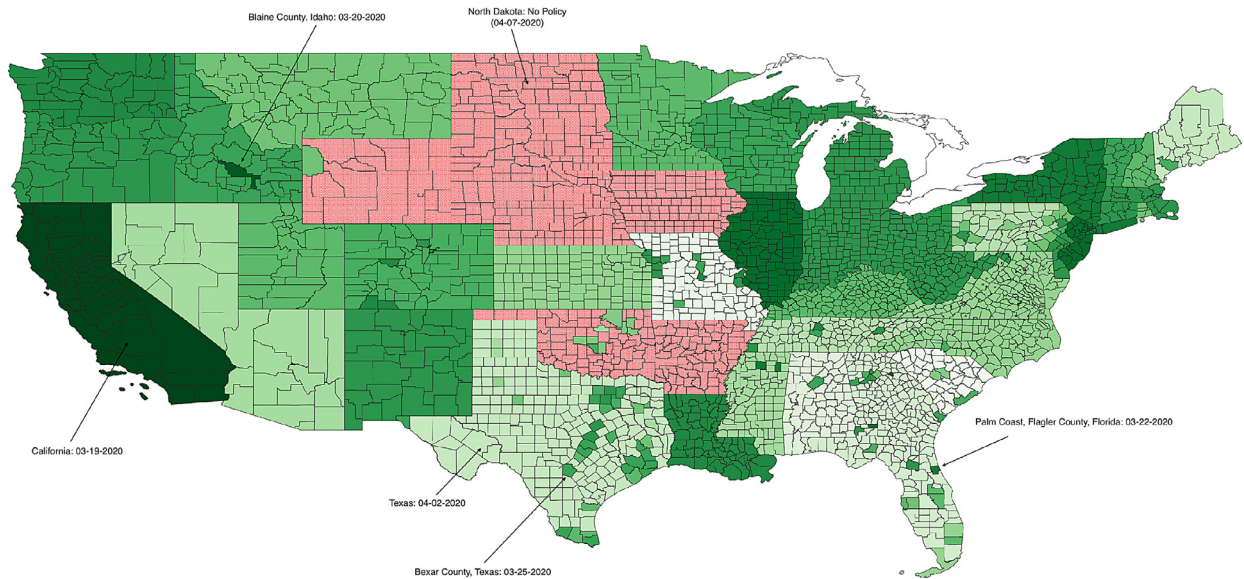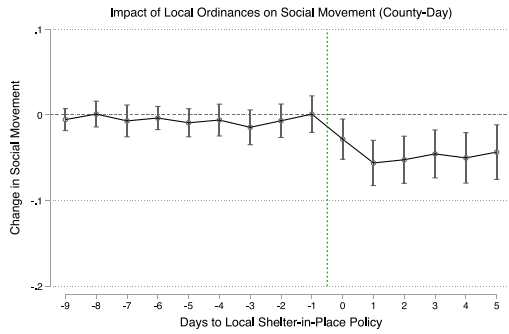

Fig. 2. Staggered introduction of shelter-in-place ordinances. *Notes*: Map of staggered introduction of shelter-in-place policies across the United States. Darker shades of green indicate earlier policy dates, beginning on March 19 (California). Localized policies were adopted in more than 125 counties. Data on state-level policies corrected for additional policies up until April 7. Data on county-level policies is from Painter and Qiu (2020). Additional details in Appendix. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.1. Research design

We employ a difference-in-difference design to estimate the impact of the shelter-in-place ordinances on population movement (Bertrand et al., 2004). Introduction of social distancing policies was staggered across states in time (Fig. 2). More than 125 counties adopted policies prior to state-wide mandates. We leverage this data to examine how localized population movement changes after the introduction of a county-level policy (either via county- or state-wide mandates). Local ordinances are staggered across states and, in a number of cases, counties within states that have not yet enacted state-wide mandates.

To account for this staggered variation in the introduction of the natural experiment (local policy change), we include county and day fixed effects (Goodman-Bacon, 2018). This partials out any variation in population movement
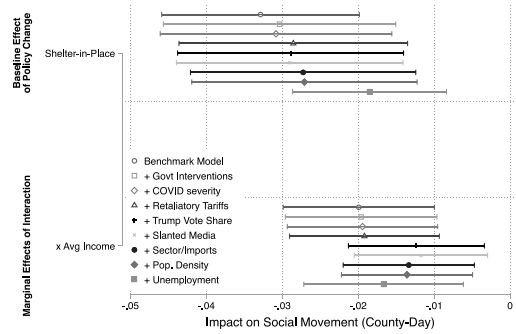
(a) Event Study Results.   (b) Staggered Diff-in-Diff Results.

Fig. 3. Staggered introduction of shelter-in-place ordinances and localized impact on population movement. *Notes*: Panel (A): Event study design using leads and lags of policy change to assess pre-treatment and post-treatment changes in population movement. No evidence of anticipation effects. Substantive and stable declines in movement estimated after the first full day of shelter-in-place. 90% confidence intervals reported. Panel (B): Difference-in-differences results for estimated impact of shelter-in-place. Baseline effect of policy plot in upper left. Heterogeneous (marginal) effects plotted below. Additional coefficient plots represent sequential regression models with additional control variables and fixed effects.

(dependent variable) that is correlated with factors that remain fixed regarding counties during the sample period (February 23 to May 1) or changes in mobility due to nation-wide shifts in policies or protocols ('common shocks').[6] These factors include, for example, population density, industrial sector shares, and federal mandates. Our benchmark specification also incorporates state × week fixed effects, which account for differential trends in population movement and policy interventions by state across time. We produce Fig. 3 Panel B by studying Eq. (2):

$$
\begin{aligned}
y_{c,d} = \alpha &+ \beta_1 active\_policy_{c,d} + \beta_2 active\_policy_{c,d} \times income\_above\_median_c^{2016} \\
&+ \omega X_{c,d} + \phi \theta_d \times X_c + \lambda_c + \theta_d + \kappa \theta_d \times income\_above\_median_c^{2016} \\
&+ \zeta_c state \times week_d + \epsilon_{c,d}
\end{aligned}
\tag{2}
$$

where $y_{c,d}$ is the measure of population movement derived from cellphone location data, standardized using location-specific day-of-the-week trends prior to March 8, 2020. *active_*k $policy_{c,d}$ indicates whether a county-specific or state-wide shelter-in-place mandate is active on a given day. We estimate the marginal effects using a threshold indicator for whether a given county is above the median level of income in 2016 ($\times income\_above\_median_c^{2016}$). λ_c indicates county-level fixed effects, θ_d indicates day-level fixed effects, and $\zeta_c state \times week$ denotes state-by-week specific fixed effects. $\theta_d \times income\_above\_median_c^{2016}$ is the day-specific effect of the median income measure. This accounts for any linear and non-linear trends in movement across counties with varying levels of income by day. This enables us to interpret the main effects as a triple difference (difference-in-difference-in-differences) design since the baseline shift is otherwise collinear with parameters in the main model (due to staggered policy introduction across counties and states) (Goodman-Bacon, 2018). $\theta_d \times X_c$ and $X_{c,d}$ are a vector of control variables described in the main text (see (a) through (h) in main text). The unit-specific time-invariant characteristics are interacted with the day-level fixed effects to allow these effects to vary flexibly across the sample period. Since the benchmark specification includes state×week specific fixed effects, we flexibly account for any state-wide trends over time. Heteroskedasticity-robust standard errors are clustered by county, day, and state × week.

The difference-in-differences design provides a causal estimate of compliance with shelter-in-place policies so long as trends in the outcome of interest in untreated counties represent a valid counterfactual for treated counties (Donald and Lang, 2007). We assess the plausibility of this identifying assumption using an event study design. This approach leverages leads and lags of local shelter-in-place policies to study variation the outcome variable prior to and following the policy change.[7] We study Eq. (3) using 10 window prior to the leads as the base period:

$$
\begin{aligned}
y_{c,d} = \alpha &+ \sum_{i=-1}^{-9} (\beta_i active\_policy_{c,t+i}) + \beta_0 active\_policy_{c,t=0} \\
&+ \sum_{i=1}^{5} (\beta_i active\_policy_{c,t+i}) + \omega X_{c,d} + \lambda_c + \theta_d + \epsilon_{c,d}
\end{aligned}
\tag{3}
$$

Notation follows the difference-in-differences design. One exception is that $active\_policy_{c,t+i}$ coincides with the date of the shelter-in-place policy change, not the policy status (0 or 1) itself. For consistency, heteroskedasticity-robust standard errors

---

[6] We choose May 1 as the end date for the sample due to a large number of states scaling back local mandates on or around this date. Using differential end dates would lead to significant imbalance in the panel, though the estimated effects are consistent with this alternative approach as well.

[7] It also enables us, when we estimate separate event studies across median thresholds, to assess the validity of this assumption specifically for the marginal effect of interest ($\times income\_above\_median_c^{2016}$).
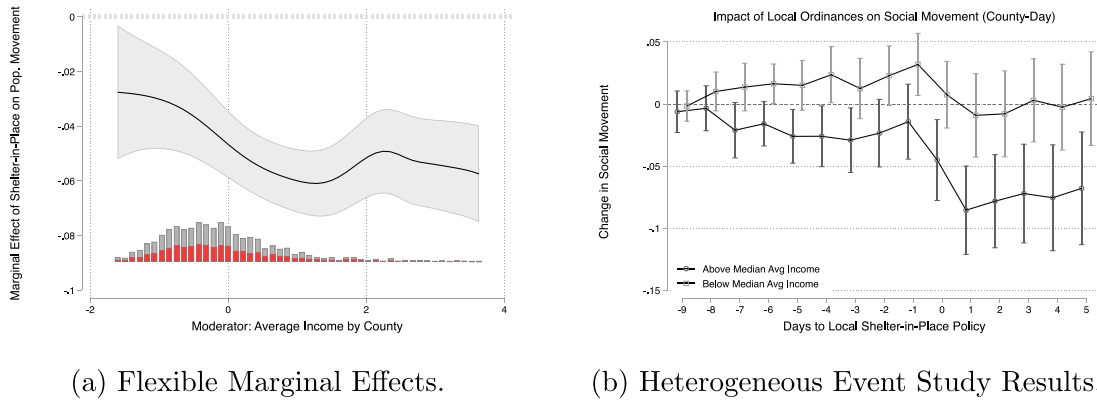
(a) Flexible Marginal Effects.   (b) Heterogeneous Event Study Results.

Fig. 4. Event study designs indicate varying compliance across local income thresholds. *Notes*: Panel (A): Flexible marginal effects of average income in standardized levels. These effects were estimated using the *interflex* package compiled by Hainmueller et al. (2019). Distribution trimmed above and below for top percentile (98% of sample remaining). Panel (B): Event study design using leads and lags of policy change in above versus below the median level of income by county. Event study results suggest large reduction population movement in above median counties; no change in below median counties. 90% confidence intervals reported.

are calculated using a three way clustering approach following the main effects in Eq. (2).[8] We calculate total effects in the threshold-separated tests in Fig. 4 Panel C via a split-sample approach as suggested by Goodman-Bacon (2018).[9]

4. Results

In this section, we present main results as well as a battery of robustness checks (Section 4.2). We present supplemental analyses, including the impact of the 2020 CARES Act, in the final Section 4.3.

4.1. Main effects

We find that shelter-in-place policies reduce movement and that these changes are influenced by economic endowments. Table A-1 presents the benchmark results and first robustness checks. Visualization of these point estimates is presented in Fig. 3 Panel B. First, on average, we find a significant decline in population movement after the local shelter-in-place policies were enacted ($\widehat{\beta}_{active\_policy} = -0.032$, $p = 0.001$). Second, counties with above median income comply with shelter-in-place policies by reducing movement by an additional 60% relative the baseline policy impact ($\widehat{\beta}_{income} = -0.0199$, $p = 0.002$). These results are also presented in Table A-1.

We assess the 'common trends' assumption in Fig. 3 Panel A using the event study design outlined above. Notice that the effect of shelter-in-place policies is statistically indistinguishable from zero prior to the onset of local mandates (green line). Population movement significantly declines after the first full day of shelter-in-place ($\widehat{\beta}_{t+1} = -0.052$, $p = 0.002$) and remains stable in subsequent days. This provides baseline support for a causal interpretation of our statistical estimates.

4.2. Robustness checks

We next consider a number of robustness checks, sequentially adding additional control models to the benchmark specification. The timing of shelter-in-place policies could be **(a)** correlated with related government interventions, including school closures, bans on social gathering, restrictions on eviction notices, and utility cutoffs. The onset of shelter-in-place might also be driven by **(b)** the severity of COVID-19 (cumulative cases and deaths). We add these parameters to the benchmark model and results remain consistent (square and diamond symbols respectively).

Next, we account for the correlation between local economic endowments and **(c)** exposure to the ongoing trade war (Fajgelbaum et al., 2020), **(d)** political partisanship, and **(e)** exposure to slanted media coverage of the COVID-19 pandemic. To study **(c)**, we rely on data from Fajgelbaum et al. (2020) measuring exposure to retaliatory tariffs. In **(d)**, we rely on county-level data on Republican vote share in the 2016 elections. In **(e)**, we gather data on Designated Market Area (DMA) penetration by the Sinclair Broadcast Group and focus on their Fox News affiliates. We then capture trends in population movement correlated with these factors through the use of interacted day fixed effects; this approach flexibly accounts for potentially non-linear effects during our sample. These factors are sequentially added to the model and results are indicated in Fig. 3 Panel C with triangle, plus, and grey cross symbols respectively.

---

[8] Results are similar if we calculate standard errors at the state level instead.

[9] We confirm the divergent compliance patterns in the split-sample design correspond to marginal effects in a fully interacted event study design.

It is also possible that local economic endowments are tied to **(f)** sector-specific dynamics, specifically agricultural and manufacturing industries. These local labor conditions might also make counties more likely to be exposed to import tariffs, which undermine access to foreign suppliers. We allow these effects to vary by day during our sample and include these parameters in the model, indicated by a solid circle symbol. We also **(g)** incorporate measures of county-specific population density and allow these effects to vary by day. These effects are noted by a solid grey diamond.

Finally, we account for **(h)** the potential correlation between income and unemployment. We do this in two ways. First, we flexibly account for unemployment during the period when our income measure is produced using interacted day-specific effects. Second, we track real time demand for information about how to file for unemployment via geocoded Google search activity (Choi and Varian, 2009). We compile DMA-day specific search intensity and use the benchmarking procedure introduced by Goldsmith-Pinkham and Sojourner (2020). We then calculate the cumulative search intensity for each region and assign this value to each corresponding county-day in our sample until the first surge in unemployment declines (April 1). We introduce these measures to the model and the corresponding main effects are denoted by solid grey square. In the most saturated model **(h)**, the main effects remain substantively meaningful and statistically precise ($\widehat{\beta}_{s-i-p} = -.0185$, $p = 0.004$; $\widehat{\beta}_{income} = -.0167$, $p = 0.02$).

4.3. Supplemental analyses

In Fig. 4 we conduct several additional analyses. In Panel A, we flexibly estimate the marginal effect of shelter-in-place policies across average income in standardized levels (with mean = 0, standard deviation = 1). There is clear evidence of an increasing marginal effect (in absolute terms) until approximately two standard deviations above the mean, where the marginal effect levels off. This suggests that economic endowments reach a saturation effect at this level in the distribution, where an additional dollar of income does not cause an additional marginal decrease population movement. Using a median-based measure absorbs this dynamic by calculating an average across counties above the average threshold value. If, instead we relied on a linear specification of the marginal effect (using average income in standardized levels), the slope might be attenuated by counties with very high income levels. We present these results in Fig. A-1 (and Table A-2). Notice that the results are highly stable in the linear specification as well, though the interpretation of the magnitude of the effects changes (to standard deviation shifts in income). In Panel B, we introduce an event study design that leverages the median split of income. Population movement in counties above the median in income decreases sharply the day after policies go into effect ($\widehat{\beta}_{t+1}^{above\_income} = -0.084$, $p = 0.001$), but the population movement below the median remains unaffected ($\widehat{\beta}_{t+1}^{below\_income} = -0.002$, $p = 0.938$).

Next, we take advantage of a large-scale economic stimulus program in the United States to estimate the impact of local transfers on social distancing patterns. The US CARES Act authorized approximately 250 billion USD to be transferred to more than 127 million taxpayers. These transfers were scaled to the gross taxable income of each recipient, with 1200 USD, 2400 USD, and 500 USD in maximum benefits per individual, married couple, and dependent child. Transfers were decreasing with taxable income. This means relatively poorer individuals received the largest benefits in absolute terms.

We use zipcode-level data from Facteus across four years to identify trends in federal bank deposits. For stimulus recipients that used an electronic banking transfer for their tax returns and payments, transfers were submitted electronically as a trackable deposit. We observe a spike in federal transfers around April 10, the first designated date when stimulus checks were electronically (and physically) allocated. We match zipcode-level transfers to the likely origin county and generate a cumulative measure of stimulus transfers by county-day, which we standardize in per capita terms. The recipient-level roll out of the stimulus program was haphazard with some beneficiaries receiving financial transfers weeks ahead of others. We take advantage of this quasi-random variation in cumulative transfers at the county level to study how local economic transfers (positive income shocks) influence social distancing.

We estimate a model consistent with Eq. (2), where the estimated effect of the stimulus transfers is now the primary quantity of interest. These results are presented in Fig. 5. Complete results are shown in Table A-3. In the benchmark model, every additional USD per capita reduces movement by approximately 1.1%. An effect of this magnitude is statistically indistinguishable from the marginal increase in compliance with local mandates associated with an above median income by county. Moreover, this effect remains robust when we replicate the sensitivity analyses presented earlier (see (a) through (g)).[10] These results suggest that a widespread economic stimulus program that targeted economically vulnerable individuals and communities meaningfully increased social distancing overall.

Taken together, these results suggest local income levels meaningfully influence social distancing during the COVID-19 pandemic. Although the event study results provide support for the common trends assumption, it is also possible that our dynamic estimates from the fully saturated benchmark model do not fully account for potential sources of bias. Even if the results presented in the supplemental analysis represent descriptive rather than causal evidence, the role that poverty plays in shaping whether and how communities engage in social distancing during the COVID-19 pandemic remains a first order public policy challenge.

---

[10] We cannot implement (h) since the period of data collection for the real time unemployment trends does not overlap with the period during which stimulus checks were being distributed.
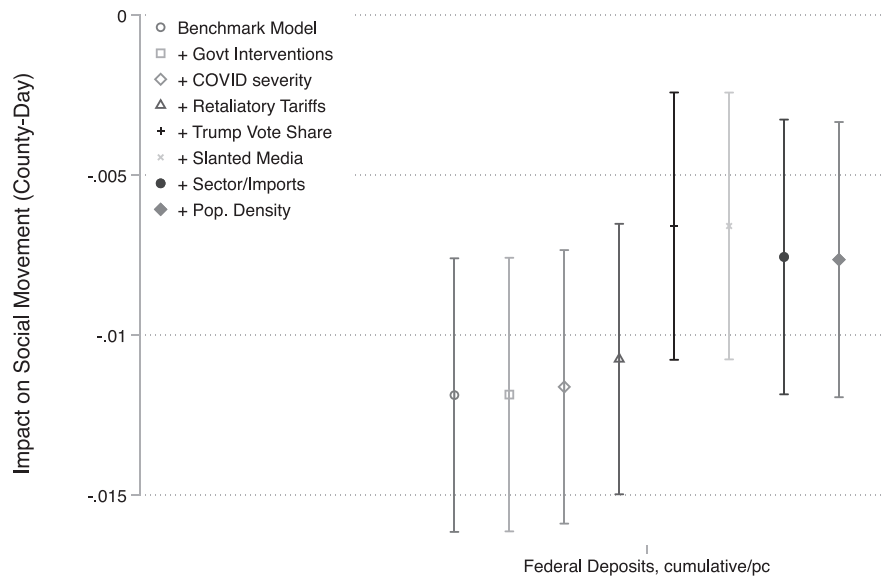
Fig. 5. Staggered roll out of stimulus checks and localized impact on population movement. *Notes*: Difference-in-differences results for estimated impact of federal deposits on social movement. Additional coefficient plots represent sequential regression models with additional control variables and fixed effects. 90% confidence intervals reported.

5. Conclusion

Shelter-in-place policies and related 'lockdown' quarantine orders are among the most economically costly and robust measures governments have taken to combat the COVID-19 pandemic. Our evidence demonstrates that these policies effectively reduce population movement overall and, as evidence in medicine and epidemiology indicate, slow the spread of COVID-19 by reducing interpersonal transmission risks. Compliance with shelter-in-place directives is individually costly and requires behavioral changes across diverse sub-populations. Variation in compliance, however, represents a local and regional risk to public health. Clarifying how income and other economic factors influence compliance provides actionable insights for policy makers and public health officials. We provide robust evidence that economic endowments significantly influence when and by how much communities engage in social distancing. We also demonstrate that a large-scale stimulus program significantly increase social distancing, consistent with our main results.

Our evidence suggests that at-risk, impoverished communities exposed to economic dislocation are the least likely to comply with shelter-in-place policies. A number of related studies have demonstrated the substantial economic shocks caused by the pandemic and public policy responses. These conditions have created new or enhanced existing forms of economic hardship. Given the prevalence of comorbidity factors and poor health insurance coverage among these populations (Akinyemiju et al., 2016), the risks of continued exposure are substantial. These communities are also more likely to lack sufficient testing and contact tracing capacity to monitor and thwart COVID-19 outbreaks. This means any spread due to non-compliance within these communities will be difficult to detect until a potentially large number of symptomatic infections or deaths occur. Consistent with our findings, targeted economic relief such as direct stimulus transfers and increased unemployment benefits may have limited potential spread of COVID-19 among economically disadvantaged populations.

References

Akinyemiju, T., Jha, M., Moore, J.X., Pisu, M., 2016. Disparities in the prevalence of comorbidities among us adults by state medicaid expansion status. Prevent. Med. 88, 196–202.

Allcott, H., Boxell, L., Conway, J., Ferguson, B.A., Gentzkow, M., Goldman, B., 2020. What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? Working Paper, 27965. doi:10.3386/w27965.

Anderson, R.M., Heesterbeek, H., Klinkenberg, D., Hollingsworth, T.D., 2020. How will country-based mitigation measures influence the course of the Covid-19 epidemic? The Lancet 395 (10228), 931–934.

Angrist, J.D., Pischke, J.-S., 2009. Mostly Harmless Econometrics. Princeton University Press, Princeton, NJ.

Bai, Y., Yao, L., Wei, T., Tian, F., Jin, D.-Y., Chen, L., Wang, M., 2020. Presumed Asymptomatic Carrier Transmission of COVID-19. JAMA Research Letter, Published Online February 21, 2020.

Baker, S.R., Bloom, N., Davis, S.J., Terry, S.J., 2020. Covid-Induced Economic Uncertainty. (preprint), Posted April 4

Banerjee, A.V., Mullainathan, S., 2008. Limited attention and income distribution. Am. Econ. Rev. 98 (2), 489–493.

Barrios, J. M., Benmelech, E., Hochberg, Y. V., Sapienza, P., Zingales, L., 2020. Civic Capital and Social Distancing During the Covid-19 Pandemic. NBER Working Papers (27320).

Barrios, J.M., Hochberg, Y.V., 2020. Risk Perception Through the Lens of Politics in the Time of the Covid-19 Pandemic. Becker Friedman Institute. (preprint), Posted April 6

Beland, L.-P., Brodeur, A., Haddad, J., Mikola, D., 2020. Covid-19, family stress and domestic violence: Remote work, isolation and bargaining power. GLO Discussion Paper.

Bernheim, B.D., Ray, D., Yeltekin, Ş., 2015. Poverty and self-control. Econometrica 83 (5), 1877–1911.
Bertrand, M., Duflo, E., Mullainathan, S., 2004. How much should we trust differences-in-differences estimates? Q. J. Econ. 119 (1), 249–275.
Brodeur, A., Grigoryeva, I., Kattan, L., 2020. Stay-at-Home Orders, Social Distancing and Trust. 553. https://ideas.repec.org/p/zbw/glodps/553.html.
Brzezinski, A., Deiana, G., Kecht, V., Van Dijcke, D., 2020. The Covid-19 Pandemic: Government vs. Community Action Across the United States. Covid Economics: Vetted and Real-Time Papers 7, 115–156.
Bursztyn, L., Rao, A., Roth, C., Yanagizawa-Drott, D., 2020. Misinformation During a Pandemic. Becker Friedman Institute. (preprint), Posted April 17
Cajner, T., Crane, L. D., Decker, R. A., Grigsby, J., Hamins-Puertolas, A., Hurst, E., Kurz, C., Yildirmaz, A., 2020. The U.S. Labor Market During the Beginning of the Pandemic Recession. NBER Working Papers (27159).
Carvalho, B., Peralta, S., Pereira dos Santos, J., 2020. What and How Did People Buy During the Great lockdown? Evidence From Electronic Payments. ULB Working Paper.
Carvalho, L.S., Meier, S., Wang, S.W., 2016. Poverty and economic decision-making: evidence from changes in financial resources at payday. Am. Econ. Rev. 106 (2), 260–284.
Chan, H.F., Brumpton, M., Macintyre, A., Arapoc, J., Savage, D. A., Skali, A., Stadelmann, D., Torgler, B., 2020. How confidence in health care systems affects mobility and compliance during the COVID-19 pandemic. CREMA Working Paper Series. https://ideas.repec.org/p/cra/wpaper/2020-11.html.
Chen, S., Yang, J., Yang, W., Wang, C., Bärnighausen, T., 2020. Covid-19 control in china during mass population movements at new year. The Lancet 395 (10226), 764–766.
Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., et al., 2020. How did Covid-19 and stabilization policies affect spending and employment? A new real-time economic tracker based on private sector data. NBER Working Paper (27431).
Choi, H., Varian, H., 2009. Predicting Initial Claims for Unemployment Benefits. Google Inc, pp. 1–5.
Chudik, A., Pesaran, M. H., Rebucci, A., 2020. Voluntary and mandatory social distancing: evidence on Covid-19 exposure rates from chinese provinces and selected countries. NBER Working Paper (27039).
Coibion, O., Gorodnichenko, Y., Weber, M., 2020. The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. NBER Working Paper (27141).
Dave, D., Friedson, A., McNichols, D., Sabia, J., 2020a. The Contagion Externality of a Superspreading Event: The Sturgis Motorcycle Rally and Covid-19. IZA DP No. 13670 (preprint).
Dave, D. M., Friedson, A. I., Matsuzawa, K., Sabia, J. J., Safford, S., 2020b. Were Urban Cowboys Enough to Control Covid-19? Local Shelter-in-Place Orders and Coronavirus Case Growth. NBER Working Paper (27229).
Donald, S.G., Lang, K., 2007. Inference with difference-in-differences and other panel data. Rev. Econ. Stat. 89 (2), 221–233.
Dupas, P., Robinson, J., 2013. Why don't the poor save more? Evidence from health savings experiments. Am. Econ. Rev. 103 (4), 1138–1171.
Fajgelbaum, P., Goldberg, P.K., Kennedy, P.J., Khandelwal, A.K., 2020. The return to protectionism. Q. J. Econ. 135 (1), 1–55.
Fowler, J. H., Hill, S. J., Obradovich, N., Levin, R., 2020. The Effect of Stay-at-Home Orders on Covid-19 Cases and Fatalities in the United States. medRxiv.
Friedson, A. I., McNichols, D., Sabia, J. J., Dave, D., 2020. Did Californias Shelter-in-Place Order Work? Early Coronavirus-Related Public Health Effects. NBER Working paper.
Gadarian, S., Goodman, S. W., Pepinsky, T. B., 2020. Partisanship, Health Behavior, and Policy Attitudes in the Early Stages of the Covid-19 Pandemic. SSRN (preprint), Posted March 27.
Gaventa, J., 1982. Power and Powerlessness: Quiescence and Rebellion in an Appalachian Valley. Appalachian Studies. University of Illinois Press, Urbana and Chicago.
Gehlbach, S., Sonin, K., 2014. Government Control of the Media. J. Publ. Econ. 118 (C), 163–171. doi:10.1016/j.jpubeco.2014.07.
Gentzkow, M., Shapiro, J.M., 2008. Competition and Truth in the Market for News. J. Econ. Perspect. 22 (2), 133–154.
Gitmez, A., Sonin, K., Wright, A., 2020. Political Economy of Crisis Response. Mimeo.
Goldsmith-Pinkham, P., Sojourner, A., 2020. Predicting Initial Unemployment Insurance Claims using Google Trends. Working Paper (preprint), Posted April 3.
Goodman-Bacon, A., 2018. Difference-in-Differences With Variation in Treatment Timing. NBER Working Paper (25018).
Goodman-Bacon, A., Marcus, J., 2020. Using Difference-in-Differences to Identify Causal Effects of Covid-19 Policies. DIW Berlin Discussion Paper No. 1870.
Goolsbee, A., Syverson, C., 2020. Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020. NBER Working Paper (24732).
Gormsen, N. J., Koijen, R. S., 2020. Coronavirus: Impact on Stock Prices and Growth Expectations. Becker Friedman Institute (preprint), Posted April 8.
Guiso, L., Paiella, M., 2008. Risk aversion, wealth, and background risk. J. Eur. Econ. Assoc. 6 (6), 1109–1150.
Gupta, S., Montenovo, L., Nguyen, T. D., Rojas, F. L., Schmutte, I. M., Simon, K. I., Weinberg, B. A., Wing, C., 2020. Effects of Social Distancing Policy on Labor Market Outcomes. NBER Working Paper.
Hainmueller, J., Mummolo, J., Xu, Y., 2019. How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice. Polit. Anal. 27 (2), 163–192.
Hsiang, S., Allen, D., Annan-Pan, S., Bell, K., Bolliger, I., Chong, T., Druckenmiller, H., Hultgren, A., Huang, L.Y., Krasovitch, E., Lau, P., Lee, L., Rolf, E., Tseng, J., Wu, T., 2020. The Effect of Large-Scale anti-Contagion Policies on the Coronavirus (Covid-19) Pandemic. MedRxiv (preprint), Posted March 31 2020. doi:10.1101/2020.03.22.20040642.
Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., Westwood, S.J., 2019. The origins and consequences of affective polarization in the United States. Annu. Rev. Polit. Sci. 22 (1), 129–146. doi:10.1146/annurev-polisci-051117-073034.
Jalan, J., Ravallion, M., 2001. Behavioral responses to risk in rural china. J. Dev. Econ. 66 (1), 23–49.
Kapoor, R., Rho, H., Sangha, K., Sharma, B., Shenoy, A., Xu, G., 2020. God is in the Rain: The Impact of Rainfall-Induced Early Social Distancing on Covid-19 Outbreaks. Available at SSRN 3605549.
Levi, M., 1997. Consent, Dissent, and Patriotism. Political Economy of Institutions and Decisions. Cambridge University Press, New York.
Lewnard, J.A., Lo, N.C., 2020. Scientific and ethical basis for social-distancing interventions against Covid-19. Lancet Infect. Dis. 20, 631–633.
Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. Microeconomic Theory. Oxford University Press.
Matrajt, L., Leung, T., 2020. Evaluating the Effectiveness of Social Distancing Interventions Against Covid-19. medRxiv (preprint).
Painter, M. O., Qiu, T., 2020. Political Beliefs Affect Compliance With Covid19 Social Distancing Orders. SSRN (preprint), Posted April 4.
Scott, J.C., 1985. Weapons of the Weak. Yale University Press, New Haven.
Stapleton, K., 2020. News Release, April 9 2020. US Department of Labor.
Tanaka, T., Camerer, C.F., Nguyen, Q., 2010. Risk and time preferences: linking experimental and household survey data from Vietnam. Am. Econ. Rev. 100 (1), 557–571.
Tyler, T.R., 1990. Why People Obey the Law. Yale University Press, New Haven.
Viner, R.M., Russell, S.J., Croker, H., Packer, J., Ward, J., Stansfield, C., Mytton, O., Bonell, C., Booy, R., 2020. School closure and management practices during coronavirus outbreaks including Covid-19: a rapid systematic review. Lancet Child Adolesc. Health doi:10.1016/S2352-4642(20)30095-X.