Social Repercussions of Pandemics

Philip Barrett*, Sophia Chen†‡

December 29, 2021

Abstract

Historical accounts suggest that epidemics may have social scarring effects, increasing the likelihood of social unrest. We investigate this view systematically in two ways: using daily timing of disease outbreak and unrest during the COVID-19 pandemic; and in a dynamic monthly panel 130 countries over the short- and medium-run. In both cases, the likelihood of unrest is lower in the immediate aftermath of epidemics. We show that limits on social interaction was an important driver of this result during COVID. In the broader sample, epidemics are are unique among disasters in reducing the likelihood of unrest, consistent with disease-specific limits on social contact. We argue that the opposite narrative from historical accounts likely derives from a non-systematic selection of cases and long-run relationships.

JEL-Codes: I15, H56, N40

Keywords: Epidemics, Social Unrest, COVID-19

*International Monetary Fund, Research Department, email: pbarrett@imf.org

†International Monetary Fund, Research Department, email: ychen2@imf.org

‡We thank Antonio Spilimbergo, Deniz Igan and conference and seminar participants at the Executive Council on Diplomacy, IMF, United Nations, University of Cambridge Money, Macro, and Finance Society Conference, University of Essex, University of Leicester, and World Bank for their helpful comments and suggestions. We thank Luisa Calixto for excellent research assistance. The views expressed here are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

1 Introduction

In 1832, the cholera epidemic hit Paris. In just a few months, the disease killed almost 20,000 of the city's 650,000 population (Francke and Korevaar (2021)). Most fatalities occurred in the heart of the city where many poor workers lived in squalid conditions, drawn to Paris by the Industrial Revolution. The spread of the disease heightened class tensions: the rich blamed the poor for spreading the disease and the poor thought they were being poisoned. Animosity and anger were soon directed at the unpopular King. Following the funeral of General Lamarque – pandemic victim and defender of popular causes – a large anti-government demonstration occurred on the barricaded streets. This event was known subsequently as the June Rebellion[1] and was immortalized in Victor Hugo's novel *Les Misérables*.

The temporal proximity of a severe disease outbreak and unrest in this instance prompts a natural question: did the former cause the latter? More broadly, do epidemics cause social unrest? A natural presumption is that they do. Epidemics are traumatic events, so it does not seem implausible that they would generate considerable social discontent. Indeed, several authors have argued in favor of such a link in historical accounts. For example, Snowden (2019) argues that the cholera epidemic of 1832 exacerbated pre-existing tensions and was a principal cause of the June Rebellion. Jedwab et al. (2021) present a series of case studies including the plagues of antiquity, the black death, the influenza of 1918, and the spread of HIV/AIDs, arguing that the role of minorities as scapegoats is particularly important.[2]

Yet despite plentiful historic accounts of the social impact of epidemics, quantitative evidence on the link between epidemics and social unrest – large protests, riots, or other forms of civil disorder – is scarce. Existing evidence is predominantly from case studies, qualitative in nature, and focuses on chains of events in the long-run. To our knowledge, no systematic evidence existed before COVID-19 on the short-run social repercussion during or in the immediate aftermaths of epidemics. This paper fills that gap. We systematically investigate the historic view that epidemics leads to increased unrest. We ask whether this view is consistent with global evidence of epidemics in recent decades including COVID-19, especially over the short- and medium-run horizons – something particularly important for trying to understand any ongoing epidemic.

Conceptually, epidemics could affect the likelihood of social unrest in several ways. On the one hand, epidemics, like other threats to human health such as natural disasters, can subvert the social order. Mishandling of epidemics may reveal deep social problems such as insufficient social safety nets, incompetent government, or the public's lack of trust in institutions. Outbreaks of contagious diseases have historically caused "fear of the other" and backlash against certain groups (Deverell (2004); Hogarth (2017); Randall (2019); Jedwab et al. (2021)). Containment and mitigation efforts could be seen as excessive and unnecessarily costly ex post. Economic damage from epidemics, especially if affecting disproportionately the poor, could exacerbate inequality and sow the seeds of future social unrest. These are the scarring factors of epidemics and may give rise to increased social

[1] Sometimes known as the Paris Uprising of 1832.

[2] Other writers positing a similar link include Roberts (2020), Elledge (2020), and Wade (2020).

unrest. On the other hand, epidemics are humanitarian crises that disrupt lives and livelihoods. Such disruptions may impede social infrastructure – such as communication and transportation – and increases the costs of organizing and participating in major protests. Moreover, public opinion might favor cohesion and solidarity in times of duress. In some cases, incumbent regimes may also take advantage of an emergency to consolidate power and suppress dissent. A priori, the overall effect of these forces is ambiguous, depending on the offsetting effects of the scaring and mitigating factors, which likely vary across different time horizons.

Our analysis draws on a new cross-country dataset on social unrest – the Reported Social Unrest Index (RSUI). The RSUI is a monthly index constructed based on press coverage of social unrest for 130 countries from 1985 to the present (Barrett et al. (2020)). The authors identify social unrest events using spikes of the index and show that these identified events line up very closely with consensus narrative descriptions of unrest in a large number of case studies. In a subsequent paper, Barrett et al. (2021), we identify the starting dates for a subset of 203 events from 2011 to the present.

We use two complementary approaches to analyze these data. First, we use the daily data to estimate the short-run response of social unrest to the COVID-19 pandemic. We merge daily social unrest events with data on the timing and severity of COVID outbreak from Johns Hopkins University as well as data on containment policies from the Oxford COVID-19 Government Response Tracker (OxCGRT) (Hale et al. 2021). We interpret this data using a dynamic difference-in-differences (DID) model that exploits the staggered timing of COVID outbreaks at the country level. It also exploits the staggered timing of countries' adoption of COVID containment policies. It thus allows us to test the impediment of social interactions (as a result of containment policies) as a mitigating factor of social unrest.

We find that COVID outbreak is associated with a significantly lower likelihood of unrest: 55 percent lower than the pre-COVID average. Lockdowns – restrictive forms of containment policies – are also associated with a significantly lower likelihood of unrest. The magnitude of the effect persists beyond the initial COVID outbreak and extends to a week after a lockdown ends. These results are consistent with epidemics mitigating the likelihood of social unrest over the short horizon through the impediment of social interaction.

The causal interpretation of these results rests on the following grounds. First, the exact timing of COVID outbreak across countries was arguably unforeseen over a short time window, as is the timing of containment policy, which largely follows COVID outbreak. We find little evidence that the likelihood of unrest was significantly different during the month preceding COVID outbreak or a lockdown. Second, the dynamic DID model allows us to directly test reverse causation claims about social unrest affecting containment policies. For example, countries may impose restrictions in response to a surge in social unrest in order to suppress dissent or conversely loosen restrictions in order to calm the public. The former will imply a significantly higher likelihood of social unrest before a lockdown while the latter will imply a significantly higher likelihood of social unrest after the end of a lockdown. Our estimates are inconsistent with these predictions, suggesting little evidence of these sorts of reverse causality.

In our second approach, we use the monthly data for the whole sample period to estimate dynamic responses of social unrest following an epidemic via local projection in the style of Jordà (2005). We merge monthly unrest events with by a comprehensive dataset of disasters: the International Disaster Database (EM-DAT). These findings are consistent with the findings on the COVID-19 pandemic. Following an epidemic, the likelihood of social unrest in a given country declines. The monthly likelihood falls persistently by about 0.2 percentage points, or by around one fifth. To understand whether this result is specific to epidemics or it is relevant for disasters more broadly, we perform the same exercise for four types of natural disasters – droughts, earthquakes, floods, and storms. There, we find no statistically significant relationship between disasters and subsequent unrest. This rules out factors common across different disasters – interruption of communication and transport and the like – as drivers for our results. Instead, we posit that the unique feature of epidemics drives this result – that social gatherings are essential to unrest and less common during an epidemic. This might be voluntary. That is that, people simply don't want to meet in large groups during a pandemic for fear of catching the disease (in contrast, you can't "catch" an earthquake from other people). Or it might be due to public health responses which prevent social gatherings, such as lockdowns. Distinguishing between these two channels is particularly difficult in the panel, although the evidence on the COVID-19 pandemic specifically suggests that the public health response is at least an important part of the story. Finally, the comparison to other disasters also serves as a partial placebo test, showing that nothing mechanical in our empirical method or the measurement of unrest is driving these results.

One obvious identification challenge to the local projection estimate is omitted variables bias. It could be that some other third factor might drive both epidemics and social unrest. For instance, forces which increase interaction and mobility between people – either deliberate, like the loosening of travel restrictions, or incidental, like opening borders or long-distance transport infrastructure – could be popular but bad for public health. To address this concern, we also estimate the dynamic effects of epidemics on unrest using an instrumental variables (IV) approach. Specifically, we use epidemics in neighboring countries to instrument for epidemics in a country. The rationale for this instrument is that epidemics tend to spread across countries in regional waves and the spread is plausibly exogenous to domestic circumstances. This identification strategy is similar to Acemoglu et al. (2019) who use waves of regional democratization to instrument for local democratic changes. The instrumental variables estimate has the same sign as the ordinary least squares (OLS) one, with a much larger magnitude. This difference may reflect attenuation bias of the OLS estimate. It may also be that the IV estimate – by exploiting variations from domestic epidemics concurrent with international disease outbreaks – likely captures the effects of severer epidemics. Therefore the IV approach estimates not the average impact of epidemics in general, but of the most serious ones. In other words, this is a local average treatment effect (Imbens and Angrist (1994)). We do not take a strong stand on whether the OLS or IV estimates is better. Instead, we think of them as different points on a tradeoff. The former is more broadly relevant but likely less well-identified; the latter has fewer identification challenges but is less generalizable. That both show the likelihood of unrest decreases after epidemics, though, is reassuring.

In summary, the COVID-19 pandemic and other epidemics in recent decades offer robust evidence that the likelihood of social unrest *declines* during and in the immediate aftermath of an epidemic. At first glance, this is somewhat at odds with the view raised in historical accounts as we mentioned earlier. One possibility is that this difference reflects the different time horizons analyzed. Our approach is more advantageous in discerning effects over the short- and medium run whereas historical accounts may be better suited for drawing inferences for a chain of events in the long run. This comparison also points to another limitation of drawing generalized conclusions from narrative historical accounts – that the conclusion may be driven by the selection of specific episodes. We will return to this point at more length in our concluding remarks.

Our paper is related to a large literature on social and political instability. One strand of the literature examines the impact of social and political instability on growth, output, investment (Alesina and Perotti (1996); Alesina et al. (1996); Jong-A-Pin (2009); Aisen and Veiga (2013); Bernal-Verdugo et al. (2013)), and stock market performance (Barrett et al. (2021)). Miguel et al. (2004) find that economic growth is strongly negatively related to civil conflict: a negative growth shock of five percentage points increases the likelihood of conflict by one-half the following year. A separate strand of the literature examines the determinants of social unrest. Ponticelli and Voth (2020) find a positive correlation between fiscal austerity and social unrest in Europe in the period between World War I and the Global Financial Crisis. Fearon and Laitin (2003) and Collier et al. (2009) find that poor policies and institutions are important determinants of social unrest among low-income countries (LICs) or other emerging and developing economies (EMDEs).

Althought empirical evidence on the relationship between epidemics and social unrest are scant, there is some literature the connection between *violence* and epidemics. Censolo and Morelli (2020) argue that other episodes of cholera epidemics in the 19th century sowed seeds of conflict. And Cervellati et al. (2017) find that exposure to multi-host vector pathogen (e.g., Malaria, zika, and yellow fever) affects the likelihood of civil wars – a particularly extreme form of social instability. In another study, Cervellati et al. (2018) find that a higher exposure to malaria in African countries increases the incidence of civil violence, which include riots and protest as well as severe forms of conflicts such as battles, killings, and violence against civilians. Both these last two papers focus on severer forms of social instability in Africa and explore cross-section variation in the exposure to a specific type of contagious diseases. In comparison, our paper offers broad-based evidence on social unrest – a less severe form of social instability – and explores both cross-sectional and time-series variations from the short- to medium run horizons.

The rest of the paper is organized as follows. Section 2 describes our data sources. Section 3 studies social unrest during the COVID-19 pandemic. Section 4 investigates the relationship between epidemics and unrest more generally and Section 5 concludes.

2 Data

2.1 Sources

We use a newly constructed data on social unrest events based on the Reported Social Unrest Index (RSUI) of Barrett et al. (2020). The authors use articles from major international news sources to create country specific RSUI. For each country, the RSUI is constructed using the number of articles on social unrest as a fraction of total articles. The authors use text-based criteria to identify articles on social unrest. For example, relevant articles must include words such as "protest", "riot" or "revolution". They must also exclude certain terms to avoid counting reports about previous events or revolution-themed movies. The selected articles must be at least 100 words long and must mention the name of the country in question. At the country level, the RSUI exhibits very large spikes that are associated with major episodes of social unrest. The authors develop quantitative criteria to formalize these spikes and identify a list of social unrest events at the country-month frequency. They compare these events against consensus narratives for a number of case studies, showing that they align closely, and conclude that this method captures actual major social unrest events. The final dataset consists of 569 events in 130 countries from mid-1980s to the present.

We merge the social unrest data with EM-DAT – a comprehensive database of international epidemics and natural disasters, with information on the timing and location of more than 11,000 events since the 1980s. Although our focus is on epidemics, we also use data on four other types of disasters. Three types – floods, storms, and earthquakes – are, along with epidemics, the most common events in the sample (online appendix Table 6). the fourth type of disaster, droughts, is only the seventh most common type of event but have particularly extensive impacts, on average affecting ten times as many people as the next most far-reaching disaster, storms. The comparison to other disaster types is useful because, as we discuss earlier, epidemics and natural disasters are both humanitarian crises that may present similar challenges to social orders.

In our analysis in Section 3, we merge the daily social unrest data with data on the timing and severity of COVID outbreak from Johns Hopkins University as well as data on containment policies from the OxCGRT.

2.2 Stylized facts

Here we summarize some basic facts in our data. First, we consider country-level variation in social unrest and disasters (Table 1). Notably, whereas social unrest seems to be distributed similarly across geographic and income groups, epidemics are not. In particular, social unrest seems similarly distributed across regions and income groups, with around half of all countries experiencing 4 to 5 unrest events within the sample. The only exception is Advanced Economies, which seem to have notably fewer unrest events. Epidemics, however, seem to be concentrated in low-income countries, particularly in Africa, although countries in Asia-Pacific also have an increased likelihood of epidemics. In Africa, the median country experiences 17 epidemics in our sample – in Europe the equivalent number is just 1. This geographic spread is not entirely surprising given cross country differences in climate and public health systems. It suggests that country-specific factors are important to control for in subsequent analysis.

	Social Unrest Percentile			Epidemics Percentile		
	10	50	90	10	50	90
Region						
Africa	2.9	5.5	8.0	8.0	17.0	35.6
Asia-Pacific	1.0	4.0	6.0	1.0	6.0	17.0
Europe	1.0	4.0	6.2	0.0	1.0	2.2
Middle East and Central Asia	1.0	5.0	6.3	0.0	1.0	13.1
Western Hemisphere	1.7	5.0	7.3	1.7	6.0	9.3
Income group						
Advanced Economies	0.0	2.0	5.0	0.0	1.0	3.7
Middle-Income Countries	2.0	5.0	7.0	0.0	2.0	10.0
Low-Income Countries	2.0	5.0	7.3	2.7	11.5	34.0

Table 1: Number of events by country: Unrest and Epidemics, 1990-2019

The cross-sectional relationship between unrest and epidemics, shown in Figure 1a, seems to accord with the view from historical accounts – countries with more epidemics also have more unrest (the same is true for the other disasters we look at, see Appendix Figure 7). Interestingly this continues to be the case even after controlling for income group, suggesting that cross-sectional relationship is unlikely to be driven by lower income countries having more epidemics and more unrest. However, the aggregate time series, shown in Figure 1b, paints a very different picture. There, waves of epidemic and unrest seem to be negatively correlated. This could be for many reasons – causal and non-causal – but hints at the results to come in more sophisticated analysis later.

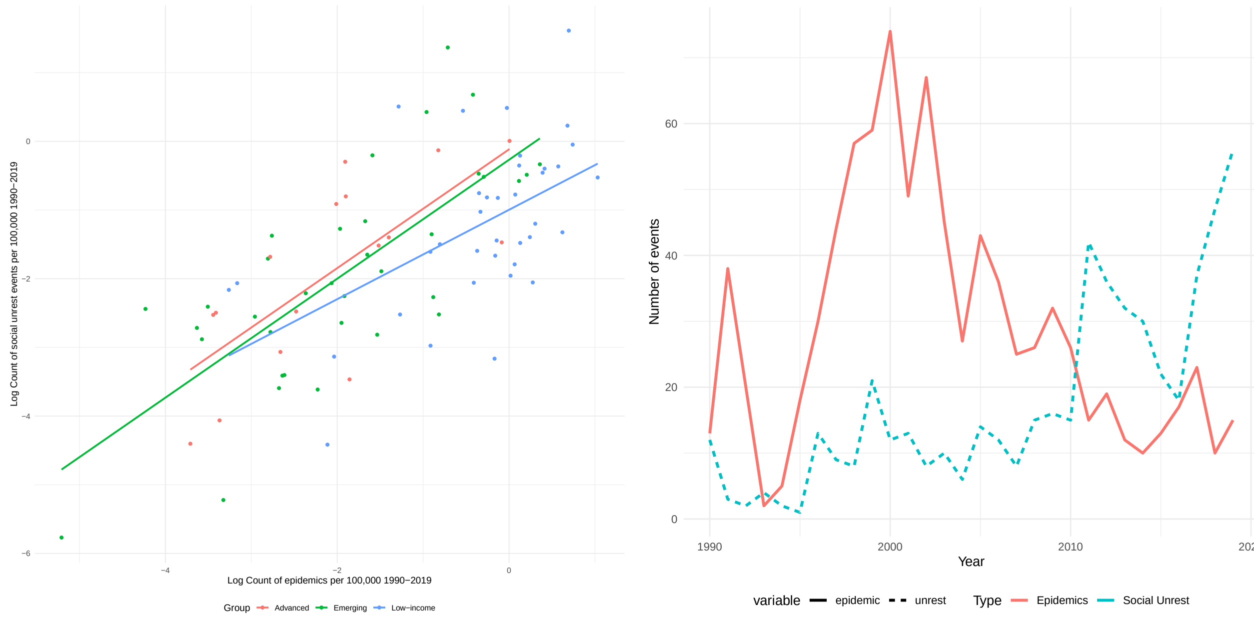

(a) Cross-section, per 100,000 population

(b) Time series, number of events

Figure 1: Correlations of social unrest and epidemic events, 1990-2019.

3 Social Unrest During the COVID-19 Pandemic

We start by looking at a particularly acute recent pandemic: the COVID-19 pandemic beginning in 2020. This is a useful starting point for several reasons. First, this is the biggest global pandemic in at least a century. If our results do not hold in this case, the relevance of our work is somewhat undermined (and conversely, if they do, it makes them more pertinent). Second, that the nature of this episode eases some identification problems which might otherwise be very challenging. One might be concerned about reverse causality – that because large gatherings of people are a frequent component of social unrest they may act as super-spreader events – or that if pandemics are to some extent anticipated, the timing of cause and effect is hard to pin down. As we discuss later, a case study of the COVID-19 pandemic cannot entirely eliminate these concerns, but the fine (daily) resolution of the data can at least mitigate them.

We present our analysis of this episode in two steps. First, a broad-brush narrative account of unrest during the pandemic which, despite its crudeness, conveys the basic point clearly. Second, a more careful and quantitative investigation, exploiting variation in the precise timing of the pandemic across different countries, provides more precise estimates. In both cases, we find that the likelihood of unrest *falls* concurrent with the onset of the pandemic, remaining low for at least a month. Longer-run conclusions are harder to deduce, simply because this is a recent event.

3.1 Narrative account

Social unrest pre-pandemic. Pre-pandemic, unrest was higher than usual. Between January 2019 and January 2020, there were 55 unrest events in forty countries, up from 45 events the year before. Major protests in late 2019 to early 2020 most notably included those in the Arabic-speaking world and South America. None of these events appear to be directly linked to major natural disasters or epidemics. Instead, most of the unrest events were motivated by political factors, such as pro-democracy demonstrations in Sudan and Algeria, and anti-government protests in Chile, Columbia, Ecuador, and others. This recent wave of social unrest events was the continuation of a longer trend since 2016 (Figure 2), which itself reversed a gradual decline in unrest following a peak after the Arab Spring of 2011 (see Figure 1b).

Social unrest during the COVID-19 pandemic. After the COVID-19 outbreak, the number of major unrest events worldwide fell sharply and in March 2020 reached its lowest level in almost five years (see Figure 2). The decline in social unrest matches remarkably closely the global spread of the pandemic in Spring 2020. This is illustrated in Figure 2 which includes the global averages of two series of community mobility from Google Community Mobility Reports. These series use cell phone location to measure activities in specific categories of locations, of which we include two: activities in retail and recreation spaces and transit stations, although others look very similar. The timings of the decline in protest and the abrupt cessation of social activities are almost identical. We interpret this time series correlation as suggestive evidence that social distancing – both mandatory (e.g., shelter-in-place orders, school closures, etc.) and voluntary – interrupted not just ordinary commercial and social activities but also an extraordinary one, social unrest. In other words, the mitigating effects of the latest epidemic have likely outweighed any factors for increased unrest.

Of course, social unrest did not disappeared completely, most notably in the United States and Lebanon, both of which saw mass protests in the first half of 2020. Yet even in these cases, the largest protests were related to issues that long preceded the COVID-19 epidemic: racial injustice in the United States and governance in Lebanon.

The protests in the United States are a good cross-check of the media-based approach to measuring unrest. An important challenge to our method is that the media coverage of other high-profile issues, such as an ongoing pandemic, may "crowd out" the coverage of unrest. In Figure 3 we replicate the search criteria of Barrett et al. (2020) at a daily frequency for the United States. It shows that press articles related to unrest increase sharply at almost exactly the same time as major street protests broke out across the United States despite a severe and continuing pandemic. This gives us some confidence that media coverage remains a good indicator of unrest even during a pandemic.

3.2 Quantitative Analysis

In this subsection, we use a dynamic DID model to quantify the effect of the COVID-19 pandemic on social unrest. We use two key variables to capture the staggered timing of the pandemic across

Figure 2: Global unrest and mobility since 2016. Mobility data comes from Google Community Mobility Reports, sample covers 130 countries in the RSUI.

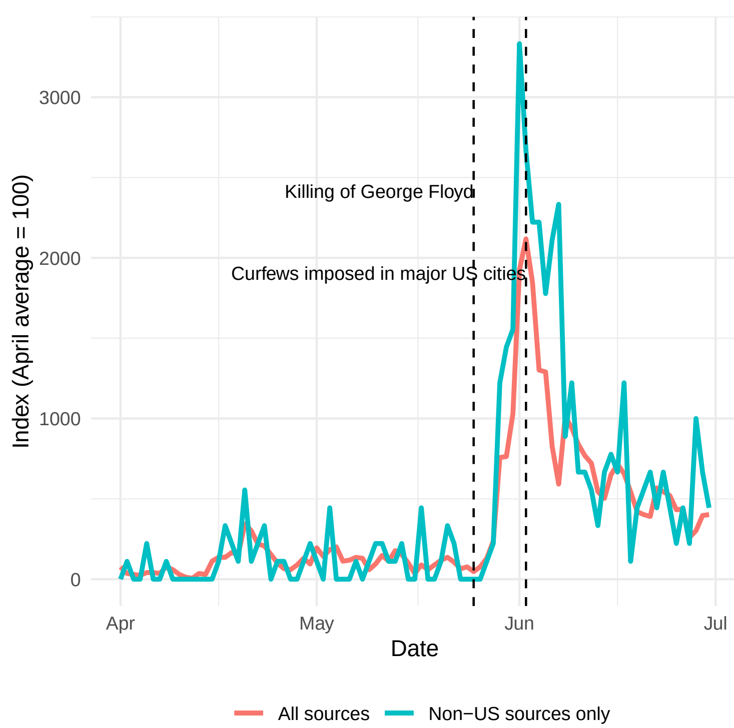

Figure 3: Daily media mentions of unrest in the United States, early 2020, collected from 18 English-language news sources in the UK, USA, and Canada (see Barrett et al. (2020) for details).

countries. The first is disease outbreak measured by the timing of a country's first confirmed COVID case. This variable aims to capture behavioral changes in response to an infectious disease as a result of, for example, fear and voluntary social distancing (Chen et al. (2020)). The second is a country's adoption of containment policies. This variable aims to capture behavioral changes in response to mandatory restrictions on mobility and social gathering. Specifically, we estimate the following regression:

$$y_{it} = \alpha_i + \eta_t + \sum_j \beta_j COVID_{it-j} + \sum_j \gamma_j Lockdown_{it-j} + \varepsilon_{it},$$

where y_{it} is an indicator variable for a social unrest event in country i on date t. Because the average daily likelihood of social unrest is small (about 3 basis points in our sample), we multiple the indicator variable y_{it} by 100 so it can be interpreted as a likelihood in percentage point. α_i is country fixed effects to account for any time-invariant country characteristics that may affect social unrest, such as institution, culture, and political systems. η_t is date fixed effects to account for any common shocks that may affect social unrest, such as a global wave of unrest as a result of spillover or simply seasonality.

$COVID_{it-j}$ $(j \in J)$ is a set of indicator variables that take a value of 1 if country i's COVID outbreak (i.e., the date of first confirmed COVID case) occurred on date $t-j$. The set of indicators $j \in J$ covers the following time window: the day of COVID outbreak, between 1 and 7 days after the outbreak, between 8 and 30 days after the outbreak, between 31 and 180 days after the outbreak, between 1 and 7 days before the outbreak, and between 8 and 30 days before the outbreak. These indicators captures time windows around the outbreak with varying intervals: shorter intervals near the outbreak and longer intervals further away. This approach is advantageous to including daily indicators because it significantly reduces the number of indicators in the regression and increases the power of the estimation.[3] It achieves this by preserving information at daily frequency and accounting for potential nonlinear effects over time. The two indicators covering the pre-event period are included to test a pre-trend, that is, whether the likelihood of unrest started changing before the outbreak. Similarly, $Lockdown_{it-j}$ $(j \in J)$ is a set of indicator variables that take a value of 1 if country i implemented a lockdown policy on date t. This set of lockdown indicators cover the same time window as the set of COVID indicators. We estimate the regression over the period of January 1, 2019 to March 31, 2021. The coefficients β_j and γ_j thus captures the likelihood (in percentage points) of social unrest during the time interval j relative to the pre-COVID period from January 1, 2019 to 31 days before the event. We cluster standard errors at the country level to allow for potential serial correlation within a country.

Our data source for COVID outbreak is Johns Hopkins University's COVID-19 Dataset, which reports the daily number of COVID cases by country. Our data source for lockdown policies is OxCGRT, which contains information on various categories of policy responses collected from

[3]The estimates of daily indicators would be extremely noise due to the very low likelihood of unrest on a daily basis. We return to this point in Section 4

official sources and media reports. We use information on the OxCGRT category on containment and closure policies, which includes eight measures: the closing of school, workplaces, and public transport, the cancellation of public events, restrictions on public gatherings, domestic travel, and international travel, and stay at home requirements. Each measure is assigned a scale between 0 (for no policy in place) to 4 (for the most restrictive policies). We construct a stringency index as a weighted average of the 8 measures, each weighted by the maximum value of the measure:

$$SI_{it} = \frac{1}{8} \sum_{p=1}^{8} \frac{S_{it}^{P}}{\max S^{P}},$$

where S_{it}^{P} is the score for measure p in country i on date t. This scaling accounts for the fact that various measures may have a different gradient of restrictiveness. For example, the closing of public transport has 3 levels of restrictiveness (no measure, recommend closing, and required closing) whereas restrictions on gatherings has 5 levels (from no measure to restrictions on various gathering sizes).[4] The stringency index SI_{it} ranges from 0 (for no policy) to 1 (for the most stringent policy). It is similar to the OxCGRT stringency index with two exceptions: we only include the 8 policies in the containment and closure category whereas OxCGRT also includes a measure on public information campaigns; we only include national measures whereas OxCGRT also includes subnational measures. We assign the indicator $Lockdown_{it}$ to 1 if SI_{it} is above a threshold. We use a threshold of 0.8 in our baseline and alternative thresholds of 0.65 and 0.95 in robustness tests.

Given that we aim to identify the effects of the pandemic on the likelihood of unrest, a threat to our identification strategy is that the timing of COVID outbreak may coincide with other unobserved factors associated with social unrest. This seems unlikely because the exact timing of COVID outbreak in different countries was unforeseen. Although the timing of a country's first confirmed case may be related to its global connection, time-invariant factors of this sort are absorbed by country fixed effects. For lockdown policies, our data suggest that the timing of its adoption follows closely with the disease outbreak. Nevertheless, it is possible that countries may impose restrictions in response to a surge in social unrest in order to suppress dissent or conversely loosen restrictions in order to calm the public. In the case of the former, one expects a significantly higher likelihood of social unrest before a lockdown. In the case of the latter, one expects a significantly higher likelihood of social unrest after a lockdown ends. As we will show, our estimates are inconsistent with these predictions, providing no support of these reversal causation arguments.

The remaining threat is from fast-moving factors that may be correlated to the likelihood of social unrest. We control for two likely candidates: the severity of the pandemic and community mobility. People may be less inclined to participate in unrest if the disease transmission or the disruption in mobility is severer (for reasons related or unrelated to our outbreak and lockdown measures). We measure the severity of the pandemic in a country on a given date by the log

[4] Policy measures with a finer gradient of strictness have a larger maximum value than measures with a coarser gradient. In the calculation of the stringency index, each score is weighted by the maximum value of the measure so that each measure contributes equally to the index.

of total confirm COVID cases. We measure community mobility using total visits to categorized places from Google's Community Mobility Report. This measure captures changes in total visits relative to the pre-COVID period.

Table 2 reports the results. Column 1 and 2 report results with the COVID indicators and lockdown indicators estimated separately. Column 3 reports results with them estimated jointly. Column 4 adds the outbreak severity control and column 5 adds the community mobility control.

We have two main findings. First, we consistently find that COVID outbreak is associated with a significantly lower likelihood of unrest. The effect is also economically meaningful. For example, column 3 suggests that the likelihood of unrest on the day of COVID outbreak is 55 percent lower than the pre-COVID average, 1.6 basis points compared to 3.6 basis points. Results from other columns are quantitatively similar. We also find that the effect remains negative with a week of the outbreak, but it is only statistically significant when the lockdown policy is not controlled for. We also find a significantly negative likelihood of unrest in the week before the COVID outbreak. This is not surprising for several reasons. A small lag may exist between the actual COVID outbreak and the confirmation of the first COVID case in the official statistics. People's behavior may start to change in response to COVID-like symptoms even before the case was officially confirmed—which took days in many countries during the early phase of the pandemic. Our 1-week pre-COVID indicator thus allows for some ambiguity in the exact timing of COVID outbreak. Nevertheless, the same ambiguity should not indicate a significant result during a longer pre-COVID period, the lack of which is exactly what we find. The estimate for the 8 to 30 days before the COVID outbreak indicator is not only statistically insignificant but also very small—in the order of one hundredth of the magnitude of the estimate for the date of the COVID outbreak. Overall, these results point to a negative effect of COVID outbreak on the likelihood of social unrest with no anticipation effect.

Second, lockdown is also associated with a significantly lower likelihood of unrest. The effect is again quantitatively similar across the specifications. We also find a significantly negative lagged effect within a week after a lockdown ends. Importantly, we do not find evidence of significant changes in the likelihood of unrest before a lockdown—this holds for the 1-7 days before and 8-30 days before lockdown indicators across the specifications. The lack of pre-lockdown result is very reassuring because, unlike the timing of COVID outbreak, the timing of lockdown is subject to less ambiguity. Lockdown policies are usually announced publicly and implemented shortly after the announcement.

Even though the timing of lockdown policies can be clearly identified from public information, the definition of a lockdown may be somewhat subjective. As the OxCGRT data show, many governments implemented a package of containment and closure measures, with different degrees of stringency. Our stringency index is meant to capture the variation in stringency. To investigate the robustness of our results with respect to the stringency of policies, we alternatively define the lockdown indicator to take a value of 1 if the stringency index is above 0.65 or 0.95, in additional to the baseline threshold of 0.8. Because governments tend to tighten containment and closure measures when the disease spread is in the upswing stage and loosen them in the downswing stage, a lockdown dummy defined with a lower threshold tend to capture a longer time window than one defined with a higher threshold. For example, when a government loosens a very restrictive policy in phases, the first (in terms of calendar day) lockdown indicator to turn from 1 to 0 is based on the 0.95 threshold; the second, the 0.8 threshold; the third, the 0.65 threshold. As a result, we expect the effect of the lockdown indicator to be stronger based on the higher threshold and weaker based on the lower threshold. This is what we find in Appendix Tables 7 and 8. Consistent with the baseline, the coefficients of lockdown indicators based on alternative thresholds are significantly negative. The absolute value of the coefficients tend to be smaller with the 0.65 threshold and larger with the 0.95 threshold. Furthermore, the indicators for one month after and before lockdown are not significant with the 0.65 threshold but they are negatively significant with the 0.95 threshold. This is to be expected again as a result of a shorter time window of lockdown under the 0.95 threshold: Less stringent measures may exist some days after the lockdown indicator turns from 1 to zero, whose effect is captured by the 1-month post lockdown indicator under the 0.95 threshold but by the lockdown indicator under the 0.65 threshold.

Overall, the results robustly point to a lower likelihood of unrest after the COVID outbreak relative to the pre-COVID period. This is consistent with the notion that mitigating factors from the disease outbreak outweigh scarring factor in the immediate aftermath of COVID-19. Although disentangling potential factors is out of the scope of this paper, our results offer some suggestive evidence. The fact that both the COVID and lockdown indicators have a negative sign suggests that behavioral changes for both mandatory and voluntary reasons are at force.

4 Estimating the Dynamic Effect of Epidemics on Unrest

We now turn to the estimation of the effects of epidemics on unrest from the full sample. Our principal analytical device is local projections in the style of Jordà (2005), which estimates the dynamic impulse response of unrest to an epidemic.

Dynamics are interesting in this case in their own right. Popular sentiment may respond to pandemics differently at different horizons, for example if fatigue at public health policies changes over time. Moreover, the outcome variable is a rare event – social unrest occurs in around 1 percent of observations. This means that realized unrest is typically a rather noisy measure of average unrest. So a framework which allows for changes in the average likelihood of unrest over multiple periods (i.e., the dynamic response) has the benefit of greater statistical power.

One further benefit of the local projection approach is that it allows us to easily address the key challenge to identification – omitted variables bias. One might be concerned that some other variable, such as government competence, drives both unrest and epidemics. The local projection setting allows us to address this by use of an instrument. Specifically, we use regional pandemic waves to compute a local average treatment effect.

	(1)	(2)	(3)	(4)	(5)
			Likelihood of unrest (%)		
COVID (first case)	-0.0151*		-0.0197*	-0.0196*	-0.0181*
	(0.009)		(0.010)	(0.010)	(0.010)
COVID (1-7 days after first case)	-0.0112*		-0.0150	-0.0150	-0.0138
	(0.007)		(0.009)	(0.009)	(0.009)
COVID (8-30 days after first case)	-0.0129		-0.0143	-0.0144	-0.0140
	(0.008)		(0.010)	(0.010)	(0.010)
COVID (31-180 days after first case)	-0.0032		0.0007	0.0006	0.0002
	(0.012)		(0.013)	(0.013)	(0.013)
COVID (1-7 days before first case)	-0.0272*		-0.0278**	-0.0274*	-0.0266*
	(0.014)		(0.014)	(0.016)	(0.014)
COVID (8-30 days before first case)	-0.0002		-0.0006	-0.0004	-0.0000
	(0.017)		(0.017)	(0.017)	(0.017)
Lockdown (during)		-0.0510***	-0.0513***	-0.0515***	-0.0539***
		(0.017)	(0.018)	(0.020)	(0.018)
Lockdown (1-7 days after lockdown)		-0.0448**	-0.0450**	-0.0452**	-0.0463**
		(0.018)	(0.018)	(0.020)	(0.019)
Lockdown (8-30 days after lockdown)		0.0311	0.0310	0.0307	0.0301
		(0.048)	(0.048)	(0.048)	(0.048)
Lockdown (31-180 days after lockdown)		-0.0146	-0.0147	-0.0150	-0.0149
		(0.018)	(0.018)	(0.020)	(0.018)
Lockdown (1-7 days before lockdown)		0.1006	0.1028	0.1027	0.1023
		(0.135)	(0.136)	(0.136)	(0.136)
Lockdown (8-30 days before lockdown)		-0.0257	-0.0228	-0.0229	-0.0233
		(0.018)	(0.018)	(0.018)	(0.018)
Log of confirmed COVID cases				0.0001	
				(0.001)	
Mobility					-0.0002
					(0.000)
Constant	0.0309***	0.0354***	0.0361***	0.0360***	0.0352***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)
Observations	155,169	155,169	155,169	155,169	155,169
R-squared	0.007	0.007	0.007	0.007	0.007
Country FE	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes

Table 2: COVID-19 and social unrest

Notes: This table shows results of a dynamic DID model for the sample period of January 1, 2019 to March 31, 2021. The dependent variable is the daily likelihood (in percentage) of social unrest in a country. COVID (first case) is a indicator that takes a value of 1 on the day of a country's first confirmed COVID case. COVID (1-7 days after first case) is an indicator that takes a value of 1 between 1 and 7 days after the first confirmed case. Lockdown (during) is an indicator that takes a value of 1 during a lockdown (i.e., when the stringent index is above 0.8). Other indicators are defined similarly. Standard errors clustered at the country level are shown in parentheses. *p<0.1; **p<0.05; ***p<0.01.

4.1 Local projection: Ordinary Least Squares

We consider two specifications for our local projection analysis. For each horizon $h = -H, \dots, 0, \dots, H$, we estimate separately:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i + \beta^h x_{i,t} + \sum_{k=1}^{K} \gamma_k^h y_{i,t-k} + \sum_{k=1}^{M} \eta_k^h x_{i,t-k} + \sum_{k=0}^{N} \lambda_k^h \tilde{y}_{i,t-k} + \nu^{h\prime} W_{i,t} + \epsilon_{i,t} \tag{1}$$

$$\frac{1}{1+h}(Y_{i,t+h} - Y_{i,t-1}) = a_i + b^h x_{i,t} + \sum_{k=1}^{K} d_k^h y_{i,t-k} + \sum_{k=1}^{M} f_k^h x_{i,t-k} + \sum_{k=0}^{N} g_k^h \tilde{y}_{i,t-k} + p^{h\prime} W_{i,t} + e_{i,t} \tag{2}$$

where $y_{i,t}$ is an indicator for a social unrest event in country i in period t and $Y_{i,t} = \sum_{s=0}^{t} y_{i,s}$ its cumulative sum over the period 0 to t. $x_{i,t}$ is an indicator for an epidemic, $\tilde{y}_{i,t}$ an indicator for unrest in any neighboring country, and $W_{i,t}$ a vector of controls. All data are monthly, so the period length is also a month.

The interpretation of the first equation is relatively straightforward. The dependent variable, $y_{i,t+h} - y_{i,t-1}$ is simple the change in the likelihood of unrest in period $t + h$ relative to period $t - 1$. So the coefficient β^h is the difference in the conditional average likelihood of unrest between h months after an epidemic versus those h months after a non-epidemic period, conditional on the other explanatory variables.

The interpretation of the second specification is a little less obvious from the algebra, but ultimately rather intuitive. Substituting in the definition of $Y_{i,t}$, we can see that the dependent variable can be written as $\bar{y}_{i,t}^h = \frac{1}{1+h} \sum_{j=0}^{h} y_{i,t+j}$. In other words, this is the *average* likelihood of social unrest in period t to $t + h$. So the coefficient b^h has the interpretation of the change in the average likelihood of unrest in the h months following a pandemic. This is a particularly useful quantity to compute when the monthly sequence is noisy.

Given that the dependent variable of equation (2) is just an average, why do we write it as the (scaled) difference between cumulative sums? It is for two reasons. First, to emphasize that there is no special magic when estimating a local projection for the average over h periods. Equation (2) is structurally identical to equation (1). It is a regular local projection, simply replacing the variable $y_{i,t}$ with its scaled cumulative sum. The relationship between equations (1) and (2) is identical to that between estimating local projections for inflation and the price level. Second, this formulation generalizes neatly to the case where $h < 0$. There, the resulting difference is simply the average likelihood in the h periods prior to an epidemic. By writing this as a single formula we aim to make clear that the before- and after-epidemic average outcomes are measured the same way.[5]

We estimate equations (1) and (2) by OLS. This means that equation (1) is a linear probability model rather than a logit or probit model. A linear probability model has the advantage of being easy to estimate with many controls including a battery of fixed effects. Moreover, because the

[5] There is a special case, where $h = -1$. There, we adopt the convention that $b^h = 0$, consistent with the fact that $\beta^h = 0$.

dependent variable in specification (2) is a fraction rather than a 0/1 indicator, logit or probit models would not work. Thus, OLS is a flexible and simple way to estimate both the contemporaneous and average dynamic effects.

Figure 4 presents the estimates for the simplest case, without controlling for past epidemics or unrest. As is evident from Figure 4a, the frequency of unrest seems to be generally lower following an epidemic, but the estimates are so noisy that it is rather hard to draw any firm conclusions. If the coefficients at each horizon were all estimated in a single framework one could simply perform a joint hypothesis test of multiple coefficients being negative. However, such a joint hypothesis is impossible to test using local projection, simply because each horizon is estimated in a separate regression. The average effect coefficient, b^h, provides a solution to this problem. Because the dependent variable in equation (2) is just the $h-$period average of that in equation (1), a similar relationship holds between the estimated coefficients: b^h is just the average of $\beta_0, \ldots, \beta_h$.[6] Indeed, even a casual inspection of the two panels of Figure 4 confirms what one would expect: that the average effect is just a smoothed version of the contemporaneous one.

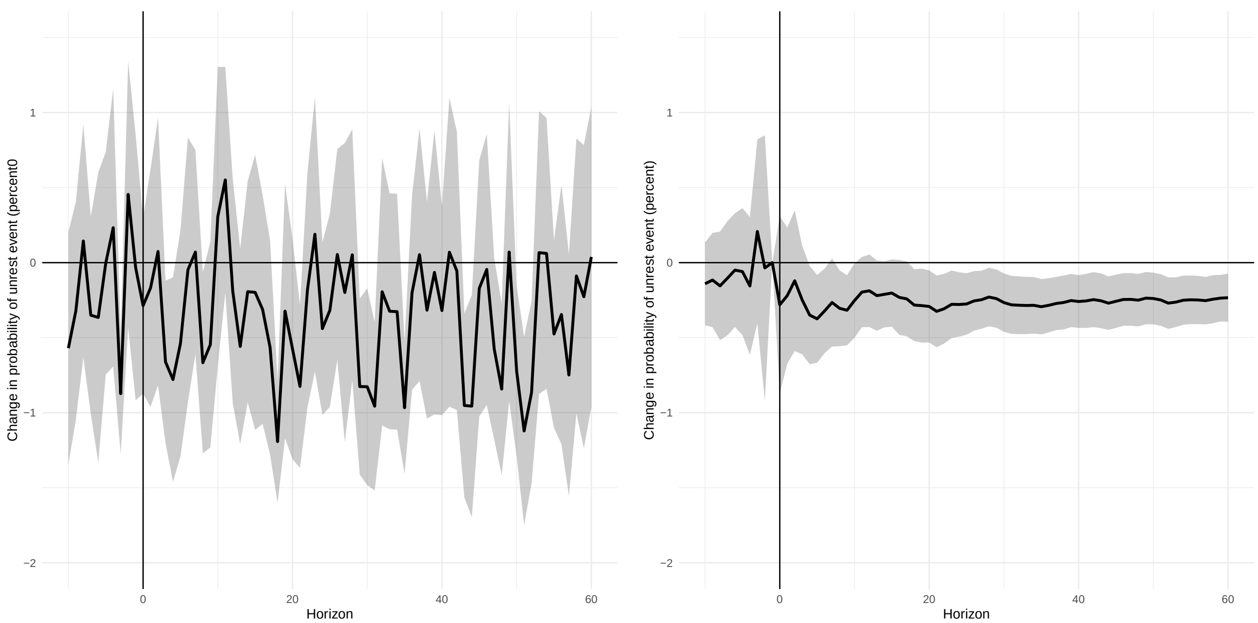

(a) Contemporaneous effect, β^h

(b) Average cumulative effect, b^h

Figure 4: Local projection impulse responses, effect of an epidemic event on the probability of social unrest, $K = N = M = 0$. Shaded regions show 95 percent confidence intervals using standard errors which are double-clustered by both country and time period.

This second panel allows us to look through the noise of the first, showing that there is a clear and persistent negative response of unrest to the incidence of an epidemic. Contemporaneous with the epidemic (i.e., $h = 0$) unrest declines, although it is not statistically significant. In the periods following the shock, the magnitude of the response is similar. It is also highly persistent and as

[6] It is straightforward to show that the linear projection of the sum of two variables, say $Y_1 + Y_2$, onto X, is just the sum of the linear projections of Y_1 onto X and Y_2 onto X.

the horizon extends the average effect is more precisely identified. After 18 months the effect is consistently statistically different from zero. The magnitude of this decline is in the order of 0.2 percentage points – a not insubstantial change given that the average unconditional probability of unrest is around 1 percent per month. To state more informally, following an epidemic the average likelihood of unrest declines by around 0.2 percentage points per month.

Being able to compute the average response is particularly useful in this case because of the noise in the monthly response. But why does this noise arise? Our interpretation is that it occurs because unrest is a relatively rare event and the average effect is small. A 0.2 percentage increase in the likelihood of unrest at a fixed horizon after an epidemic, say six months, requires a large number of observations to detect, especially when most of the time one will not observe unrest (as it is rare). This is compounded by the fact that epidemics are also rare, resulting in relatively few post-pandemic periods that can be used in the estimation. But by taking the average rate of unrest over the first six months after an unrest event, one smooths through the inevitable noise is trying to measure a rare event. This approach reduces the noise of the dependent variable and increases the precision of the estimates. The cost of this approach is a less clear picture of the month-to-month precise dynamics of unrest. But these were not very well-estimated to start with so it seems like a price worth paying. In other words, we choose to be more precise about something less detailed (the average response over multiple periods) rather than less precise about something more detailed (the exact response in each period).

There are two main challenges to any causal interpretation of these results. The first is that there may be omitted variables bias. Common factors, such as government (in)competence, could lead to both epidemics and unrest. The second is that there might be some sort of selection of the incidents with epidemics based on pre-existing characteristics.[7] One rebuttal to these points is the sign of the estimates – one would presumably expect such factors to cause a positive association between epidemics and unrest, something at odds with our findings. We postpone a more detailed discussion of general omitted variables bias to the next section (where we address this using an instrument). But we can say a little more about selection specifically even within the context of OLS.

One way to assess the extent of any selection bias is by studying the pre-shock period. If there were any obvious pre-existing trend, it would show up there. A simple test suggests that the average likelihood of unrest is not different in the period prior to an epidemic. An more formal way to address this issue is to include controls for lagged unrest, which soak up any pre-existing trends, i.e. $K > 0$. Thus, in Table 3 we also present results at select horizons for a variety of alternate specifications, including with lagged domestic and neighboring unrest. Our choice of additional controls is informed by analysis in Barrett et al. (2020) and Redl and Hlatshwayo (2021) who show that domestic and neighboring unrest are the strongest predictors of unrest, and that social and

[7] Strictly speaking, this is just a type of omitted variables bias. Selection just means that some omitted variables explain the selection of particular periods into the epidemic/non-epidemic categories. However, for purposes of exposition we distinguish it from omitted variables bias more generally.

economic factors are only very weakly correlated with unrest.[8] Including past epidemic indicators also adjusts for the fact that epidemics may come in tightly packed waves. In general, the results are consistent with the findings in the simplest specification (column (1) is the same as Figure 4b): the average effect across columns (1)-(5) is in the order of a 0.2 to 0.3 percent reduction in the probability of unrest; this effect begins to be significant after 6 months, and is uniformly so after 24. In other words, any concerns about selection appear to have a relatively minor quantitative effect on the results.

Table 3 also allows us to address a further issue: the role of common shocks which might affect all countries simultaneously. This is potentially a concern if there are other factors which drive both epidemics and unrest on average across all countries. In other words, if there are omitted variables in the time dimension. One could normally address this through the inclusion of time fixed effects. However, much of the variation in both epidemics and unrest is in waves, and for reasons plausibly unrelated to third factors. It is in the nature of epidemics to be correlated across countries at specific times. Moreover, anecdotal evidence suggests that unrest is also regionally correlated due to common cultural or linguistic differences – unrest in 2019 in Spanish-speaking Latin America, and the Arab Spring of 2011 both spread long distances through countries with a common tongue, but barely at all in other neighbors. Including time fixed effects would absorb much of the variation that we want to exploit. Nevertheless, to check the importance of this concern, we also estimate in equation (6) a specification using time and country fixed effects. There, the sign of the estimates remains negative, and the magnitude is not enormously different from the estimates with only country fixed effects. But the lack of residual variation after including time fixed effects mean that the statistical significance is harder to establish.

A natural question at this stage is whether the decline in unrest is specific to epidemics or if it is a property of disasters more generally. To investigate this question, we report the results of a similar exercise for four other types of disaster recorded in the EM-DAT dataset, in Figure 5 for the simplest case and in Appendix Tables 9-12. These show that epidemics do seem to be special. Overall, the response of unrest to these other disasters are smaller and rarely statistically significant. Significant effects only show up in the more complex specifications. This comparison hints at what might be driving our results on epidemics. Many factors which could affect unrest after epidemics are likely the same, or at least similar, across different types of disaster. Examples include public dissatisfaction at government performance, a "rally-round-the-flag" effect where a sense of common struggle deters dissent, and damage to transport and communication infrastructure which might impede the organization of protests. However, epidemics are distinct in one obvious way – the presence of a dangerous and communicable disease. Therefore a driving factor in the reduction in unrest following an epidemic might be a public distaste for gathering in large groups. Of course, public health restrictions may also play a role.

[8]This is consistent with theoretical work, such as Barbera and Jackson (2019) who model unrest as a co-ordination game. Such frameworks imply a minor role for "fundamentals" and a larger one for coordinating signals in explaining unrest. Past unrest is presumably a good coordinating signal as it informs individuals that others are willing to protest.

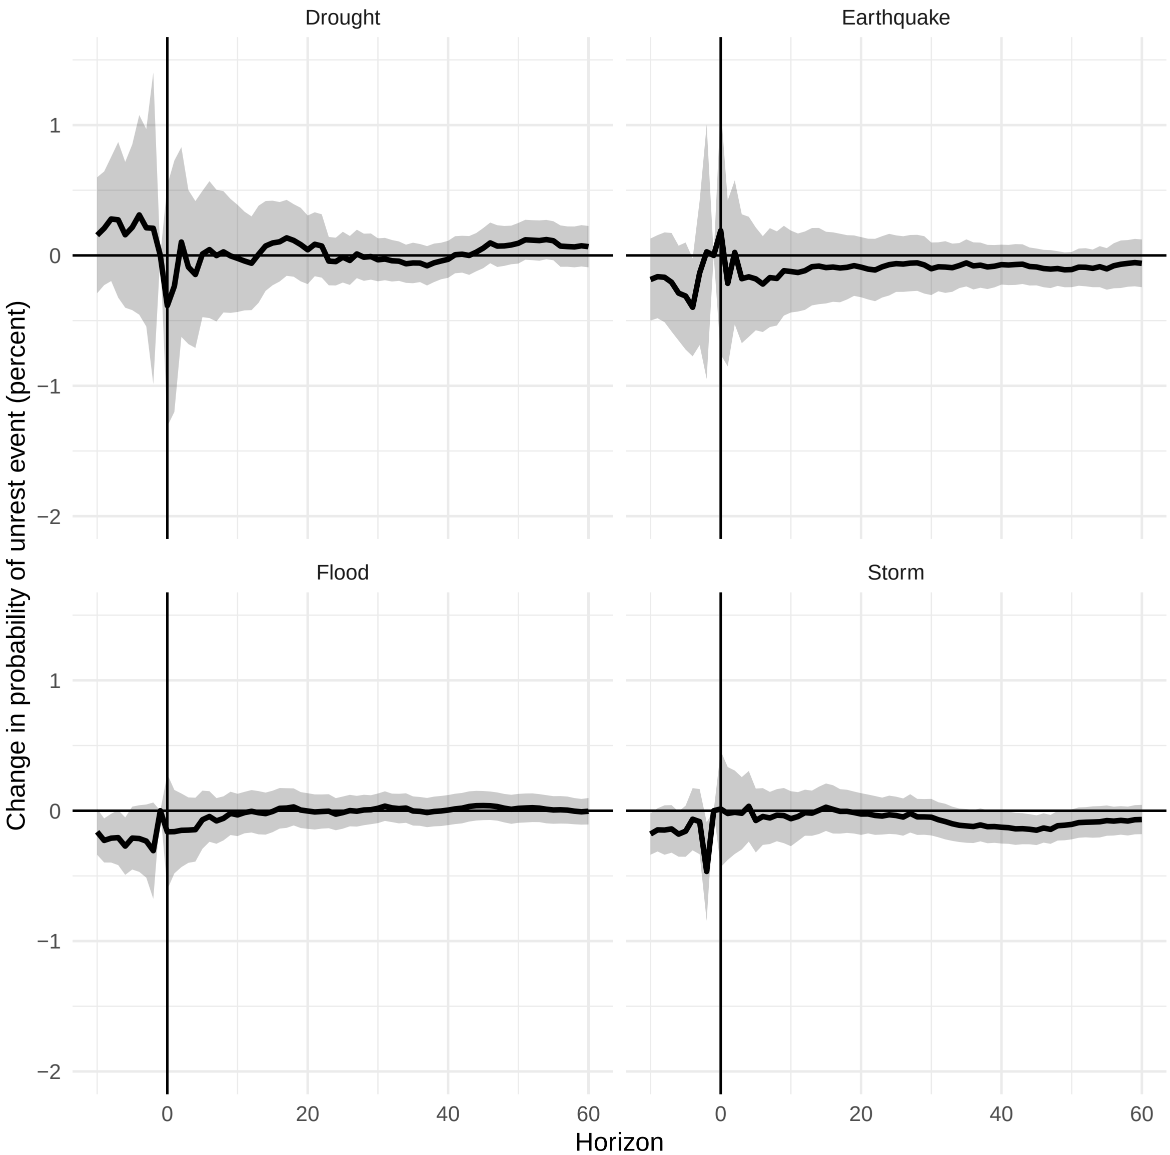

Figure 5: Local projection impulse responses, effect of various disasters on the probability of social unrest, $K = N = M = 0$. Shaded regions show 95 percent confidence intervals using standard errors which are double-clustered by both country and time period.

	(1)	(2)	(3)	(4)	(5)	(6)
$h=0$	−0.281	−0.257	−0.257	−0.312	−0.103	0.013
	(0.302)	(0.311)	(0.31)	(0.316)	(0.474)	(0.487)
$h=1$	−0.22	−0.189	−0.177	−0.236	−0.224	0.032
	(0.232)	(0.225)	(0.222)	(0.226)	(0.352)	(0.351)
$h=3$	−0.249*	−0.212	−0.204	−0.269*	−0.353	−0.167
	(0.185)	(0.174)	(0.174)	(0.166)	(0.285)	(0.287)
$h=6$	−0.322**	−0.295**	−0.29**	−0.326***	−0.349**	−0.206
	(0.144)	(0.13)	(0.129)	(0.126)	(0.21)	(0.207)
$h=12$	−0.188*	−0.174*	−0.171*	−0.199**	−0.215	−0.1
	(0.123)	(0.112)	(0.11)	(0.115)	(0.213)	(0.214)
$h=24$	−0.279***	−0.261***	−0.26***	−0.308***	−0.339**	−0.229*
	(0.109)	(0.099)	(0.098)	(0.098)	(0.174)	(0.169)
$h=48$	−0.25***	−0.232***	−0.232***	−0.264***	−0.289**	−0.162*
	(0.089)	(0.078)	(0.079)	(0.08)	(0.132)	(0.117)
Months lagged Epidemic	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Table 3: Impact of Epidemics on social unrest, Local Projection Average Effect

4.2 Local projection: instrumental variables

When it comes to a causal interpretation of our results, the issue of omitted variables bias remains only partly addressed. The problem, in general, is that epidemics and unrest may be negatively correlated for non-causal reasons. For example, changes in government policies which reduce restrictions on travel or social interactions may be popular but could also increase the likelihood of an epidemic.

One way to address this is through the use of an instrument[9] – a variable correlated with the occurrence of epidemics but without a direct impact on social unrest. Our main idea for such a variable is to exploit the variation resulting from one of the natural features of epidemics: that they spread. Specifically, we use epidemic status in neighboring countries as a proxy for regional waves of epidemics. So our instrument is given by:

$$z_{i,t} = \begin{cases} 1 & \text{A neighboring country has an epidemic} \\ 0 & \text{Otherwise} \end{cases} \tag{3}$$

Given that average effects seem better-identified than contemporaneous ones, we estimate the following two stage least squares model:

$$\frac{1}{1+h}(Y_{i,t+h} - Y_{i,t-1}) = a_i + b^h \hat{x}_{i,t} + \sum_{k=1}^{K} d_i^h y_{i,t-k} + \sum_{k=1}^{M} f_i^h x_{i,t-k} + \sum_{k=0}^{N} g_i^h \tilde{y}_{i,t-k} + p^{h\prime} W_{i,t} + e_{i,t} \tag{4}$$

where $\hat{x}_{i,t}$ are the fitted values from :

$$x_{i,t} = \chi_i + \sum_{k=0}^{L} \theta_k z_{i,t-l} + \sum_{k=1}^{K} \eta_k y_{i,t-k} + \sum_{k=1}^{M} \phi_k x_{i,t-k} + \sum_{k=0}^{N} \psi_k \tilde{y}_{i,t-k} + \pi' W_{i,t} + u_{i,t} \tag{5}$$

Specifically, we instrument epidemics in country i with current *and* lagged epidemics in neighboring countries. This seems reasonable given that transmission of disease is not immediate.

Our main definition of "neighbors" is two countries sharing a land border. This definition rules out island countries (Australia, New Zealand, Japan, Madagascar, Jamaica, etc.), but the sample attrition is limited.

A valid instrument requires two criteria to be satisfied. First, relevance requires that the instrument is a sufficiently strong predictor of the explanatory variable. Relevance is an empirical issue, addressed below. Second, exclusion restriction requires that the instrument does not have a direct effect on the outcome variable except through the explanatory variable. In other words, it requires that social unrest in country A is affected by epidemics in its neighbor, country B, via only the increased risk of an epidemic in country A. There is one obvious challenge to this assumption. An epidemic in country B might affect domestic unrest, and we have ample evidence that unrest

[9] Another way is to argue that the nature of a specific shock makes some element of the variation quasi-exogenous. For example, in Section 3 we argue that the cross-country timing of the COVID pandemic was independent of individual country circumstances.

also spills over internationally. Thus, an epidemic in country B might induce (or suppress) unrest in country B, which is turn is correlated with unrest in country A. Such a link would invalidate the exclusion restriction. To address this, in all our instrumental variable specifications, we also include unrest in neighboring countries. As a result, we satisfy a conditional exclusion restriction: the instrument only explains variation after accounting for the effect via this channel. In this case, the identifying assumption is that there is no other direct link from epidemics in country B to unrest in country A. Put differently, in order to disbelieve our identifying assumptions, one would need to identify some channel by which epidemics in country B affect unrest in country A *without* being correlated with unrest in country B. Such a link seems highly improbably, if not implausible.

The idea of using regional waves as a proxy for exogenous movements in the explanatory variable is not new. For example, Acemoglu et al. (2019) use regional waves of democratization as an instrument when assessing whether democratic reforms have a positive impact on growth. There, similar issues of spillovers there are a challenge to the exclusion restriction (e.g., a more democratic neighbor may increase demand for a country's exports, causing a direct channel from instrument to outcome) and are addressed similarly using controls for these channels.

We now turn to the issue of instrument strength. As is well known, instrumental variables estimates are unreliable when instrument is not a sufficiently strong predictor in the first stage regression (see Isaiah et al. (2018) for a recent overview). Table 4 reports the first stage regression results, which estimate equation (5) for $L = 0$. That is, we instrument using only concurrent neighboring epidemics. The relationship between epidemics in neighboring countries seems to be stable and statistically highly significant. In almost all cases, an epidemic in a neighboring country increases the likelihood of a domestic epidemic by around 8 percent. This is true even when including the full set of controls used in the main regression. A formal test of the strength of instruments is given in the two lines labeled "Instrument F-stat". These report F-statistics of the hypothesis test that coefficient on the instrument is zero. The first shows the conventional (homoskedastic) F-statistic often used as a test of instrument strength. The second corrects for correlated errors using a double-clustered measure of the variance-covariance matrix. The standard rule of thumb (Stock and Yogo (2002)) is that the F-statistic should be in excess of 10, and that above 20 is preferred. This is clearly satisfied in almost all cases: the only exception is when time fixed effects (in specification 6) are included. Unsurprisingly, time fixed absorb much of the variation. This should not be surprising – it is the nature of epidemics is that they affect multiple countries at the same time.

As a check of the strength of our instrument for longer lags, we also report Table 5, which shows that the instrument is slightly weaker when using more lags of epidemics in neighboring countries. Accordingly, our preferred instrumental variables specification is to use $L = 0$, although we investigate alternative in robustness exercises.

Figure 6 presents two instrumental variables estimates of the dynamic effect of epidemics on unrest. We include 3 lags of past epidemics and neighboring unrest, and 6 lags of domestic unrest. This guarantees that pre-pandemic responses are statistically indistinguishable from zero. That is,

Table 4: Intrumental variables, 1st stage regression for $L = 0$

	Dependent variable: Contemporaneous epidemic indicator					
	(1)	(2)	(3)	(4)	(5)	(6)
Epidemic in neighboring country	0.082***	0.080***	0.080***	0.080***	0.083***	0.053***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.015)	(0.014)
Domestic epidemic, 1 lag		0.028**	0.028**	0.028**	0.021	0.018
		(0.014)	(0.014)	(0.014)	(0.020)	(0.021)
Domestic epidemic, 2 lags		0.014*	0.014*	0.014*	0.006	0.003
		(0.008)	(0.008)	(0.008)	(0.013)	(0.013)
Domestic epidemic, 3 lags		0.037**	0.037**	0.037**	0.028	0.027
		(0.017)	(0.017)	(0.017)	(0.027)	(0.026)
Domestic unrest, 1 lag			0.002	0.002	0.002	0.004
			(0.008)	(0.008)	(0.008)	(0.008)
Domestic unrest, 2 lags			0.007	0.007	0.008	0.008
			(0.007)	(0.007)	(0.008)	(0.008)
Domestic unrest, 3 lags			−0.014***	−0.014***	−0.017***	−0.015***
			(0.003)	(0.004)	(0.004)	(0.004)
Unrest in neighbor				−0.004	−0.003	−0.001
				(0.003)	(0.003)	(0.003)
Unrest in neighbor				0.001	0.0002	0.002
				(0.003)	(0.004)	(0.004)
Unrest in neighbor, 1 lag				−0.001	−0.0001	0.002
				(0.004)	(0.005)	(0.005)
Unrest in neighbor, 2 lags				0.002	−0.001	0.002
				(0.003)	(0.003)	(0.004)
Unrest in neighbor, 3 lags					−0.00003	−0.00000
					(0.00003)	(0.00003)
Months since last domestic unrest					−0.00003	−0.00004
					(0.00003)	(0.00004)
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes
Instrument F-stat (homoskedastic)	789.9	740.5	740.8	739	361.2	133.5
Instrument F-stat (heteroskedastic)	47.6	45.1	45	45	31.6	14.4
Observations	40,768	40,432	40,432	40,432	20,506	20,506
R^2	0.019	0.022	0.022	0.022	0.020	0.009
Adjusted R^2	0.016	0.019	0.019	0.019	0.014	−0.015

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered by time and country shown in parenthesis. Instrument F-stats show tests for hypothesis that epidemic in the neighboring country is zero. Homoskedastic F-test uses a spherical variance-covariance matrix, whereas the heteroskedastic one uses a double-clustered estimator.

we have sufficient controls to avoid selection on the outcome variable. The two panels of Figure 6 present results for different numbers of lagged instruments. The main result common to both versions is that the likelihood of unrest declines statistically significantly following an epidemic, but that the magnitude of the response is much large than for the OLS estimates. The decline in the likelihood of unrest is around 2.5 percent at the longest horizons – around ten times the estimated OLS impact. The duration of the impact mirrors the findings for the OLS case, with effects decreasing only slightly at the 5-year mark.

The large difference between the two estimates likely reflects the fact that the IV approach estimates a local average treatment effect whereas the OLS approach estimates an average treatment effect. The intuition for this is that by projecting the independent variable $x_{i,t}$ onto the instrument $z_{i,t}$ we not only select plausibly exogenous (domestic) epidemic shocks, but also ones that are just severer. Epidemics which affect multiple contiguous countries are presumably larger and more dangerous than those which do not. In contrast, the OLS approach estimates the average effect across all epidemics in the sample. This is not to say that one approach is right and the other wrong, just that they estimate slightly different things. The OLS estimates are perhaps more broadly relevant but less convincingly identified, whereas the IV estimates are better-identified but only relevant to the severer epidemics. A further piece of evidence in favor of this interpretation is that the sign and approximate duration of the two estimates are similar between the two methods, just that the y-axis units are different.

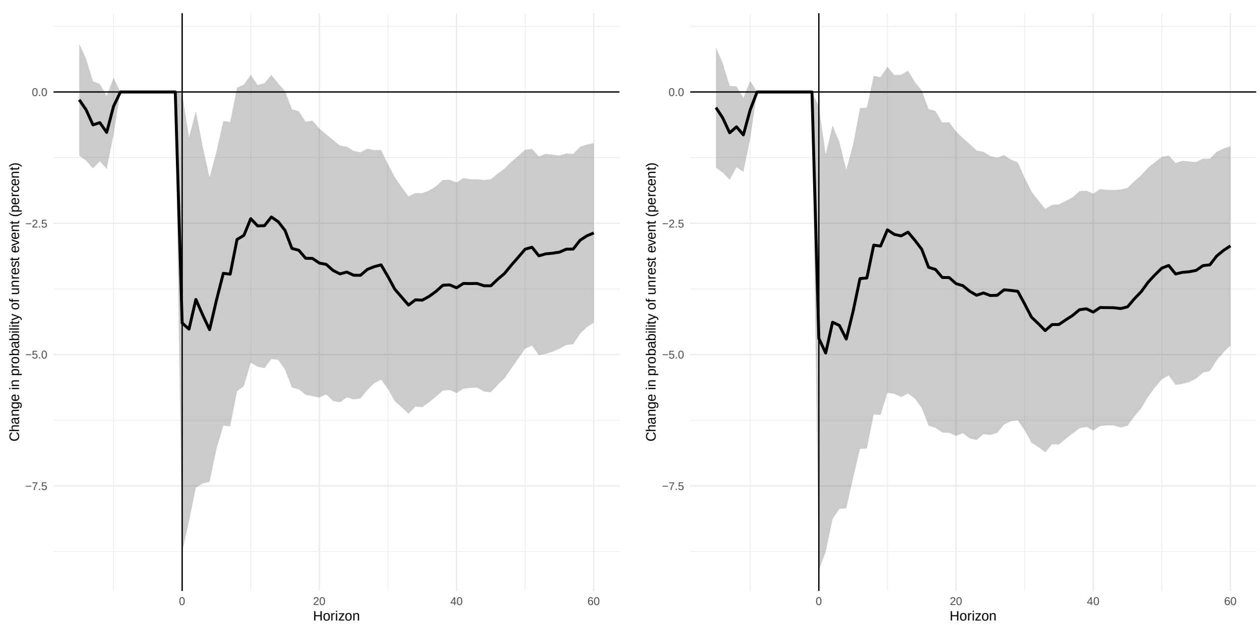

(a) Preferred specification, $L = 0$

(b) Alternative specification, $L = 2$

Figure 6: Local projection impulse responses using instrumental variables, effect of an epidemic event on the probability of social unrest, $K = 6$, $N = M = 3$. Shaded regions show 95 percent confidence intervals using standard errors which are double-clustered by both country and time period.

	(1)	(2)	(3)	(4)	(5)	(6)
$L = 0$	47.6	45.1	45	45	31.6	14.4
$L = 1$	28.5	26.4	26.3	26.3	17.5	8.2
$L = 2$	19.7	18.5	18.4	18.4	12.7	7.8
$L = 3$	16.3	14.8	14.8	14.7	10.8	9.8
Months lagged Epidemic	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Note: Table reports F-statistics for the hypotheis that all lags of the instrument are zero. F-statistics are computed using a double cluster robust heteroskedastic variance estimator.

Table 5: Heteroskedastic Instrument F-statistics

5 Conclusion

We show that, perhaps contrary to expectations, epidemics are followed by a persistently lower likelihood of social unrest in the short- and medium run. This result holds true for a large sample of epidemics since the 1980s and for the COVID-19 episode. It also seems to be a unique feature of epidemics. Natural disasters do not exhibit a similar relationship with unrest.

At first glance, this finding of a negative relationship between epidemics and unrest seems to run contrary to a recurrent view in historic accounts – that epidemics have social scarring effects, sowing seeds of social instability. We posit two likely reasons for this difference. First, it reflects the different time horizons analyzed. Our approach is more advantageous in discerning effects over the short- and medium run whereas historical accounts may be better suited for drawing inferences for a chain of events in the long run. Second, we approach the subject systematically with a large dataset of events. We are not selecting particular episodes for narrative interest or other properties. This differentiates our work from studies of individual events, where one is concern about selection of the individual events. They might be studied because that happen to be the few where unrest and epidemics go hand-in-hand, driven not by a simple causal link but instead by other factors, and so one cannot draw broader conclusions from them.

The great Parisian cholera of of 1832, where we began, serves to illustrate this point. The June Rebellion was but a relatively minor disturbance during a tumultuous century of French history. Between 1790 and 1879, the French government was overthrown at least nine times.[10] That there was unrest in this time is not surprising. That there was a cholera outbreak should not be either – sanitation in cities in the 1830s was rudimentary at best. But to say that one caused the other – while a plausible narrative in isolation – might not be as convincing in a more systematic analysis which accounts for other drivers, such as country-specific fixed factors.

[10] Our count of violent or extra-constitutional changes of power is: 1) the end of the monarchy in 1792 and the establishment of the first republic, 2) the directory in 1795, 3) Napoleon's coup 1899, 4) the Bourbon restoration of 1814, their 5) brief overthrow and 6) re-restoration when Napoleon returned from exile during the hundred days, 7) the revolution of 1848 and proclamation of the second republic, 8) the establishment of the third empire in 1852, and 9) its replacement by the third republic in 1870.

References

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2019). Democracy does cause growth. *Journal of political economy 127*(1), 47–100.

Aisen, A. and F. J. Veiga (2013). How does political instability affect economic growth? *European Journal of Political Economy 29*, 151–167.

Alesina, A., S. Özler, N. Roubini, and P. Swagel (1996). Political instability and economic growth. *Journal of Economic growth 1*(2), 189–211.

Alesina, A. and R. Perotti (1996). Income distribution, political instability, and investment. *European economic review 40*(6), 1203–1228.

Barbera, S. and M. O. Jackson (2019). A model of protests, revolution, and information. *Revolution, and Information (October 2019).*

Barrett, P., M. Appendino, K. Nguyen, and J. de Leon Miranda (2020). Measuring social unrest using media reports. *IMF Working Papers 20*(129).

Barrett, P., M. Bondar, S. Chen, M. Chivakul, D. O. Igan, and M. S. M. Peria (2021). Pricing protest: The response of financial markets to social unrest. *IMF Working Papers 21*(079).

Bernal-Verdugo, L. E., D. Furceri, and D. M. Guillaume (2013). *The dynamic effect of social and political instability on output: the role of reforms*, Volume 13. International Monetary Fund.

Censolo, R. and M. Morelli (2020). Covid-19 and the potential consequences for social stability. *Peace Economics, Peace Science and Public Policy 26*(3).

Cervellati, M., E. Esposito, U. Sunde, and S. Valmori (2018). Long-term exposure to malaria and violence in africa. *Economic Policy 33*(95), 403–446.

Cervellati, M., U. Sunde, and S. Valmori (2017). Pathogens, weather shocks and civil conflicts. *The Economic Journal 127*(607), 2581–2616.

Chen, S., D. Igan, N. Pierri, and A. F. Presbitero (2020). Tracking the economic impact of covid-19 and mitigation policies in europe and the united states. *IMF Special Series on COVID-19*.

Collier, P., A. Hoeffler, and D. Rohner (2009). Beyond greed and grievance: feasibility and civil war. *oxford Economic papers 61*(1), 1–27.

Deverell, W. F. (2004). *Whitewashed adobe.* University of California Press.

Elledge, J. (2020, April). Revolts and revolutions: how history reveals the ways coronavirus could change our world forever. *Prospect*.

Fearon, J. D. and D. D. Laitin (2003). Ethnicity, insurgency, and civil war. *American political science review 97*(1), 75–90.

Francke, M. and M. Korevaar (2021). Housing markets in a pandemic: Evidence from historical outbreaks. *Journal of Urban Economics 123*, 103333.

Hogarth, R. A. (2017). *Medicalizing blackness: Making racial difference in the Atlantic world, 1780-1840.* UNC Press Books.

Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica 62*(2), 467–475.

Isaiah, A., S. James, S. Liyang, et al. (2018). Weak instruments in iv regression: Theory and practice. *Annual Review of Economics*.

Jedwab, R., A. M. Khan, J. Russ, and E. D. Zaveri (2021). Epidemics, pandemics, and social conflict: Lessons from the past and possible scenarios for covid-19. *World Development 147*, 105629.

Jong-A-Pin, R. (2009). On the measurement of political instability and its impact on economic growth. *European Journal of Political Economy 25*(1), 15–29.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review 95*(1), 161–182.

Miguel, E., S. Satyanath, and E. Sergenti (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of Political Economy 112*(4), 725–753.

Ponticelli, J. and H.-J. Voth (2020). Austerity and anarchy: Budget cuts and social unrest in europe, 1919–2008. *Journal of Comparative Economics 48*(1), 1–19.

Randall, D. K. (2019). *Black death at the Golden Gate: The race to save America from the bubonic plague*. WW Norton & Company.

Redl, C. and S. Hlatshwayo (2021). Forecasting social unrest: A machine learning approach. *IMF Working Papers 21*(263).

Roberts, N. F. (2020, Mar). History - and psychology - predict riots and protests amid coronavirus pandemic lockdowns. *Forbes*.

Snowden, F. M. (2019). *Epidemics and society*. Yale University Press.

Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear iv regression.

Wade, S. (2020, June). Uprisings after pandemics have happened before just look at the english peasant revolt of 1381. *The Conversation*.

Disaster type	Number	Avg. deaths	People Affected	Damage (USD)	Mortality (%)
Flood	4096	48.76	751467.45	187812.29	0.01
Storm	2942	139.67	321082.16	483299.87	0.04
Epidemic	1235	163.64	18993.94	0.00	0.86
Earthquake	819	1007.25	173415.43	903272.18	0.58
Landslide	523	50.38	13025.84	15187.00	0.39
Extreme temperature	522	335.38	198391.21	105692.99	0.17
Drought	472	51.64	3671603.49	322349.18	0.00
Wildfire	340	6.89	19438.90	366170.39	0.04
Volcanic activity	155	16.61	48054.07	14949.31	0.03
Insect infestation	30	0.00	93406.67	7640.00	0.00
Mass movement (dry)	25	41.80	543.20	320.00	7.70
Animal accident	1	12.00	5.00	0.00	240.00
Impact	1	0.00	301491.00	33000.00	0.00

Table 6: EM-DAT disasters since 1990

Figure 7: Disasters and social unrest events, 1990-2019.

	(1)	(2)	(3)	(4)
	Likelihood of unrest (%)			
COVID (first case)		-0.0229	-0.0228	-0.0216
		(0.015)	(0.015)	(0.015)
COVID (1-7 days after first case)		-0.0178*	-0.0178*	-0.0169
		(0.011)	(0.011)	(0.011)
COVID (8-30 days after first case)		-0.0121	-0.0122	-0.0118
		(0.011)	(0.011)	(0.011)
COVID (31-180 days after first case)		0.0013	0.0010	0.0009
		(0.014)	(0.013)	(0.014)
COVID (1-7 days before first case)		-0.0363*	-0.0346	-0.0353*
		(0.021)	(0.022)	(0.021)
COVID (8-30 days before first case)		-0.0036	-0.0024	-0.0031
		(0.015)	(0.015)	(0.015)
Lockdown (during)	-0.0356**	-0.0361**	-0.0384*	-0.0389**
	(0.016)	(0.017)	(0.022)	(0.016)
Lockdown (1-7 days after lockdown)	0.0831	0.0830	0.0809	0.0816
	(0.087)	(0.087)	(0.087)	(0.087)
Lockdown (8-30 days after lockdown)	0.0040	0.0037	0.0017	0.0027
	(0.035)	(0.035)	(0.039)	(0.035)
Lockdown (31-180 days after lockdown)	-0.0085	-0.0087	-0.0105	-0.0092
	(0.019)	(0.019)	(0.022)	(0.019)
Lockdown (1-7 days before lockdown)	-0.0179	-0.0148	-0.0159	-0.0148
	(0.015)	(0.018)	(0.019)	(0.018)
Lockdown (8-30 days before lockdown)	0.0144	0.0202	0.0198	0.0204
	(0.038)	(0.040)	(0.040)	(0.040)
Log of confirmed COVID cases			0.0004	
			(0.002)	
Mobility				-0.0002
				(0.000)
Constant	0.0358***	0.0365***	0.0368***	0.0359***
	(0.005)	(0.005)	(0.006)	(0.006)
Observations	155,169	155,169	155,169	155,169
R-squared	0.007	0.007	0.007	0.007
Country FE	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes

Table 7: COVID-19 and social unrest: robustness with a lower lockdown threshold

Notes: This table shows results of a dynamic DID model for the sample period of January 1, 2019 to March 31, 2021. The dependent variable is the daily likelihood (in percentage) of social unrest in a country. COVID (first case) is a indicator that takes a value of 1 on the day of a country's first confirmed COVID case. COVID (1-7 days after first case) is an indicator that takes a value of 1 between 1 and 7 days after the first confirmed case. Lockdown (during) is an indicator that takes a value of 1 during a lockdown (i.e., when the stringent index is above 0.65). Other indicators are defined similarly. Standard errors clustered at the country level are shown in parentheses. *p<0.1; **p<0.05; ***p<0.01.

	(1)	(2)	(3)	(4)
	Likelihood of unrest (%)			
COVID (first case)		-0.0154*	-0.0159*	-0.0155*
		(0.009)	(0.009)	(0.008)
COVID (1-7 days after first case)		-0.0116*	-0.0118*	-0.0117*
		(0.007)	(0.007)	(0.007)
COVID (8-30 days after first case)		-0.0120	-0.0117	-0.0120
		(0.008)	(0.008)	(0.008)
COVID (31-180 days after first case)		-0.0001	0.0008	-0.0000
		(0.012)	(0.012)	(0.013)
COVID (1-7 days before first case)		-0.0277*	-0.0306*	-0.0278*
		(0.014)	(0.016)	(0.014)
COVID (8-30 days before first case)		-0.0011	-0.0032	-0.0012
		(0.017)	(0.017)	(0.017)
Lockdown (during)	-0.0568***	-0.0574***	-0.0563***	-0.0570***
	(0.020)	(0.020)	(0.021)	(0.021)
Lockdown (1-7 days after lockdown)	-0.0348*	-0.0354*	-0.0343*	-0.0352*
	(0.018)	(0.018)	(0.019)	(0.018)
Lockdown (8-30 days after lockdown)	-0.0455**	-0.0458**	-0.0447**	-0.0457**
	(0.020)	(0.020)	(0.021)	(0.020)
Lockdown (31-180 days after lockdown)	-0.0230	-0.0232	-0.0219	-0.0231
	(0.025)	(0.025)	(0.026)	(0.025)
Lockdown (1-7 days before lockdown)	-0.0302**	-0.0306**	-0.0301**	-0.0305**
	(0.015)	(0.015)	(0.015)	(0.015)
Lockdown (8-30 days before lockdown)	-0.0313**	-0.0299**	-0.0297**	-0.0299**
	(0.015)	(0.015)	(0.015)	(0.015)
Log of confirmed COVID cases			-0.0006	
			(0.001)	
Mobility				0.0000
				(0.000)
Constant	0.0322***	0.0329***	0.0334***	0.0330***
	(0.002)	(0.003)	(0.003)	(0.003)
Observations	155,169	155,169	155,169	155,169
R-squared	0.007	0.007	0.007	0.007
Country FE	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes

Table 8: COVID-19 and social unrest: robustness with a higher lockdown threshold

Notes: This table shows results of a dynamic DID model for the sample period of January 1, 2019 to March 31, 2021. The dependent variable is the daily likelihood (in percentage) of social unrest in a country. COVID (first case) is a indicator that takes a value of 1 on the day of a country's first confirmed COVID case. COVID (1-7 days after first case) is an indicator that takes a value of 1 between 1 and 7 days after the first confirmed case. Lockdown (during) is an indicator that takes a value of 1 during a lockdown (i.e., when the stringent index is above 0.95). Other indicators are defined similarly. Standard errors clustered at the country level are shown in parentheses. *p<0.1; **p<0.05; ***p<0.01.

	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	−0.385	−0.383	−0.401	−0.302	−0.566	−0.471
	(0.473)	(0.476)	(0.466)	(0.503)	(0.843)	(0.846)
$h = 1$	−0.236	−0.233	−0.26	−0.16	−0.248	−0.177
	(0.492)	(0.499)	(0.492)	(0.535)	(1.003)	(0.983)
$h = 3$	−0.089	−0.083	−0.111	−0.162	−0.083	0.054
	(0.301)	(0.311)	(0.307)	(0.301)	(0.572)	(0.574)
$h = 6$	0.045	0.051	0.026	0.052	0.214	0.295
	(0.268)	(0.277)	(0.272)	(0.287)	(0.568)	(0.559)
$h = 12$	−0.06	−0.054	−0.074	−0.047	0.084	0.145
	(0.183)	(0.191)	(0.188)	(0.201)	(0.379)	(0.37)
$h = 24$	−0.047	−0.044	−0.051	−0.027	0.065	0.144
	(0.094)	(0.1)	(0.1)	(0.11)	(0.192)	(0.18)
$h = 48$	0.073	0.083	0.079	0.081	0.173	0.171
	(0.078)	(0.083)	(0.084)	(0.089)	(0.15)	(0.145)
Months lagged Drought	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Table 9: Impact of Droughts on social unrest, Local Projection Average Effect

	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.19	0.226	0.258	0.381	0.662	0.494
	(0.485)	(0.485)	(0.481)	(0.543)	(0.865)	(0.848)
$h = 1$	−0.214	−0.202	−0.168	−0.091	−0.121	−0.173
	(0.325)	(0.322)	(0.324)	(0.359)	(0.614)	(0.606)
$h = 3$	−0.179	−0.167	−0.142	−0.113	−0.037	−0.172
	(0.252)	(0.249)	(0.251)	(0.272)	(0.509)	(0.508)
$h = 6$	−0.22	−0.214	−0.204	−0.22	−0.039	−0.134
	(0.188)	(0.183)	(0.181)	(0.182)	(0.35)	(0.364)
$h = 12$	−0.117	−0.114	−0.105	−0.129	0.151	0.188
	(0.153)	(0.149)	(0.146)	(0.148)	(0.288)	(0.274)
$h = 24$	−0.071	−0.07	−0.066	−0.085	0.11	0.107
	(0.121)	(0.118)	(0.114)	(0.116)	(0.204)	(0.187)
$h = 48$	−0.101*	−0.102*	−0.101*	−0.09*	0.088	0.111
	(0.068)	(0.066)	(0.065)	(0.07)	(0.136)	(0.119)
Months lagged Earthquake	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Table 10: Impact of Earthquakes on social unrest, Local Projection Average Effect

	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	−0.162	−0.148	−0.143	−0.092	−0.082	−0.102
	(0.226)	(0.224)	(0.225)	(0.247)	(0.433)	(0.438)
$h = 1$	−0.161	−0.146	−0.133	−0.15	−0.111	−0.143
	(0.163)	(0.155)	(0.155)	(0.167)	(0.286)	(0.297)
$h = 3$	−0.149	−0.137	−0.125	−0.132	−0.264*	−0.325**
	(0.128)	(0.118)	(0.117)	(0.123)	(0.191)	(0.191)
$h = 6$	−0.044	−0.039	−0.03	−0.043	−0.15	−0.191
	(0.099)	(0.092)	(0.091)	(0.097)	(0.155)	(0.16)
$h = 12$	−0.004	−0.007	0.001	−0.004	−0.115	−0.091
	(0.083)	(0.075)	(0.075)	(0.078)	(0.122)	(0.121)
$h = 24$	−0.027	−0.03	−0.025	−0.021	−0.188**	−0.153**
	(0.063)	(0.057)	(0.056)	(0.058)	(0.087)	(0.084)
$h = 48$	0.02	0.015	0.016	0.027	−0.148**	−0.128**
	(0.056)	(0.05)	(0.05)	(0.046)	(0.066)	(0.066)
Months lagged Flood	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Table 11: Impact of Floods on social unrest, Local Projection Average Effect

	(1)	(2)	(3)	(4)	(5)	(6)
$h = 0$	0.013	0.022	0.001	0.053	0.333	0.276
	(0.228)	(0.226)	(0.227)	(0.264)	(0.493)	(0.517)
$h = 1$	−0.021	−0.047	−0.051	−0.019	0.148	0.164
	(0.182)	(0.171)	(0.171)	(0.197)	(0.353)	(0.367)
$h = 3$	−0.02	−0.033	−0.036	−0.021	0.193	0.171
	(0.142)	(0.134)	(0.131)	(0.151)	(0.287)	(0.299)
$h = 6$	−0.045	−0.046	−0.045	−0.122	−0.197	−0.222
	(0.111)	(0.107)	(0.105)	(0.117)	(0.202)	(0.197)
$h = 12$	−0.015	−0.016	−0.013	−0.039	−0.103	−0.12
	(0.09)	(0.087)	(0.086)	(0.097)	(0.172)	(0.15)
$h = 24$	−0.031	−0.028	−0.027	−0.065	−0.112	−0.11
	(0.075)	(0.073)	(0.072)	(0.09)	(0.14)	(0.124)
$h = 48$	−0.115**	−0.106**	−0.105**	−0.141**	−0.183**	−0.154**
	(0.058)	(0.055)	(0.055)	(0.066)	(0.096)	(0.085)
Months lagged Storm	0	3	3	3	3	3
Months lagged unrest	0	0	3	3	3	3
Months lagged unrest, neighbors	0	0	0	3	3	3
Months since last unrest	No	No	No	No	Yes	Yes
Months since last unrest, neighbors	No	No	No	No	Yes	Yes
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes

Table 12: Impact of Storms on social unrest, Local Projection Average Effect