

# Is Ethnic Violence Self-Perpetuating? Quasi-Experimental Evidence from Hindu-Muslim Riots in India

Sam van Noort\*      Tanushree Goyal†

July 15, 2025

## Abstract

Ethnic riots tend to occur in the same places over time. We study whether this serial correlation exists because ethnic riots tend to be self-perpetuating or because both past and future riots are caused by the same underlying factors that persist through time. To answer this question, we leverage the fact that the timing of major Hindu festivals in India is exogenously determined by the lunar calendar and that when a major Hindu festival happens to fall on a Friday—the principal day Muslims attend mosque—the likelihood of a Hindu-Muslim riot increases significantly. Using this instrument, we find that the well-documented serial correlation in Hindu-Muslim riots disappears entirely (T= 1950-2006). This suggests that the observed recurrence of riots is not driven by the riots themselves, but by underlying conditions that remain unaddressed. Once these confounding factors are accounted for, we find no “additional” effect of past riots on future riots.

**Keywords:** Ethnic Violence, India, Hindu-Muslim relations, Natural Experiment

**Words:** 10,629 (excl. abstract)

---

\*Lecturer at Princeton University. E-mail address: [samvannoort@princeton.edu](mailto:samvannoort@princeton.edu). Website: [sites.google.com/view/samvannoort](https://sites.google.com/view/samvannoort).

†Assistant professor at Princeton University. E-mail address: [tgoyal@princeton.edu](mailto:tgoyal@princeton.edu). Website: [www.tanushreegoyal.com/](http://www.tanushreegoyal.com/).

# Introduction

Ethnic riots are a frequent and devastating occurrence in many developing countries. While many ethnic riots go unreported or undercounted, a reasonable estimate is that between 300,000 and 2.5 million people have died in ethnic riots since World War II. The Partition riots in India (1947) alone are estimated to have claimed between 200,000 and 2 million lives. In Rwanda, anti-Tutsi riots between 1959 and 1964 resulted in over 20,000 deaths. In Indonesia, a wave of ethnic riots between 1998 and 2002 caused an estimated 8,500 to 15,000 deaths, while in Nigeria, ethnic riots have claimed between 10,000 and 20,000 lives since independence. More recently, ethnic riots in Myanmar (2017) and Ethiopia (2021–2022) led to at least 16,700 deaths.

Although there is little consensus about the causes of ethnic riots, one of the most robust findings in the existing literature is that the risk of ethnic riots is substantially higher in communities that have experienced ethnic riots in the past (Horowitz, 2001; Varshney, 2001, 2002; Wilkinson, 2009).

We ask whether this serial correlation in ethnic riots reflect a causal effect—meaning that ethnic riots in themselves make future riots more likely—or whether the serial correlation is entirely driven by other factors, which are not alleviated through rioting, and which therefore also tend to cause future riots.

Knowing whether ethnic riots are self-perpetuating is important because if ethnic riots themselves increases the likelihood of future riots, the implications would be profound. It would suggest, for example, that addressing the original triggers of ethnic riots will not in itself be sufficient to prevent further violence. Instead, communities may become trapped in a “riot trap”—a self-perpetuating cycle of violence similar to the “conflict trap” observed in civil wars (Collier et al., 2003; Collier, 2008).

Studying whether ethnic riots are self-perpetuating is challenging given that the same underlying conditions that ignited riots initially may persist and therefore also cause subsequent riots, creating a classic confounding problem (Collier and Sambanis, 2002; Hazlett,

2019).<sup>1</sup> Given our imperfect knowledge of what all the relevant confounding factors are, and how to measure them sufficiently precisely, and given the fact that randomly assigning ethnic violence ourselves is both unethical and impossible, causal identification has to effectively come from a credible quasi-experiment, which accounts for confounding by design.

We are, to the best of our knowledge, the first to bring such quasi-experimental evidence to bear on the fundamental question whether ethnic riots tend to be self-perpetuating.

To do so, we focus on the important case of Hindu-Muslim riots in India.<sup>2</sup> Hindu-Muslim riots in India have left more than 40,000 citizens dead or injured since 1947. Hindu-Muslim riots not only comprise an important case in the ethnic violence literature (Nellis, Weaver and

---

<sup>1</sup>The existing literature highlights a large number of relatively persistent regular causes of ethnic riots. The most important of which are: poverty (Collier and Hoeffler, 1998, 2002), competition for jobs between migrants and native populations (Bhavnani and Lacina, 2015), negative economic shocks (Bohlken and Sergenti, 2010; Fearon and Laitin, 2003; Miguel, Satyanath and Sergenti, 2004), geospatial segregation that undermines interethnic trust and facilitates intraethnic mobilization (Eubank et al., 2019; Kasara, 2017), the absence of institutions for the peaceful expression of grievances (De Mesquita et al., 2008), socio-economic inequalities between ethnic groups (Mitra and Ray, 2014), the presence of reinforcing—as opposed to cross-cutting—cleavages (Paler et al., 2020), discriminatory government policies (Abbs, 2021), and the strategic interests of political actors in provoking ethnic violence (Brass, 1996, 2011; Wilkinson, 2004).

<sup>2</sup>There are several definitions of riots. We define an ethnic riot as: “[...] an intense, sudden, though not necessary wholly unplanned, lethal attack by civilian members of one ethnic group on civilian members of another ethnic group, the victims chosen because of their group membership.” (Horowitz, 2001, p. 1) We see ethnic riots as one type of ethnic violence. Other types of ethnic violence are: violent protest demonstrations, pogroms, feuds, lynchings, genocides, terrorist attacks, gang assaults, and ethnic warfare. See Horowitz (2001) for what distinguishes these different forms of ethnic violence.

---

Rosenzweig, 2016), but the reoccurrence of Hindu-Muslim riots in India has also regularly been explained by a self-perpetuating logic (e.g., Bohlken and Sergenti (2010), Varshney (2002), and Horowitz (2001)).

Crucially, we can credibly identify the potentially self-perpetuating effect of Hindu-Muslim riots in India by employing Iyer and Shrivastava's (2018) innovative instrumental variable design, which exploits that the incidental co-occurrence of an important Hindu festival falling on a Friday—the holy day for Muslims—introduces plausibly exogenous variation in Hindu-Muslim riots. The probability of Hindu-Muslim riots in India rises sharply when a major Hindu festival falls on a Friday. On such days large numbers of Hindus and Muslims tend to congregate in public spaces for the celebration of religious ceremonies. This leads to increased contestation of public space and alleviates some of the collective action problems inherent in mobilizing large crowds. This in turn increases the probability of the occurrence of a Hindu-Muslim riot by 39 percentage points. Importantly, the dates of major Hindu festivals are exogenously determined by the Hindu lunar calendar, which is based on the position of the sun, moon and earth relative to each other. Following Iyer and Shrivastava (2018) we exploit this plausibly exogenous process to obtain a causal estimate of the self-perpetuating effect of ethnic violence by regressing future Hindu-Muslim riots on past Hindu-Muslim riots, while instrumenting the latter with whether in that same year an important Hindu festival occurred on a Friday, or not.

Our design relies on the identifying assumption that Hindu festivals occurring on a Friday (as opposed to any other day of the week) do not affect future Hindu-Muslim riots through channels other than Hindu-Muslim riots in the current year. Like with any instrumental variable design it is not possible to verify this exclusion restriction empirically. However, we find this assumption plausible. We find that years with and without a Hindu festival on Friday do not differ substantively or statistically significantly on a range of other factors which may affect riots but should not be correlated with Friday festivals. Also, to the best of our knowledge, there is no existing evidence that suggests that Friday festivals may directly

affect future riots other than through current riots.

Initially, we find that Hindu-Muslim riots are strongly serially correlated across 339 Indian districts over the 1950 to 2006 period. When we however use the Hindu festival instrument the effect of past riots is null, or, if anything, slightly negative. We establish that this finding is robust to different lag specifications, is not driven by riot under-reporting or weak instrumental variable bias, and is robust to alternative operationalizations of Hindu-Muslim riots. We interpret this result as that Hindu-Muslim riots in India do *not* in itself cause further violence. Rather, violence is caused by other factors which are not alleviated through rioting and therefore also tend to cause future riots. Once these confounding factors are controlled for there is no “additional” effect of ethnic riots itself.

These results have a number of important implications. Most significantly, our results imply that communities that have experienced ethnic riots in the past are not destined to a higher likelihood of riots in the future *due* to their violent past. This is a hopeful message. The past is after all the past, and if it were to be true that ethnic riots are self-perpetuating, this would imply that little could be done to bring down the risk of violence once ethnic violence has found its way into society (i.e., besides trying to change the collective memory of the past, or hoping that violence control measures and time would slowly dissipate the forces that lead to self-perpetuation). The results also imply, however, that in many cases the underlying causes of ethnic riots are insufficiently addressed after ethnic violence has broken out, so that the same factors are allowed to cause additional violence in the future (hence leading to the serial correlation that we observe in the data).

Throughout the paper we acknowledge that our results are based on one large, but nonetheless, specific country and one specific type of ethnic violence. Although there is some evidence that different types of violence have similar effects (Bateson, 2012), and that the causes and institutional machinery underlying riots are often quite similar across countries (Wilkinson, 2009), future empirical research is necessary to establish how our findings travel to other countries, other times, and other types of ethnic violence.

Studies that use ethnic riots as the independent variable are rare, and our paper joins a small but flourishing body of scholarship that exploit quasi-experiments to examine the consequences of ethnic violence (Aidt and Leon, 2016; Beber, Roessler and Scacco, 2014; Hager, Krakowski and Schaub, 2019). While exposure to ethnic violence can deepen divisions and encourage revenge seeking, our findings resonate with research which suggests that these type of responses may not always predominate. Exposure to political violence can instead also lead to community leadership (Blattman, 2009), moderation and a willingness to appease and make concessions to the outgroup (Gould and Klor, 2010; Hazlett, 2019; Hartman and Morse, 2018; Lyall, 2009), and previous occurrences of riots may lead to investments in better riot-control measures which in turn may help suppress future riots (Nellis and Siddiqui, 2018). It is plausible that Hindu-Muslim riots in India trigger some of these mechanisms which are then sufficient to offset their self-perpetuating tendencies.

The paper is structured as follows. In the following section, we develop a conceptual framework that shows how—based on theory and existing empirical evidence—one can reasonably arrive at contrasting predictions about whether ethnic riots are self-perpetuating, or not. In the third section, we discuss our data and identification strategy. We then report the results. In the last section, we conclude with suggestions for future research.

## Theory and Literature Review

In this section, we start by outlining the reasons why ethnic riots may or may not be self-perpetuating.

Although the existing literature has developed a rich and rigorous conceptual framework for understanding the self-perpetuating nature of civil wars—often referred to as the “conflict trap” literature (e.g., Collier et al. (2003); Collier (2008))—we lack similarly well-developed theoretical frameworks for ethnic riots.

The distinction between civil wars and ethnic riots is important in this context because the

“conflict trap” literature typically explains the recurrence of violence through mechanisms such as the buildup of organizational capacity, institutional erosion, and the intensification of conditions favorable to insurgency (Collier et al., 2003; Collier, 2008). These mechanisms, however, are largely inapplicable to ethnic riots. Furthermore, the conflict trap framework often assumes the presence of a formal agreement or ceasefire that temporarily halts violence—conditions that are rarely relevant in the case of riots, where violence tends to erupt spontaneously and without a clear endpoint.<sup>3</sup>

Based on theoretical reasoning and a review of the existing literature, we conclude that it is a priori unclear whether ethnic riots are, on average, self-perpetuating.

## Ethnic Riots and Self-Perpetuating Dynamics

We identify three key causal mechanisms that may render ethnic riots self-perpetuating: (1) shifts in the strategic incentives of political entrepreneurs to sustain violence, (2) the formation of durable social networks that facilitate mobilization, and (3) the psychological imprint of fear and grievance.

First, an initial riot may shift the calculations of political actors in ways that could encourage recurrence. Ethnic violence can be politically useful: it polarizes society along group lines (Allie, 2024), allowing ethno-nationalist entrepreneurs to mobilize support by stoking ethnic cleavages. In India, for instance, electoral incentives have repeatedly led politicians to foment or tolerate Hindu–Muslim riots as a strategy to consolidate their ethnic vote bases (Brass, 1996, 2011). When within-group economic inequality is low or inter-group inequality is diminishing—conditions under which ethnicity loses salience—leaders may have a heightened incentive to instigate riots to refocus public identity on ethnic loyalties (Bulutgil and Prasad, 2023). In such scenarios, violence may serve as a tool to reinforce the salience of ethnicity and improve an ethnonationalist party’s electoral prospects.

---

<sup>3</sup>See Bulutgil and Prasad (2023) on how the mechanisms underlying communal riots differ from those associated with mass violence, such as civil wars, genocide, and ethnic cleansing.

More generally, the decision of the state to allow or prevent riots may often hinge on political interests: when leaders benefit from polarization, they may permit riots to occur rather than stopping them. As [Wilkinson \(2004\)](#) showed in his study of ethnic violence in India, officials often stand down the police when ruling elites calculate that riots will aid their political fortunes, whereas they clamp down on violence when their incentives favor peace. Once one riot has paid dividends, the temptation is to reuse the tactic.

Over time, a pattern of orchestration may emerge: parties, militias, and strongmen learn that communal violence “has functional utility and benefits for a wide range of groups and organizations in society,” and they develop ready templates for instigating unrest. In his seminal study of northern India, [Brass \(1996\)](#) observes that riots there are essentially a planned, orchestrated and institutionalized phenomenon—not spontaneous outbursts, but repeatedly manufactured by networks of politicians, extremists, and criminals who form an “institutionalized riot system” in locales with endemic violence. In short, once political entrepreneurs realize that riots can serve their aims, they may invest in the capacity to spark violence again. The result is that each riot increases the supply of would-be instigators and erodes any normative or institutional barriers to future clashes.

Indeed in the context of India, [Iyer and Shrivastava \(2018\)](#) have convincingly showed that the occurrence of a Hindu-Muslim riot increases the vote share of the Hindu nationalist Bharatiya Janata Party (BJP). The BJP thus arguably has an interest in inciting riots. It is an interesting fact in this context that [Nellis, Weaver and Rosenzweig \(2016\)](#) finds that if power shifts to the Indian National Congress party—which he argues has a strong electoral incentive to halt Hindu-Muslim riots—that Hindu-Muslim riots tend to become significantly less likely.

Second, ethnic rioting may forge enduring social networks and organizational structures that may facilitate its own recurrence ([Aidt, Leon-Ablan and Satchell, 2022](#)). Riots are not only political events but also deeply social ones: they bring together mobs, vigilantes, and militia groups in collective violence. These interactions may solidify into durable networks.



Participants in a riot—local strongmen, gang leaders, youth activists, and their followers—may establish relationships and a reputation for violence that outlast any one episode. Once an ethnic clash has occurred, a community is often left with organized cadres (e.g. armed youth wings or neighborhood defense leagues) that can be quickly re-mobilized in the next crisis.

At the same time, riots reshape the broader social fabric in ways that can make future unrest more likely. Communal bloodshed may drive groups apart: frightened minorities often retreat into homogeneous enclaves, while mixed localities may undergo de facto ethnic segregation after riots as families move to “safer” areas. The loss of day-to-day interethnic contact means the civic networks that once bridged communities are severed, removing a key brake on violence. Indeed, Varshney (2002) found that Indian towns with strong interethnic civic associations were largely able to avert riots, as their networks of trust dampened rumors and brokered peace, whereas towns lacking such ties suffered recurrent clashes. When a riot tears those networks apart, future sparks may find little resistance. Ethnic violence may not only reduce trust and prosocial behavior across ethnic groups, but even within ethnic groups, potentially more generally affecting societies’ capacity to peacefully resolve societal tensions and distributional issues (Hager, Krakowski and Schaub, 2019).

Third, ethnic rioting may imprint deep psychological scars—fear, hatred, and grievance—that can fuel a self-perpetuating logic of vengeance. Riots leave behind more than physical destruction; they also create collective memories of trauma that can harden ethnic identities and biases. In the immediate wake of violence, communities on both sides experience pervasive insecurity and anger. People who have witnessed or survived an ethnic riot may come to view their rivals as implacable enemies and live in fear of future attacks. Horrific anecdotes, whether true or exaggerated, become part of the group narrative—each atrocity recounted is a spur to hatred.

As Horowitz (1985, 2001) documented across dozens of cases, riots commonly begin with rumors of aggression by the other group, feeding a climate of fear and righteous fury

that makes violent “retaliation” seem acceptable. After one riot, those same rumors and memories remain circulating in the community, priming people to expect the worst from their neighbors (Bhavnani, Findley and Kuklinski, 2009). The result is a palpable tension: each community believes it must be ready to strike if provoked. Empirical studies confirm that exposure to ethnic violence can entrench hostile attitudes. In northern Sudan, for example, northerners who lived through a deadly riot became markedly more antagonistic toward the southern minority—they showed increased support for separation and less willingness to coexist peacefully (Beber, Roessler and Scacco, 2014). Similarly, survey research in Kenya found that after the 2007 electoral violence, trust between communities plummeted (Dercon and Gutiérrez-Romero, 2012).

Fear for personal safety can also translate into political preferences that make the self-perpetuation of riots more likely: threatened populations often rally behind hardline leaders and policies. Experimental evidence from India indicates that when people feel endangered by prior violence, they develop stronger prejudice against out-groups and greater support for extremist measures, along with heightened internal cohesion born of a siege mentality (Schutte, Ruhe and and, 2023). In this way, the emotional aftermath of a riot empowers the most militant voices on each side—precisely those actors likely to provoke fresh violence.

A riot inflicts enduring wounds and resentments: families remember their dead, communities recall the indignities suffered, and these grievances can fester into a desire for retribution. Each outbreak of violence may thus sow the seeds of the next, as today’s embittered victims may become tomorrow’s perpetrators seeking to “settle the score.”

Furthermore, recurrent riots may lead to the normalization of violence in society (Horowitz, 2001). Psychological research—particularly among children and young adults—finds that repeated exposure to violence may make people less sensitive to violence, may make people more aggressive themselves, and may increase the likelihood that violence is seen as a legitimate mechanism to solve distributional issues with “outgroups” (Tarabah et al., 2016; Niwa et al., 2016). Nisbett (2018) argues that ethnic violence can lead to a “culture of

honor”, where communities unable to rely on outside protection must develop reputations of toughness and must seek vengeance as a means of deterring future attacks.

## Challenges to the Self-Perpetuation Hypothesis

While there are internally coherent reasons to believe that ethnic riots may be self-perpetuating, there are also compelling theoretical and empirical reasons to expect that this may not always be the case.

We identify three key causal mechanisms that may counteract the self-reinforcing logic of riot recurrence: (1) emotional responses that reduce support for future violence, (2) the emergence of community leadership and social capital that promotes peacebuilding, and (3) institutional responses that enhance the state’s riot-prevention capacity.

First, individual emotional and attitudinal responses to ethnic violence may, under certain conditions, reduce the likelihood of future riots. Contrary to the assumption that exposure to ethnic violence necessarily radicalizes attitudes or hardens intergroup hostility, a growing body of research suggests the opposite effect. For example, quasi-experimental evidence from Sierra Leone by [Hazlett \(2019\)](#) finds that individuals more directly exposed to violence are often *less* supportive of future violence or retaliation. Similarly, studies such as [Gould and Klor \(2010\)](#), [Lyal \(2009\)](#), and [Hartman and Morse \(2018\)](#) show that the experience of violence may make individuals more empathetic toward out-groups or more supportive of concessions to prevent renewed violence. These emotional responses—such as fear, grief, and a desire to avoid further harm—may shift preferences toward conflict mitigation and compromise.

While such emotional reactions may not directly address the structural conditions that make some regions more prone to riots—such as poverty, political exclusion, or intergroup inequality—they may help explain why ethnic riots do not always produce a reinforcing cycle of violence. In other words, even in the presence of structural risk factors, riots themselves may not exert an *additional* causal effect on future outbreaks of violence.

Second, ethnic riots may stimulate new forms of civic engagement and political mobilization that work against their recurrence. As Blattman (2009) shows, individuals exposed to violence may experience personal growth and become more politically active—often in ways that channel grievances through peaceful, institutionalized avenues. This includes not just victims of violence but also former perpetrators who come to reject further conflict. In some communities, the trauma and moral exhaustion of communal violence lead to a kind of “riot fatigue,” where the horror of previous events generates local revulsion and a strong collective desire to avoid future unrest.

Over time, this response can foster interethnic reconciliation, especially when supported by deliberate peacebuilding efforts. Initiatives such as intergroup dialogue, economic interdependence, shared civic institutions, and peace committees can help rebuild trust and create social firewalls against future violence (Jha, 2013; Varshney, 2001). These new ties are not just symbolic; they can serve as practical deterrents to mobilization by increasing the cost of renewed riots and by creating cross-cutting allegiances that reduce polarization.

In addition, the same political incentives that sometimes encourage ethnic violence can, under different conditions, lead to its suppression. When political leaders no longer benefit from ethnic polarization—whether due to changing electoral coalitions, international scrutiny, or reputational risks—the strategic supply of riot-inciting rhetoric and mobilization may decline.

Third, ethnic riots may induce governments to adopt stronger riot-prevention measures, even in places historically prone to communal violence. As Nellis, Weaver and Rosenzweig (2016) and Nellis and Siddiqui (2018) show, incumbents who are electorally incentivized to maintain interethnic peace often invest in policing, early warning systems, and legal deterrents that often successfully reduce the incidence of riots, even in riot-prone areas. In other cases, even governments without direct electoral incentives may be compelled to act—due to pressure from higher levels of government (especially when controlled by different parties), legal institutions such as courts, or political opposition groups that seek to capitalize

on the ruling party’s failure to maintain order.

## Conclusion of Theory and Literature Review

We have shown that both theoretical reasoning and existing empirical findings yield indeterminate predictions about whether ethnic riots, in and of themselves, increase the likelihood of future riots—beyond the influence of the underlying conditions that triggered the initial ethnic riot in the first place.

Simply observing a positive correlation in the timing or frequency of riots over time can also not resolve this question. While such a correlation may reflect a genuine causal relationship, it may equally result from the persistence of structural factors—such as economic inequality, political polarization, or institutional weakness—that independently cause both past and current riots.

The objective of this paper is to address this endogeneity problem by leveraging a quasi-experimental research design in the context of Hindu-Muslim riots in India from 1950 to 2006. We do not claim that evidence from a single country and a single time period can conclusively prove or disprove the broader hypothesis that ethnic riots are self-perpetuating. Nonetheless, this case is especially instructive. Not only is it one of the most extensively studied cases of ethnic rioting—with many accounts suggesting a reinforcing logic of violence (Bohlken and Sergenti, 2010; Varshney, 2002; Horowitz, 2001)—but it also offers a unique empirical opportunity to causally test the self-perpetuating riot hypothesis, using the novel identification strategy designed by Iyer and Shrivastava (2018).

# Empirical Strategy

## Data

To test the effect of ethnic riots over time we use Varshney and Wilkinson's (2006) data on the location and date of Hindu-Muslim riots across India from 1950 until 1995.<sup>4</sup> We complement this data with data from Iyer and Shrivastava (2018), who have extended Varshney and Wilkinson's (2006) data until 2006, using the exact same coding scheme. 2006 is, to the best of our knowledge, the latest year for which there is reliable publicly available data on Hindu-Muslim riots in India.

We chose to aggregate the data up to the district-level for our main analysis, but provide robustness checks with alternative aggregation levels (e.g., state-level). We choose to focus on districts, rather than towns/cities, to reduce the probability of spatial spillover effects that would violate the stable unit treatment value assumption that we require for causal identification (see below).

In terms of the time dimension, we choose to analyze the data on the year level because this is the level of temporal granularity that the existing ethnic riot literature has used to analyze the potentially self-perpetuating nature of ethnic violence (e.g., Hegre and Sambanis (2006) and Bohlken and Sergenti (2010)). In appendix A we analyze the results on the six month level and find the results unchanged.<sup>5</sup>

To measure ethnic violence in the Indian context we use as our baseline measure a dichotomous variable that takes the value 1 if a particular district-year has experienced a Hindu-Muslim riot, and 0 otherwise. We focus on this dichotomous variable as the distribu-

---

<sup>4</sup>Varshney and Wilkinson (2006) have collected this data using the Mumbai edition of the Times of India.

<sup>5</sup>Disaggregating further runs into issues related to data sparsity and single riots spanning multiple time periods, causing spurious correlation. We discuss these issues in further detail in Appendix A.

tion of Hindu-Muslim riots on the district level is almost completely dichotomous in nature (98.1% of district-years have experienced either 0 or 1 riots).<sup>6</sup> In the robustness checks we use two alternative operationalizations of Hindu-Muslim riots: (1) the total number of Hindu-Muslim riots in a particular district-year; and (2) a dummy that takes the value 1 if a riot has occurred within a 100 km radius from a district's largest city, and 0 otherwise (to account for spatial spillovers).

## Instrument

To identify the causal effect of ethnic riots over time we follow the innovative identifications strategy of Iyer and Shrivastava (2018) by regressing future Hindu-Muslim riots on past Hindu-Muslim riots, while instrumenting the latter with whether in that same year an important Hindu festival occurred on a Friday, or not.

This identification strategy relies on five key assumptions. First, Friday festivals are correlated with Hindu-Muslim riots (relevance assumption). Second, Hindu festivals do not fall on Friday because of other factors that affect Hindu-Muslim riots, and Hindu festivals do not fall on Friday because of past Hindu-Muslim riots (exogeneity assumption). Third, Friday festivals do not affect future Hindu-Muslim riots through any other channel than present Hindu-Muslim riots (excludability assumption). Fourth, Friday festivals do not lead to a decrease in the probability of Hindu-Muslim riots in any district within our sample (monotonicity assumption). Last, Friday festivals in one area do not affect the probability of Hindu-Muslim riots in any other area within the sample (stable unit treatment value assumption).

We now discuss why we believe that these five identifying assumptions are likely to hold in our specific case.

**Relevance:** The incidental concurrence of a Hindu festival falling on a Friday is strongly

---

<sup>6</sup>This fact further substantiates our decision to focus on the year-level in our baseline specifications.

correlated with Hindu-Muslim riots in India. This is likely because on such days large crowds of Hindus and Muslims tend to congregate in public spaces for the celebration of religious ceremonies. These religious ceremonies are often highly public and vocal events which can easily spark tensions between the two communities, ultimately leading to an increased risk of ethnic violence.<sup>7</sup> We are hardly the first to point out the relationship between Hindu festivals falling on Friday and Hindu-Muslim riots in India. This correlation was first hinted at in qualitative work by Horowitz (2001) and Jaffrelot (1999), among others, and established quantitatively over the 1981 to 2001 period by Iyer and Shrivastava (2018).<sup>8</sup> We provide another test of this first-stage relationship by regressing Hindu-Muslim riots on Friday festivals over the 1950 to 2006 period. Following the common standard of Staiger and Stock (1997) we regard the instrument as sufficiently strong if the Kleibergen and Paap (2006) F statistic of this regression exceeds 10.

**Exogeneity:** The date of Hindu festivals is determined by the Hindu lunisolar calendar, which is based on the interaction between monthly moon cycles and yearly solar cycles. While different subnational regions within India tend to use slightly different types of lunisolar calendars all major Hindu festivals occur on the same day throughout the country, and the specific day is always set by the position of the sun and moon as observed from earth. Given

---

<sup>7</sup>The Times of India (2018) provides an illustrative example of how the incidental co-occurrence of a Hindu festival on Friday can lead to a Hindu-Muslim riot: “According to police officials, a Holi procession was to begin on Friday around 2pm from Chaudharana Devta temple to Nagina city. When preparations for the procession were being made around 1pm, two youths threw colour on two members of another community on their way to prayers. After the violence escalated, police had to launch lathi charge and fire tear gas shells to disperse the mob.”

<sup>8</sup>See Allie (2024) for closely related work showing how incompatible ritual holidays where one religion’s rituals are at odds with another (e.g., sacrificing cows or engaging in processions with idolatry) increases the probability of a Hindu-Muslim riot in India.



that moon and sun cycles cannot be affected by Hindu-Muslim riots and cannot be co-determined together with Hindu-Muslim riots by third factors, the instrument is exogenous.

**Excludability:** We foresee two potential ways how Hindu festivals on Friday could affect future Hindu-Muslim riots through channels other than inciting riots in the current period. First, it could be the case that Friday festivals affect future Hindu-Muslim riots not through riots in the current year, but through riots in the following year(s). We account for this possibility by running robustness checks where we incrementally increase the lag length up until 20 years. Second, it could be the case that some riots go underreported and that the instrument affects future riots through these underreported riots. To account for this possibility we primarily focus on the reduced form effect (i.e., the effect of a Hindu festival on a Friday in the previous year on Hindu-Muslim riots in the current year), which would be unaffected by this potential exclusion restriction violation. Focusing on the reduced form does not sacrifice the possibility of interpreting the size of substantive effects in this case, as we are inclined to show that the self-perpetuating effect of Hindu-Muslim riots in India is essentially null.<sup>9</sup>

**Monotonicity:** We are not aware of any empirical evidence, or plausible theoretical mechanisms, through which the probability of a Hindu-Muslim riot could in some cases

---

<sup>9</sup>Note that we do not deny that it could be possible that political parties, or other actors in Indian society, are aware of the positive effect Hindu festivals on Friday have on the probability of ethnic violence, and that these actors use this strategically to incite riots. Importantly, however, such a scenario would not invalidate our identification strategy as also in that case the Hindu-Muslim riots that we use to estimate the effect on would be the result of the exogenously determined co-occurrence of Hindu festivals on Friday (working through lowering the cost of strategically inciting a riot) and would not plausibly lead to future riots through any other channel. Such a scenario would have consequences for external validity, however, if strategically incited riots have different effects on future riots than other type of riots.

systematically decrease as a result of a major Hindu festival falling on a Friday.

**Stable unit treatment value:** It could be the case that Friday festivals lead to tensions between Hindu and Muslim communities close to district borders and then cause Hindu-Muslim riots to occur in adjacent districts. Such a scenario would violate the stable unit treatment value assumption. We regard this as relatively unlikely as Hindu-Muslim riots in India tend to be highly localized events (Varshney, 2002). To nonetheless account for this possibility we run robustness checks on the state level (which are very large units, comparable with countries in other areas of the world) and analyze effects in a radius of 100 km from a district’s largest city. As shown below, this leaves our results unchanged.

## Estimation

Our baseline specification is a simple time series model with district and time fixed effects. More in particular we estimate the following model with two-stage least squares (2SLS):

$$Y_{it} = \beta_0 + \beta_1 \hat{Y}_{it-1} + \delta_i + \phi_t + \epsilon_{it} \quad (1)$$

$$Y_{it-1} = \alpha_1 + \alpha_2 F_{it-1} + \delta_i + \phi_t + v_{it} \quad (2)$$

where  $Y$  is a dummy that takes the value 1 if a Hindu-Muslim riot occurred in district  $i$  and year  $t$ , and 0 otherwise,  $\delta$  are a full set of district fixed effects,  $\phi$  are a full set of 5-year time fixed effects, and  $F$  is a dummy that takes the value 1 if at least one major Hindu festival in district  $i$  and year  $t$  fell on a Friday, and 0 otherwise.

Table 1 shows the Hindu festivals that Iyer and Shrivastava (2018) use to construct the Friday Hindu festivals instrument. As can be seen, major Hindu festivals only vary across states and over time. We nonetheless choose to focus on the district level in our baseline specification as this allows for the inclusion of unit fixed effects on a much lower level of aggregation—the average population of an Indian state is approximately 65 million, while the average population of an Indian district is approximately 1 million.<sup>10</sup> This does mean,

---

<sup>10</sup>Note that we require unit fixed effects for identification to control for any potentially

however, that our effective sample size is smaller than the absolute number of district X year observations suggests.<sup>11</sup> We take this into account by using the Driscoll and Kraay (1998) covariance matrix estimator, which adjusts the standard errors and first-stage F statistics for the cross-sectional autocorrelation that using a state-level instrument introduces.<sup>12</sup> As a robustness check we run our analysis on the state-level and find substantively similar results.

Given that two major Hindu Festivals are shared by all states (i.e., Dushehra and Diwali), our instrument varies relatively little across states at the same point in time.<sup>13</sup> This prevents us from including year fixed effects without running into a weak instrument problem. Although there appears to be no reason to believe that the dates of Hindu festivals (which are fully determined by the Hindu lunar calendar) are somehow correlated with the timing of nationwide riot shocks, we do acknowledge that this is a limitation of our empirical approach. To reduce the possibility that our results are driven by common shocks we include 5 year time fixed effects. In appendix B we take an alternative approach by adding time confounding factors that determine both why particular Hindu festivals are important in particular Indian states and why Hindu-Muslim riots occur more in some states than in others. Given that the state-level differences in the importance of different Hindu festivals has been determined long before our period of observation these factors are time-invariant within our sample, and thus can be fully accounted for by unit fixed effects.

<sup>11</sup>Decomposing the variance of our instrument in a cross-sectional and temporal dimension suggests that 12.4 percent of the variance comes from the cross-sectional dimension (i.e., because not all states following the same festivals and therefore some districts experience a Friday festival in a particular year and others not) and 87.6 percent comes from the over-time dimension (i.e., because the same festival tends to fall on different days of the week in different years).

<sup>12</sup>The Driscoll and Kraay (1998) estimator also adjusts for serial autocorrelation and heteroscedasticity.

<sup>13</sup>Note that the Driscoll and Kraay (1998) estimator also adjusts for the cross-sectional autocorrelation that occurs due to states sharing some of the same festivals.

Table 1: Major Hindu festivals per state.

State	Festivals				
Andhra Pradesh	Ramnavami	Durga ashtami	Navami	Dushehra	Diwali
Assam	Janmashtami	Durga ashtami	Navami	Dushehra	Diwali
Bihar	Holi	Ramnavami	Navami	Dushehra	Diwali
Gujarat	Holi	Ramnavami	Navami	Dushehra	Diwali
Haryana	Shivratri	Holi	Janmashtami	Dushehra	Diwali
Jammu and Kashmir	Shivratri	Ramnavami	Janmashtami	Dushehra	Diwali
Karnataka	Shivratri	Ganesh Chaturthi	Navami	Dushehra	Diwali
Kerala	Shivratri	Janmashtami	Navami	Dushehra	Diwali
Madhya Pradesh	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
Maharashtra	Ramnavami	Ganesh Chaturthi	Navami	Dushehra	Diwali
Odisha	Holi	Durga ashtami	Navami	Dushehra	Diwali
Punjab	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
Rajasthan	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
Tamil Nadu	Janmashtami	Ganesh Chaturthi	Navami	Dushehra	Diwali
Uttar Pradesh	Ramnavami	Janmashtami	Navami	Dushehra	Diwali
West Bengal	Holi	Durga ashtami	Navami	Dushehra	Diwali

*Notes:* Festivals are selected based on public holidays as declared by each state government. Data on which day of the week individual festivals fell in the 1976 to 2001 period is collected by [Lyer and Shrivastava \(2018\)](#). Data on which day of the week individual festivals fell in the 1950 to 1976 and 2001 to 2006 period has been collected by ourselves.

fixed effects for the most riot-prone years (nationwide), and find the results unchanged.

## Balance on Observables

Assuming that our instrument is valid, we would expect that the probability of Hindu-Muslim riots in India rises sharply in district-years with a major Hindu festival on a Friday, *and* that all other potential determinants of riots are balanced across years with and without a Friday Hindu festival. We test this expectation empirically by regressing our riot variables and a number of covariates on our instrument, together with district and time fixed effects.<sup>14</sup>

Table [2](#) reports the results. In support of the validity of the instrument we find that district-years with or without a major Hindu festival falling on a Friday do not differ signifi-

<sup>14</sup>Data on the covariates used for the balance on observables test is taken from the Indian Census. This data is linearly interpolated between Census years. Data for the covariates is only available from 1970 onwards.

Table 2: Balance on observables.

	(1) Riot (0-1)	(2) Muslim pop (%)	(3) Urban pop (%)	(4) Literate pop (%)	(5) % with electricity
Friday festival $t_{-1}$	0.012** (0.004)	-0.031 (0.024)	-0.021 (0.047)	-0.161 (0.196)	-0.243 (0.451)
Constant	0.031*** (0.002)	10.841*** (0.020)	17.744*** (0.038)	26.340*** (0.158)	7.251*** (0.363)
District FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Districts	339	339	339	339	339
Observations	19323	12204	12204	12204	12204

*Notes:* OLS regressions. [Driscoll and Kraay \(1998\)](#) standard errors robust against cross-sectional autocorrelation, serial autocorrelation up to lag 20, and heteroskedasticity are reported in parentheses.

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

cantly in terms of our covariates (both in terms of point estimate and p-value), but do differ substantively and statistically significantly in terms of riot prevalence. More in particular riots are 1.2 percent more likely to occur in years with at least one important Hindu festival occurring on a Friday. This is a 39 percentage points increase relative to the control group probability of 3.1 percent (as depicted by the constant in table [2](#)).<sup>15</sup>

## Results

### Main Results

Table [3](#) reports our baseline estimates. The OLS estimate in column (1) is in line with the hypothesis that ethnic riots are self-perpetuating and suggests that district-years that experience at least one Hindu-Muslim riot are 7.5 percentage points more likely to experience at least one more Hindu-Muslim riot in the next year. Given that the baseline probability of a riot in any particular district-year is only 5 percent this is a very strong correlation indeed.

Although district and 5-year time fixed effects may well control for some important factors

<sup>15</sup>Note that based on the same fixed effects regression as in table [2](#) the absolute number of Hindu-Muslim riots roughly doubles in years with a Hindu festival on Friday.

Table 3: Baseline results.

	(1) Baseline (OLS)	(2) First stage (2SLS)	(3) Second stage (2SLS)	(4) Reduced form (OLS)
Riot $t-1$	0.075*** (0.016)		-0.094 (0.201)	
Friday festival $t-1$		0.013*** (0.003)		-0.001 (0.003)
F statistic	—	15.0	—	—
District FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Districts	339	339	339	339
Observations	18984	18984	18984	18984

*Notes:* OLS and 2SLS regressions. Data is for Indian districts (1950–2006). The dependent variable in columns (1), (3), and (4) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the current year, and 0 otherwise. The dependent variable in column (2) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the previous year, and 0 otherwise. Friday festival  $t - 1$  is a dichotomous variable that takes the value 1 if at least one major Hindu festival fell on a Friday in the previous year, and 0 otherwise. The reported F statistic is the Kleibergen-Paap F statistic for the excluded instrument. Standard errors are reported in parentheses. Standard errors and F statistic are robust against cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the [Driscoll and Kraay \(1998\)](#) estimator.

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

that cause ethnic violence in both the current and the previous year, it would be premature to interpret these fixed effects estimates as reflecting a causal effect. This is because many time-varying unobservables are potentially correlated with both past and current ethnic violence, hence biasing unit fixed effects estimates through confounding. Examples of such time-varying confounders could be: government policies, election results, economic shocks, crime events, the radicalization of civil society leaders, among many others. Similarly, 5-year time fixed effects only account for factors that affect all districts at the same point in time, while we know from the existing literature that the causes of Hindu-Muslim riots in India are often highly place- and time-specific ([Varshney, 2002](#)).

Our current understanding of the causes of ethnic violence is arguably very far from sufficient to know for sure what all the relevant confounding factors are, how to measure them sufficiently precisely, and how to control for them in an unbiased fashion. For this reason we follow [Iyer and Shrivastava \(2018\)](#) by instrumenting Hindu-Muslim riots in the

past year with whether a major Hindu festival happened to fall on a Friday in that year. This plausibly controls for all potentially confounding factors because the date of Hindu festivals is exogenously determined by the lunar calendar, and because it is arguably unlikely that Hindu festivals on Fridays (as opposed to any other day of the week) are associated with future riots through other channels.

Column (2) reports the first stage results. The Kleibergen and Paap (2006) F statistic is 15, which substantially exceeds Staiger and Stock's (1997) commonly used threshold of 10.<sup>16</sup> We interpret this result as indicating that the instrument is strong and thus that asymptotic bias due to weak instrumentation is of little concern in our case.

In column (3) we report the second-stage 2SLS results. We find that the effect of current riots on future riots becomes *negative* (and statistically insignificant) when we account for endogeneity through instrumentation. This suggests that Hindu-Muslim riots do *not* tend to be self-perpetuating.

While 2SLS estimates are important to consider we in this case prefer to focus on the reduced form estimates reported in column (4) (i.e., the effect of a Hindu festival on a Friday in the previous year on Hindu-Muslim riots in the current year). Reduced form estimates are in our view preferable in this case because they do not rely on the assumption that Friday festival-induced riots are equally likely to go unreported as other riots, and because they do not rely on a correctly specified functional form in the first stage. Similarly as the 2SLS results we find the effect of a Friday festival on riots in the next year to be null, or if anything very slightly negative, when focusing on the reduced form.

Taken together the results from table 3 suggest that Hindu-Muslim riots in India do *not* in themselves tend to cause more riots in the future. Rather, riots are caused by other factors

---

<sup>16</sup>Note that we report significantly weaker first-stage F statistics than Iyer and Shrivastava (2018) do. This is because we take into account that the instrument is correlated within and across states at the same point in time, and hence that districts-years cannot be seen as entirely independent of each other.

which are not alleviated through riots and therefore also tend to cause future violence. Unit and time fixed effects do not control for all these confounding factors indicating that at least some of these unobserved factors are time-variant and/or do not affect all districts at the same point in time.

## Robustness

We examine the robustness of our baseline results across two dimensions.

First, we check whether our results are robust to: (1) using alternative measures of Hindu-Muslim riots; (2) controlling for the time-varying covariates from table 2;<sup>17</sup> (3) excluding particularly riot-prone district-years; (4) including state-specific time trends; and (5) assessing the effect on the state (rather than district) level. To save space we focus on the reduced form effect and report all results in one single table. As can be seen in table 4 our results are highly robust across (1)–(5).


Second, we examine the robustness of our results to taking deeper lag lengths, thereby taking into account that the effect of our instrument may operate over a longer time horizon. To simplify our analysis we examine the reduced form effect of Friday festivals at lag 2 to 20. As can be seen in figure 1 taking longer time lags makes no substantive difference to our conclusion—i.e., the reduced form effect remains in all cases small and statistically indistinguishable from zero. Taking the average of all 20 points estimates of lag 1 to 20 leads to a point estimate of -0.00128607, with 9 being slightly positive and 11 being slightly negative. This thus again suggests that the effect of riots over time is simply indistinguishable from zero.

---

<sup>17</sup>We include these covariates at  $t - 1$ . All results hold when (also) including them at  $t - 2$ .



Figure 1: Robustness to different lag lengths.



*Note:* Plot of point estimates and confidence intervals derived from 20 separate OLS regressions. District and 5-year time fixed effects are included in all regressions. The dependent variable is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot at time  $t$ , and 0 otherwise. The independent variable is a dichotomous variable that takes the value 1 if at least one major Hindu festival fell on a Friday  $\#$  years before  $t$ , and 0 otherwise. Lines refer to a 95% confidence interval. Confidence intervals are robust again cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.

Table 4: Robustness to alternative measurement and specification.

	(1) Baseline estimate	(2) Time-varying Controls	(3) Excluding No. of Riots > 2	(4) Time Trend	(5) State-level Panel
<b>Panel A:</b> Dependent variable is Riot at $t$					
Friday festival $t_{-1}$	-0.001 (0.005)	-0.006 (0.006)	-0.001 (0.004)	-0.000 (0.005)	-0.003 (0.005)
<b>Panel B:</b> Dependent variable is No. of Riots at $t$					
Friday festival $t_{-1}$	-0.002 (0.012)	-0.003 (0.024)	-0.000 (0.005)	0.004 (0.017)	-0.001 (0.020)
<b>Panel C:</b> Dependent variable is Riot within 100 km at $t$					
Friday festival $t_{-1}$	-0.018 (0.019)	-0.020 (0.023)	-0.018 (0.018)	-0.019 (0.020)	-0.031 (0.019)
Unit FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Districts	339	339	339	339	—
Observations	18984	11865	18857	18984	896

*Notes:* OLS regressions of reduced form effect. Standard errors are reported in parentheses. Standard errors are robust against cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## Generalizability of the Local Average Treatment Effect

Our 2SLS estimates capture the local average treatment effect of past Hindu-Muslim riots on current Hindu-Muslim riots. This parameter is estimated exploiting variation from those district-years which treatment status has changed due to the Friday festival instrument (Imbens and Angrist, 1994). This opens up the possibility that our instrumental variable estimates differ so starkly from our “naive” OLS estimates because Hindu-Muslim riots that are due to the coincidence of a major Hindu festival falling on a Friday are simply different from other Hindu-Muslim riots (Dunning, 2012).<sup>18</sup>

Although this problem is a result of our design, so that we cannot rule it out per definition, we find it unlikely that in our case the local average treatment effect differs significantly from the (true/unbiased) average treatment effect. First, we find that riots that occur in a year

<sup>18</sup>Here we mean “different” in a way that is relevant for riots being self-perpetuating, or not.

with or without a Hindu festival on Friday do not differ in terms of riot deaths and duration (Appendix C). Riots that do and do not co-occur with a major Hindu festival falling on a Friday thus do not appear to differ in terms of their intensity. Second, Iyer and Shrivastava's (2018) results are based on the same design and they find a large effect of Hindu-Muslim riots, induced by the coincidence of a major Hindu festival falling on Friday, on voting for the Hindu nationalist BJP. The “type” of Hindu-Muslim riots our instrumental variable estimates rely on are thus certainly not inconsequential in other important political-economic domains. Last, we have read the notes of all Hindu-Muslim riots in Varshney and Wilkinson's (2006) original dataset and have found no clear difference between riots that appeared to have been the result of public rituals/festivities, and Hindu-Muslim riots that were triggered by other events.<sup>19</sup>

Having said the above we can naturally not know for sure that Hindu-Muslim riots that are triggered by a Friday Hindu festival do not differ on relevant unobserved variables from riots with other causes. What we ultimately can be more sure off is that at least the type of riots that are Friday Hindu festival induced—which are an important and politically consequential subset of Hindu-Muslim riots (Iyer and Shrivastava, 2018)—do not appear to be self-perpetuating.

---

<sup>19</sup>Note that Varshney and Wilkinson (2006) specifically code when a riot appeared to be the result of “public rituals/festivities.” Examples of other triggers that Varshney and Wilkinson (2006) record are: “animal slaughter (cow),” “desecration of religious site/symbol,” and “economic interest (land).” As expected, riots in a district-year with a Hindu festival on a Friday are 1.3 times more likely (P-value: 0.000) to be coded by Varshney and Wilkinson (2006) as being caused by public rituals/festivities. Note that the correlation with Varshney and Wilkinson's (2006) public rituals/festivities coding is not perfect as riots sometimes result from public rituals/festivities that don't fall on a Friday (Allie, 2024).

## Measurement Error

Clearly any data on Hindu-Muslim riots in virtually the whole of India over 56 years does contain at least some measurement error. But can our results be explained away by measurement error?

We believe that it is highly unlikely that measurement error is driving our results. First, Varshney (2002) has shown, using a number of robustness checks, that the Times Mumbai does not tend to underreport deathly riots in parts of India further away from Mumbai. Most importantly, Varshney (2002) shows that the ratio of small (one to three deaths) to major riots (four deaths or more) is not correlated with an individual states' distance to Mumbai. Assuming that major riots are highly likely to be picked up by any national newspaper, and assuming that the ratio between major and small riots does not differ significantly across states, this indicates that the Mumbai Times does not underreport riots further away from its head office. Second, we find the results unchanged when using only district-years in the state of Maharashtra (the state of Mumbai), when dropping all district-years in states not directly bordering Maharashtra, and when excluding all riots with less than four reported deaths. Third, systematic measurement error in the form of the Mumbai Times being less informed or less interested in riots further away from Mumbai is likely to be (relatively) time invariant and thus controlled for by district fixed effects. Last, any type of measurement error is in any case unlikely to be systematically associated with whether Hindu festivals fall on a Friday or not (which is fully determined by the lunar calendar).

## Why Are Ethnic Riots Not Self-Perpetuating? Suggestive Evidence from Treatment Heterogeneity

Our results naturally raise the question *why* Hindu-Muslim riots in India are on average not self-perpetuating. Answering this question beyond a reasonable doubt is beyond the scope of this article. Here we do, however, shortly address one potential hypothesis.

The seminal work of Wilkinson (2004) and Brass (1996, 2011) suggests that where and

when Hindu-Muslim violence breaks out in India is often the result of strategic organization by politicians/political parties that benefit electorally from ethnic riots. Based on this prior literature we hypothesize that Hindu-Muslim riots are more likely to become self-perpetuating when *those in power have an interest in such a development*.

To test this hypothesis we study whether our instrumental variable estimates change depending on whether the BJP are in power at  $t_0$  on the state-level.<sup>20</sup> The logic is as follows: policing in India is mostly organized on the state level, so that state-level politicians have significant control over riot control measures (Nellis, Weaver and Rosenzweig, 2016) and the BJP tends to tailor to a Hindu nationalist constituency and tends to electorally benefit from Hindu-Muslim riots (Iyer and Shrivastava, 2018). The BJP therefore arguably has little incentive to lessen ethnic tensions, and may in fact actively benefit from politicizing Hindu-Muslim cleavages.<sup>21</sup>


In figure 2 we report the reduced form results when interacting the Friday Hindu festival instrument with a dummy that takes the value 1 if the Chief Minister of a district-year's state is a member of the BJP, and 0 otherwise. As can be seen, and in line with the hypothesis laid out above, we do find a statistically and economically significant positive effect of our instrument among district-year's which are ruled by a BJP state Chief Minister.

---

<sup>20</sup>This analysis thus estimates whether the effect of a Hindu festival falling on Friday in the previous year on Hindu-Muslim riots in the current year differs depending on whether the BJP is in power in the current year, or not (but regardless of whether the BJP was in power in the previous year). Appendix D shows that the results remain substantively the same when requiring the BJP to be in power in both years.

<sup>21</sup>Although this is hard to proof beyond a reasonable doubt it is regularly alleged that members of the BJP actively instigate Hindu-Muslim riots and that BJP politicians regularly do less that they could to stop Hindu-Muslim riots after they have started (see, for example, the BBC documentary "India: The Modi Question," on the role of Narendra Modi in the 2002 Gujarat riots).

Figure 2: Treatment heterogeneity with BJP in state executive power.



*Note:* Plot of how the reduced form effect varies with whether the Chief Minister of a district-year' state is a member of the BJP, or not. Dots refer to point estimates. Lines refer to a 95% confidence interval. Confidence intervals are robust to cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the [Driscoll and Kraay \(1998\)](#) estimator.

We interpret this as suggestive evidence that Hindu-Muslim riots in India *can* be self-perpetuating, but only if politicians that hold state executive power have an interest in it being so. We are however hesitant in interpreting the estimates in figure 2 as causal because whether the BJP is in state executive power, or not, is likely to be endogenous.<sup>22</sup>

Combining the results in figure 2 with the findings of Iyer and Shrivastava (2018) does, however, suggest the possibility of a troubling feedback loop: Hindu-Muslim riots may make it more likely that the BJP wins elections and holds power, and once the BJP is in power, Hindu-Muslim riots are more likely to reoccur, possibly further improving the BJP’s electoral fortunes.

## Conclusion

We are, to the best of our knowledge, the first to use a credible quasi-experimental design to test whether ethnic riots are self-perpetuating. Focusing on the substantively important case of Hindu-Muslim riots in India we find no evidence that ethnic riots are self-perpetuating, lending support to the idea that the widely documented correlation between ethnic riots over time is spurious.

Our findings raise several important questions for future research. First, how well do our results travel to other countries, other times, and other types of conflict? The major strength of our paper is that it is able to recover a plausible estimate of the causal effect of ethnic riots on future riots in a substantively important case over a relatively long period of time. Clearly, however, the self-perpetuating riots hypothesis is much more general than

---

<sup>22</sup>We are unable to address this endogeneity problem with a standard close election regression discontinuity design because of a lack of statistical power (note that the BJP only came into widespread power at the end of the 1990s and that such an analysis would be “double local” in the sense that it would estimate the local average treatment effect of the instrument among that small subset of cases that experienced a very close election involving the BJP).

the specific case of Hindu-Muslim riots in India alone. Future research is therefore required to shed light on to what extent our findings are generalizable beyond Hindu-Muslim riots in India over the 1950 to 2006 period. Equally important, Hindu-Muslim riots in India represent only one type of ethnic violence. It remains unclear whether the absence of self-perpetuation we observe in the case of Hindu-Muslim riots in India also applies to other forms of ethnic violence, such as pogroms, ethnic warfare, or genocides. Advancing quasi-experimental or otherwise well-identified research designs in these other domains is an important area for future research.

Second, our preliminary evidence suggests that Hindu-Muslim riots are more likely to become self-perpetuating when the Bharatiya Janata Party (BJP) holds state executive power. This conditional finding points to the need for further investigation into the political and institutional factors that shape the self-reinforcing dynamics of ethnic riots. Under what conditions, and through which mechanisms, do ethnic riots become more or less likely to generate subsequent riots? Understanding these contingencies may help explain both spatial and temporal variation in riot recurrence.

Finally, the observed interaction with BJP rule aligns with the broader insights of [Wilkinson \(2004\)](#) and [Nellis, Weaver and Rosenzweig \(2016\)](#), who argue that ethnic riots in India are closely tied to elite political incentives. In particular, their work suggests that when political leaders face strong electoral incentives to prevent ethnic violence riots are less likely to (re)occur. Our findings reinforce this view and underscore the need for further research on how political incentives can be structured to prevent the escalation of ethnic violence. If the recurrence of ethnic riots is not inevitable, but conditional on the strategic behavior of political actors, then designing institutions that align elite incentives with peace becomes all the more urgent.



## References

- Abbs, Luke. 2021. “The Language of the Unheard? Ethno-Political Exclusion and Ethnic Riots in Africa.” *Journal of Global Security Studies* 6(2):1–19.
- Aidt, Toke, Gabriel Leon-Ablan and Max Satchell. 2022. “The Social Dynamics of Collective Action: Evidence from the Diffusion of the Swing Riots, 1830–1831.” *Journal of Politics* 84(1):209–225.
- Aidt, Toke S. and Gabriel Leon. 2016. “The Democratic Window of Opportunity: Evidence from Riots in Sub-Saharan Africa.” *Journal of Conflict Resolution* 60(4):694–717.
- Allie, Feyaad. 2024. “Sacred Time and Religious Violence: Evidence from Hindu-Muslim Riots in India.” *Journal of Conflict Resolution* 68(10):1968–1993.
- Bateson, Regina. 2012. “Crime Victimization and Political Participation.” *American Political Science Review* 106(3):570–587.
- Beber, Bernd, Philip Roessler and Alexandra Scacco. 2014. “Intergroup Violence and Political Attitudes: Evidence from a Dividing Sudan.” *Journal of Politics* 76(3):649–665.
- Bhavnani, Ravi, Michael G Findley and James H Kuklinski. 2009. “Rumor Dynamics in Ethnic Violence.” *Journal of Politics* 71(3):876–892.
- Bhavnani, Rikhil R and Bethany Lacina. 2015. “The Effects of Weather-Induced Migration on Sons of the Soil Riots in India.” *World Politics* 67(4):760–794.
- Blattman, Christopher. 2009. “From Violence to Voting: War and Political Participation in Uganda.” *American Political Science Review* 103(2):231–247.
- Bohlken, Anjali T. and Ernest J. Sergenti. 2010. “Economic Growth and Ethnic Violence: An Empirical Investigation of Hindu-Muslim Riots in India.” *Journal of Peace Research* 47(5):589–600.
- Brass, Paul R. 2011. *The Production of Hindu-Muslim Violence in Contemporary India*. University of Washington Press.
- Brass, Paul R (ed.). 1996. *Riots and Pogroms*. Macmillan Press Ltd.
- Bulutgil, H. Zeynep and Neeraj Prasad. 2023. “Inequality, Elections, and Communal Riots

- in India.” *Journal of Peace Research* 60(4):619–633.
- Collier, Paul. 2008. *The Bottom Billion: Why the Poorest Countries are Failing and What Can Be Done About It*. Oxford University Press.
- Collier, Paul and Anke Hoeffler. 1998. “On Economic Causes of Civil War.” *Oxford Economic Papers* 50(4):563–573.
- Collier, Paul and Anke Hoeffler. 2002. “On the Incidence of Civil War in Africa.” *Journal of Conflict Resolution* 46(1):13–28.
- Collier, Paul and Nicholas Sambanis. 2002. “Understanding Civil War: A New Agenda.” *Journal of Conflict Resolution* 46(1):3–12.
- Collier, Paul, V.L. Elliot, Håvard Hegre, Anke Hoeffler, Mara Reynal-Querol and Nicholas Sambanis. 2003. *Breaking the Conflict Trap: Civil War and Development Policy*. Oxford University Press and World Bank Publications.
- De Mesquita, Ethan Bueno et al. 2008. “Terrorist Factions.” *Quarterly Journal of Political Science* 3(4):399–418.
- Dercon, Stefan and Roxana Gutiérrez-Romero. 2012. “Triggers and Characteristics of the 2007 Kenyan Electoral Violence.” *World Development* 40(4):731–744.
- Driscoll, John C. and Aart C. Kraay. 1998. “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data.” *Review of Economics and Statistics* 80(4):549–560.
- Dunning, Thad. 2012. *Natural Experiments in the Social Sciences: A Design-Based Approach*. Cambridge University Press.
- Eubank, Nicholas et al. 2019. “Social Networks and the Political Salience of Ethnicity.” *Quarterly Journal of Political Science* 14(1):1–39.
- Fearon, James D and David D Laitin. 2003. “Ethnicity, Insurgency, and Civil War.” *American Political Science Review* 97(1):75–90.
- Gould, Eric D. and Esteban F. Klor. 2010. “Does Terrorism Work?” *Quarterly Journal of Economics* 125(4):1459–1510.
- Hager, Anselm, Krzysztof Krakowski and Max Schaub. 2019. “Ethnic Riots and Prosocial

- Behavior: Evidence from Kyrgyzstan.” *American Political Science Review* 113(4):1029–1044.
- Hartman, Alexandra C. and Benjamin S. Morse. 2018. “Violence, Empathy and Altruism: Evidence from the Ivorian Refugee Crisis in Liberia.” *British Journal of Political Science* pp. 1–25.
- Hazlett, Chad. 2019. “Angry or Weary? How Violence Impacts Attitudes toward Peace among Darfurian Refugees.” *Journal of Conflict Resolution* 64(5):1–27.
- Hegre, Håvard and Nicholas Sambanis. 2006. “Sensitivity Analysis of Empirical Results on Civil War Onset.” *Journal of Conflict Resolution* 50(4):508–535.
- Horowitz, Donald L. 1985. *Ethnic Groups in Conflict*. University of California Press.
- Horowitz, Donald L. 2001. *The Deadly Ethnic Riot*. University of California Press.
- Imbens, Guido W. and Joshua D. Angrist. 1994. “Identification and Estimation of Local Average Treatment Effects.” *Econometrica* 62(2):467–475.
- Iyer, Sriya and Anand Shrivastava. 2018. “Religious Riots and Electoral Politics in India.” *Journal of Development Economics* 131:104–122.
- Jaffrelot, Christophe. 1999. *The Hindu Nationalist Movement and Indian Politics: 1925 to the 1990s*. Penguin Books India.
- Jha, Saumitra. 2013. “Trade, Institutions, and Ethnic Tolerance: Evidence from South Asia.” *American Political Science Review* 107(4):806–832.
- Kasara, Kimuli. 2017. “Does Local Ethnic Segregation Lead to Violence?: Evidence from Kenya.” *Quarterly Journal of Political Science* 11(4):441–470.
- Kleibergen, Frank and Richard Paap. 2006. “Generalized Reduced Rank Tests Using the Singular Value Decomposition.” *Journal of Econometrics* 133(1):97–126.
- Lyall, Jason. 2009. “Does Indiscriminate Violence Incite Insurgent Attacks? Evidence from Chechnya.” *Journal of Conflict Resolution* 53(3):331–362.
- Miguel, Edward, Shanker Satyanath and Ernest Sergenti. 2004. “Economic Shocks and Civil Conflict: An Instrumental Variables Approach.” *Journal of Political Economy* 112(4):725–

- Mitra, Anirban and Debraj Ray. 2014. "Implications of an Economic Theory of Conflict: Hindu-Muslim Violence in India." *Journal of Political Economy* 122(4):719–765.
- Nellis, Gareth, Michael Weaver and Steven Rosenzweig. 2016. "Do Parties Matter for Ethnic Violence? Evidence From India." *Quarterly Journal of Political Science* 11(3):249–277.
- Nellis, Gareth and Nilouffer Siddiqui. 2018. "Secular Party Rule and Religious Violence in Pakistan." *American Political Science Review* 112(1):49–67.
- Nisbett, Richard E. 2018. *Culture of Honor: The Psychology of Violence in the South*. Routledge.
- Niwa, Erika Y., Paul Boxer, Eric Dubow, L. Rowell Huesmann, Khalil Shikaki, Simha Landau and Shira D. Gvirsman. 2016. "Growing Up Amid Ethno-Political Conflict: Aggression and Emotional Desensitization Promote Hostility to Ethnic Outgroups." *Child Development* 87(5):1479–1492.
- Paler, Laura, Leslie Marshall, Sami Atallah et al. 2020. "How Cross-Cutting Discussion Shapes Support for Ethnic Politics: Evidence from an Experiment in Lebanon." *Quarterly Journal of Political Science* 15(1):33–71.
- Schutte, Sebastian, Constantin Ruhe and Niranjana Sahoo and. 2023. "How Fear of Violence Drives Intergroup Conflict: Evidence from a Panel Survey in India." *Terrorism and Political Violence* 35(2):229–247.
- Staiger, Douglas and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65(3):557–586.
- Tarabah, Asma, Lina Kurdahi Badr, Jinan Usta and John Doyle. 2016. "Exposure to Violence and Children's Desensitization Attitudes in Lebanon." *Journal of Interpersonal Violence* 31(18):3017–3038.
- The Times of India. 2018. "Communal Clash After Youth Throw Colour on Minority Men; 10 Hurt." March 4, 2018. Available at: <https://timesofindia.indiatimes.com/city/meerut/communal-clash-after-youths-throw-colour-on-minority-men-10-hurt/>

---

[articleshow/63150011.cms](#) (last accessed: February 5, 2023).

- Varshney, Ashutosh. 2001. "Ethnic Conflict and Civil Society: India and Beyond." *World Politics* 53(3):362–398.
- Varshney, Ashutosh. 2002. *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. Yale University Press.
- Varshney, Ashutosh and Steven Wilkinson. 2006. "Varshney-Wilkinson Dataset on Hindu-Muslim Violence in India, 1950-1995, Version 2." Ann Arbor, MI: Inter-University Consortium for Political and Social Research.
- Wilkinson, Steven I. 2004. *Votes and Violence: Electoral Competition and Ethnic Riots in India*. Duke University Press.
- Wilkinson, Steven I. 2009. "Riots." *Annual Review of Political Science* 12(1):329–343.

— Online Appendix —

# Is Ethnic Violence Self-Perpetuating? Quasi-Experimental Evidence from Hindu-Muslim Riots in India

By: Sam van Noort & Tanushree Goyal

## Contents

<b>A 6-month level analysis</b>	<b>1</b>
<b>B Time fixed effects for exceptionally riot-prone years</b>	<b>3</b>
<b>C Intensity of riots in years with and without a Hindu festival on Friday</b>	<b>5</b>
<b>D Alternative specification BJP interaction</b>	<b>7</b>

## A 6-month level analysis

In the main text we analyze our data on the year-to-year level. It is possible, however, that while Hindu-Muslim riots in one year do not lead to a higher probability of Hindu-Muslim riots in the next year, or the years thereafter (i.e., figure 1 in the main text), that the relationship does exist on a shorter time scale.

To consider this possibility, we analyze the effect on the six-month level. While it is theoretically possible to examine effects at increasingly granular time scales—down to months, weeks, days, hours, or even seconds—we focus on the six-month level for two practical reasons.

First, Hindu-Muslim riots in India are a relatively rare occurrence, as 225 out of 338 districts in our data set have experienced at least one riot between 1950 and 2006, but only 45 districts have ever experienced riots in consecutive calendar months, which includes instances of a single riot spanning more than one calendar month (see next point). A second, more fundamental issue is that many Hindu-Muslim riots in our data set occur over the course of several days, weeks, or even months (the maximum length of a single discrete riot in our dataset is 76 days, taking place over three distinct calendar months).<sup>1</sup> This means that when collapsing the data to 1 or even 3-month level one may in fact code the same riot in both periods and thus regress the same riot on itself (finding the autocorrelative pattern hypothesized by the self-perpetuating riots hypothesis by definition). The six-month level is in our view a reasonable compromise between a more disaggregated analysis and the aforementioned concerns.

In table A1 we report the results of running our main specifications on the six month level. As can be seen, this leaves the null result essentially unchanged.

---

<sup>1</sup>These are single riots in the sense that they take place in the same area, with a singular trigger, over a single issue, but play out over a longer period of time.

Table A1: Results when using 6 months lags (as opposed to annual lags).

	(1) Baseline (OLS)	(2) First stage (2SLS)	(3) Second stage (2SLS)	(4) Reduced form (OLS)
Riot $t-1$	0.347*** (0.068)		-0.210 (0.497)	
Friday festival $_{t-1}$		0.005** (0.002)		-0.001 (0.003)
F statistic	–	7.1	–	–
District FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Districts	338	338	338	338
Observations	38194	38194	38194	38194

*Notes:* OLS and 2SLS regressions. Data is for Indian districts (1950–2006). The dependent variable in columns (1), (3), and (4) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the current year, and 0 otherwise. The dependent variable in column (2) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the previous year, and 0 otherwise. Friday festival  $t - 1$  is a dichotomous variable that takes the value 1 if at least one major Hindu festival fell on a Friday in the previous year, and 0 otherwise. The reported F statistic is the Kleibergen-Paap F statistic for the excluded instrument. Standard errors are reported in parentheses. Standard errors and F statistic are robust against cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.


\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$



## B Time fixed effects for exceptionally riot-prone years

In the main text we include 5-year time fixed effects in all specifications. These account by design for all factors that affect all districts in the same 5-year period. Here we take a different approach to the issue of common shocks by defining the time fixed effects based on the years with the most Hindu-Muslim riots nationwide. In figure 1 we plot the total number of Hindu-Muslim riots in each year across India from 1950 until 2006.

Figure 1: Number of riots per year (in India as a whole).



To ensure that our results are not driven by a few years with unusual amounts of riots we generate a dummy that takes the value 1 if more than 50 Hindu-Muslim riots occurred in that year, and 0 otherwise.<sup>2</sup> Table A2 shows the results when using this dummy as the time fixed effect. As can be seen we find the same null effect in this alternative specification.

<sup>2</sup>The years coded as 1 are: 1986, 1987, 1990, 1992, 2002, 2003, 2004, 2005, and 2006.

Table A2: Results with fixed effects for each year with more than 50 riots nationwide.

	(1) Baseline (OLS)	(2) First stage (2SLS)	(3) Second stage (2SLS)	(4) Reduced form (OLS)
Riot $t-1$	0.079*** (0.015)		0.007 (0.134)	
Friday festival $_{t-1}$		0.015*** (0.004)		0.000 (0.002)
F statistic	–	13.8	–	–
District FE	Yes	Yes	Yes	Yes
Riot-prone year FE	Yes	Yes	Yes	Yes
Districts	339	339	339	339
Observations	18984	18984	18984	18984

*Notes:* OLS and 2SLS regressions. Data is for Indian districts (1950–2006). The dependent variable in columns (1), (3), and (4) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the current year, and 0 otherwise. The dependent variable in column (2) is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the previous year, and 0 otherwise. Friday festival  $t - 1$  is a dichotomous variable that takes the value 1 if at least one major Hindu festival fell on a Friday in the previous year, and 0 otherwise. The reported F statistic is the Kleibergen-Paap F statistic for the excluded instrument. Standard errors are reported in parentheses. Standard errors and F statistic are robust against cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## C Intensity of riots in years with and without a Hindu festival on Friday

In this section we examine whether Hindu-Muslim riots that occur in district-years with and without a Hindu festival on Friday differ in terms of their intensity. To do so we regress different measures of riot deaths and riot duration on the interaction between our riot dummy (whether a Hindu-Muslim riot occurred, or not) and our Friday festival dummy (whether a Hindu festival fell on a Friday, or not), in the same district-year. This regression captures whether district-years with a riot and a Hindu festival on Friday had more/less riot deaths and longer/shorter riots, as compared to district-years that also had a riot but without a Hindu festival falling on a Friday.

Table [A3](#) reports the results. As can be seen by looking at the interaction term, we find that Hindu-Muslim riots in district-years with and without a Hindu festival on Friday do not differ economically or statistically significantly in terms of riot deaths and duration. This is true regardless of whether one looks at the absolute number of riot deaths and riot days, the natural log of riot deaths and riot days (+1), and whether one looks at two dichotomous variables that take the value 1 if a district-year had at least one riot death and at least one riot lasting longer than 1 day, and 0 otherwise.

Table A3: Differences in riot deaths and duration in years with and without a Hindu festival on Friday.

	<b>Riot deaths</b>			<b>Riot duration</b>		
	(1) Absolute	(2) Natural log	(3) Dichotomous	(4) Absolute	(5) Natural log	(6) Dichotomous
Riot	9.606*** (2.722)	0.999*** (0.106)	0.229*** (0.044)	3.393*** (0.741)	1.097*** (0.100)	0.302*** (0.071)
Friday festival	0.014 (0.104)	-0.007 (0.006)	-0.002 (0.001)	0.002 (0.018)	0.001 (0.003)	0.001 (0.001)
Riot *	-3.325 (2.296)	-0.018 (2.296)	0.023 (0.038)	-0.779 (0.622)	-0.023 (0.079)	0.034 (0.042)
Friday festival						
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Districts	339	339	339	339	339	339
Observations	8874	8874	8439	8475	8475	8475

*Notes:* OLS regressions. Data is for Indian districts (1950–2006). Dependent variables are described in the column headings, and are measured in the current year. Riot is a dichotomous variable that takes the value 1 if a district has experienced at least one Hindu-Muslim riot in the current year, and 0 otherwise. Friday festival is a dichotomous variable that takes the value 1 if at least one major Hindu festival fell on a Friday in the current year, and 0 otherwise. Standard errors are reported in parentheses. Standard errors are robust against cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.


\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

## D Alternative specification BJP interaction

In the main text we analyze treatment heterogeneity by interacting whether a Hindu festival fell on a Friday in the previous year with whether the BJP holds state executive power in the current year.

Figure 2 re-runs this analysis but now requiring that the BJP holds state executive power in both the current and the previous year. As can be seen, this leaves the results substantively unchanged.

Figure 2: Treatment heterogeneity with BJP in state executive power (in both current and previous year).



*Note:* Plot of how the reduced form effect varies with whether the Chief Minister of a district-year' state is a member of the BJP, or not. Dots refer to point estimates. Lines refer to a 95% confidence interval. Confidence intervals are robust to cross-sectional autocorrelation, serial autocorrelation, and heteroskedasticity, using the Driscoll and Kraay (1998) estimator.

## References

Driscoll, John C. and Aart C. Kraay. 1998. “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data.” *Review of Economics and Statistics* 80(4):549–560.