DISCUSSION PAPER SERIES

DP21638

THE ECONOMIC COSTS OF RELIGIOUS RIOTS IN INDIA

Sriya Iyer, Dmitrii Petrukhin and Anand Shrivastava

DEVELOPMENT ECONOMICS AND MACROECONOMICS AND GROWTH



ISSN 0265-8003

THE ECONOMIC COSTS OF RELIGIOUS RIOTS IN INDIA

Sriya Iyer, Dmitrii Petrukhin and Anand Shrivastava

Discussion Paper DP21638
Published 18 June 2026
Submitted 17 June 2026

Centre for Economic Policy Research
187 boulevard Saint-Germain, 75007 Paris, France
2 Coldbath Square, London EC1R 5HL
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Development Economics
- Macroeconomics and Growth

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Sriya Iyer, Dmitrii Petrukhin and Anand Shrivastava

THE ECONOMIC COSTS OF RELIGIOUS RIOTS IN INDIA

Abstract

This paper estimates the economic effects of communal riots in India using a new district–year panel for 1994–2013 that combines extended riot data with nightlights, manufacturing outcomes, infrastructure, and district-level credit. To address endogeneity, we use quasi-random variation generated when major Hindu festivals fall on a Friday (Fridays draw large Muslim congregations for weekly prayers) and construct a spatially weighted instrument for local riot risk. Instrumental variable estimates show that higher riot exposure leads to sizable and persistent economic losses. A one-unit increase in distance-weighted exposure reduces nightlights by about 0.105 log points, implying a GDP decline of roughly 0.45\%. We trace several channels, including contractions in firm output and employment, fewer bank branches, less agricultural credit, and a thinner market infrastructure. Using the timing of household surveys relative to riots, we also find that trust in state institutions declines. Together, the results quantify the economic costs of communal riots and identify the mechanisms through which social conflict depresses local economic activity.

JEL Classification: Z12

Keywords: Religion, India

Sriya Iyer - si105@cam.ac.uk
University pf Cambridge and CEPR

Dmitrii Petrukhin - dp657@cam.ac.uk
University of Cambridge

Anand Shrivastava - anand.shrivastava@apu.edu.in
Azim Premji University, India

Acknowledgements
We thank Massimo Morelli, Toke Aidt for helpful comments and discussions. We are grateful to Zihan Tian and Mehar Saxena for excellent research assistance. Iyer acknowledges funding support from the Templeton Religion Trust's SCORE programme and The Keynes Fund for Applied Economics. Any remaining errors are our own. The content and opinions expressed are that of the authors, and do not necessarily reflect the views of their institutions.

The Economic Costs of Religious Riots in India

Sriya Iyer*
University of Cambridge
si105@cam.ac.uk

Dmitrii Petrukhin
University of Cambridge
dp657@cam.ac.uk

Anand Shrivastava
Azim Premji University
anand.shrivastava@apu.edu.in

June 2026

Abstract

This paper estimates the economic effects of communal riots in India using a new district–year panel for 1994–2013 that combines extended riot data with nightlights, manufacturing outcomes, infrastructure, and district-level credit. To address endogeneity, we use quasi-random variation generated when major Hindu festivals fall on a Friday (Fridays draw large Muslim congregations for weekly prayers) and construct a spatially weighted instrument for local riot risk. Instrumental variable estimates show that higher riot exposure leads to sizable and persistent economic losses. A one-unit increase in distance-weighted exposure reduces nightlights by about 0.105 log points, implying a GDP decline of roughly 0.45%. We trace several channels, including contractions in firm output and employment, fewer bank branches, less agricultural credit, and a thinner market infrastructure. Using the timing of household surveys relative to riots, we also find that trust in state institutions declines. Together, the results quantify the economic costs of communal riots and identify the mechanisms through which social conflict depresses local economic activity.

1 Introduction

Following the economic liberalisation reforms of the early 1990s, India entered a period of sustained and rapid economic expansion. India's GDP grew at an average annual rate of approximately 6–7 percent in the 1990s and 2000s, lifting hundreds of millions out of poverty and establishing the country as one of the fastest-growing large economies in the world. Yet this growth coexisted with a form of social violence that is distinctive in the global landscape of conflict: recurrent, localised episodes of communal rioting between the Hindu majority and the Muslim minority. Unlike the sustained insurgencies that have ravaged countries such as Sri Lanka, Sudan, or Colombia, and which have attracted the bulk of the economics-of-conflict literature, Hindu–Muslim riots are episodic, short-lived, and occur within a functioning democracy. Their persistence through decades of rapid growth raises a fundamental question: what is the economic cost of violence that is neither a civil war nor random crime, but a recurrent feature of local social life?

Communal violence remained a persistent feature of India's social and economic landscape throughout this high-growth period and remains one of the most important forms of local conflict in the country. For example, during the 2013 Muzaffarnagar riots in western Uttar Pradesh, curfews and insecurity brought local industrial activity to a halt, with contemporaneous reports documenting factory closures, labour shortages, and large daily production losses (Business

*We thank Massimo Morelli, Toke Aidt for helpful comments and discussions. We are grateful to Zihan Tian and Mehar Saxena for excellent research assistance. Iyer acknowledges funding support from the Templeton Religion Trust's SCORE programme and The Keynes Fund for Applied Economics. Any remaining errors are our own. The content and opinions expressed are that of the authors, and do not necessarily reflect the views of their institutions.

Standard, 2013). Such episodes are not unique to India. Ethnic and religious violence has been a recurring feature of societies across the world: Fearon (2008) documents that more than 14 percent of all ethnic minority groups were involved in violent conflict between 1945 and 1998, and estimates of deaths from ethnic violence since World War II run as high as 10 million (Horowitz, 1985). Where such violence has been studied economically, the costs have been shown to be severe and wide-ranging: linked to lower incomes (Easterly and Levine, 1997), higher corruption (Mauro, 1995), and weaker government quality (La Porta et al., 1999).

Despite extensive work on the political causes and consequences of religious riots, much less is known about their economic impact. Do riots depress local economic activity, and through which mechanisms? This paper provides new answers to these questions using a novel district–year panel that combines religious riots, night-lights, manufacturing microdata, financial and market infrastructure, and district-level credit data.

We assemble the most complete time series of Hindu–Muslim riots to date, extending the Varshney and Wilkinson (2006) and Iyer and Shrivastava (2018) datasets through 2014 using digitised newspaper archives, although our analysis focuses on 1994–2013, the period for which the economic outcome data are available. To capture the spatial footprint of violence, we construct a Gaussian distance-decay measure that aggregates riot intensity within and around each district. Because economic activity in one district can be affected by riots in nearby districts (through the erosion of trust, for example), this measure better reflects the economic exposure to violence than conventional count or binary indicators.

A key challenge in estimating the causal impact of riots is that both riot incidence and local economic conditions may be driven by unobserved time-varying factors. We address this by building on the instrumental-variable strategy of Iyer and Shrivastava (2018), who show that riots are more likely when a major Hindu festival falls on a Friday, because these coincidences generate large religious gatherings in public spaces for weekly prayers. Using variation in the lunar calendar, we construct a spatially weighted analogue of their instrument that captures a district's exposure to Friday–festival coincidences occurring in and around nearby districts. This instrument delivers exogenous cross-district and time variation while allowing us to include the year fixed effects that are crucial for night-lights analysis (Henderson et al., 2011).[1]

We study the period 1994–2013 and find that riots have large, negative, and statistically significant effects on local economic activity. Instrumenting distance-weighted riot intensity with distance-weighted festival shocks, we find that a one-unit increase in riot exposure in year $t-1$ reduces district-year night-lights, the standard proxy for local GDP, by roughly 0.105 log points. Using the elasticity of log GDP with respect to log night-lights of 0.056 estimated by Bickenbach et al. (2016) in Indian districts, our preferred IV estimate implies a roughly 0.45% decline in local GDP per one-unit increase in riot exposure, after correcting for under-reporting.

To provide aggregate context, we conduct a back-of-the-envelope aggregation of implied district GDP losses using night-lights-based economic size weights. This exercise suggests that riot exposure is associated with annual national GDP losses ranging from about 0.06% in relatively calm years to about 0.77% in peak years, and we report the full year-by-year series in Appendix Table 20. For comparison, Abadie and Gardeazabal (2003) document a persistent GDP per capita shortfall of roughly 10% associated with sustained terrorist activity in the Basque Country. In our setting, the most extreme state–year loss, Gujarat in 2003, implies a GDP decline of approximately 8.87%, which is comparable in magnitude to Abadie and Gardeazabal (2003)'s findings. More generally, implied losses are modest in most state–years: the mean state–year GDP loss is 0.23%, reflecting the episodic, localised, and short-lived nature of communal riots.[2] These results are robust across bandwidth choices, alternative night-light transformations, fixed-satellite subsamples, Conley (1999) standard errors, winsorisation, different functional forms of the outcome variables, sample restrictions and alternative GDP data.

[1] Iyer and Shrivastava (2018) use this specification in their robustness analysis.

[2] Appendix D provides the full distribution and state–year breakdown.

We next examine the mechanisms underlying the decline in local economic activity following riots using ASI data from 2000–2007. Focusing on the one-year lag of riot exposure, we find significant short-run contractions among census plants: large manufacturing establishments that are surveyed every year rather than intermittently sampled. In particular, riots reduce labour and output growth, while the effects on capital are weaker and imprecisely estimated. Level effects are limited, with a decline in labour among census firms being the only robust result. We find no evidence of responses along the extensive margin: firm entry and exit are unaffected at the one-year horizon. Overall, the manufacturing evidence points to short-run disruption operating primarily through reduced labour input and slower production growth among formal manufacturing firms.

Beyond firms, riots trigger a deterioration in the local commercial and financial environment. Districts exposed to greater riot intensity experience sizeable contractions in market infrastructure, driven primarily by sharp declines in sub-markets, with more muted responses among principal markets. Banking presence also falls substantially: the number of bank branches declines by roughly 20 percent within a year of increased riot exposure. In parallel, agricultural credit contracts by around 10 percent, even as total credit remains broadly unchanged, indicating a targeted retrenchment from riskier lending rather than a wholesale withdrawal of finance.

In addition to these real economic channels, we examine whether riots affect confidence in state institutions. While intergroup trust is difficult to measure due to data constraints, institutional confidence plays a central role in economic coordination, enforcement, and market participation. Using individual-level data from the India Human Development Survey (IHDS), we use quasi-random variation in the timing of survey interviews relative to local riot events to study short-run changes in institutional trust. Comparing respondents interviewed just before and just after a riot within the same district and survey wave, we find significant declines in confidence in institutions directly responsible for maintaining order and governance, including politicians, the police, and local government bodies. By contrast, confidence in schools, health providers, and other placebo outcomes such as asset ownership does not respond, supporting the interpretation that riots generate a selective erosion of institutional confidence rather than a general shift in survey responses.

Taken together, the evidence is consistent with riots undermining the local environment for economic exchange. We observe not only disruptions in firm activity but also a contraction in the institutions that support commerce: banks and local markets.

The findings carry broader policy relevance beyond India. Many contemporary conflicts share a common feature with the setting we study: governments are under pressure to protect civilian populations from violence, and their perceived capacity or willingness to do so is openly contested. Our results suggest that even localised, episodic violence, far short of full-scale war, can erode confidence in the state institutions responsible for maintaining order and governance. The quasi-random timing of Hindu–Muslim riots relative to household survey interviews allows us to identify this channel causally, in a way that is difficult to achieve in active war zones. The results thus offer a tractable benchmark for the institutional and economic costs of a government's failure to maintain public order, a question that is pressing well beyond the Indian context.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data. Section 4 outlines the empirical strategy and identification assumptions. Section 5 reports the main results. Section 6 explores the mechanisms. Section 7 presents a simple model of trust, participation, and firm output. Section 8 conducts a range of robustness exercises. Section 9 concludes.

2 Literature Review

Research on communal riots in India has focused overwhelmingly on their political origins and political consequences, with far less attention paid to their economic effects. Seminal work by Varshney (2002) emphasises the role of intercommunal civic networks in mitigating violence, while Wilkinson (2004) shows how electoral incentives and party strategies shape the incidence and intensity of Hindu–Muslim riots. This political economy perspective has been supported by the construction of influential event datasets, most notably the Varshney–Wilkinson dataset (Varshney and Wilkinson, 2006) and its extension by Iyer and Shrivastava (2018), which together form the backbone of empirical work on post-independence communal violence in India.

A key methodological contribution of this literature is the Friday–festival identification strategy introduced by Iyer and Shrivastava (2018), which uses quasi-random coincidences between the Hindu ritual calendar and the weekly Muslim prayer cycle to isolate exogenous variation in riot incidence.[3] While this approach has generated important insights into the political determinants and electoral consequences of riots, existing studies have largely ignored their broader economic implications. Our paper also speaks to the economics-of-religion literature in India and to recent work on the non-political consequences of communal violence; see also Iyer (2018), Iyer (2025), and Ghose and Pandey (2025).

More broadly, our work relates to the literature on conflict and economic development. A large body of research shows that violence depresses economic activity through capital destruction, displacement, and reduced investment (see Blattman and Miguel, 2010 for an overview). In the United States, Collins and Margo (2004a,b) and DiPasquale and Glaeser (1998) document that race riots led to persistent declines in employment and large losses in urban property values. In a different context, Abrahams (2022) shows that political conflict and movement restrictions in the West Bank reduce local employment and nightlights. Beyond these cases, localised ethnic violence in democratic settings (religious conflict in Indonesia (Wilson, 2005) and the Basque conflict in Spain (Abadie and Gardeazabal, 2003)) has been shown to generate persistent economic disruptions even in the absence of full-scale civil war.

Most existing work, however, focuses on civil wars or prolonged insurgencies. By contrast, religious riots that characterise settings such as India have received far less systematic attention. Our paper fills this gap by providing district-level causal estimates of how recurrent communal riots affect local economic activity in the pre-2014 period and the mechanisms through which this happens.

A central theme in our analysis concerns the relationship between violence, financial intermediation, and the broader conditions that support economic exchange. Fisman et al. (2020) show that exposure to communal conflict in India increases intergroup animosity and leads to persistent discrimination in bank lending at the micro level. More generally, Nunn and Wantchekon (2011) and Rohner et al. (2013) argue that historical and contemporary violence can weaken trust and civic norms, with lasting consequences for trade and development. While we do not directly measure trust, our findings are consistent with these mechanisms operating at a more aggregate level. In particular, we show that riot exposure is associated not only with changes in lending behaviour but also with a contraction in local financial and commercial infrastructure.

Finally, there is work by Bohlken and Sergenti (2010), who examine the reverse relationship: how economic growth affects the likelihood of Hindu–Muslim riots. In contrast, our focus is on how these riots themselves shape local economic outcomes and the mechanism through which this happens. Other related work includes Mitra and Ray (2014), who find that growth in Muslim per-capita expenditures increases the probability of future communal violence, whereas growth in Hindu expenditures has no such effect. Field et al. (2008) show that rent control in Gujarat constrained residential sorting, limiting segregation and thereby increasing riot risk,

[3] Variants of the Friday festival strategy have since been used extensively; see, for example, (Bohlken and Sergenti, 2010; van Noort and Goyal, 2025).

illustrating how local institutional features interact with economic conditions to shape violence.

A recent and closely related study examines how internet access amplifies the economic consequences of communal riots by facilitating the spread of polarising narratives along religious lines (Adhya, 2026). That paper focuses on the interaction between riots and digital connectivity, showing that greater internet penetration intensifies the negative relationship between riot exposure and economic activity. Our approach differs in several important respects. First, rather than studying amplification through information technology, we quantify the aggregate economic impact of communal riots themselves. Second, we focus explicitly on the mechanisms through which riots affect the local economy, documenting effects on firm-level outcomes, market infrastructure, banking presence, agricultural credit, and institutional confidence. Third, we use the random timing of household interviews in the India Human Development Survey to provide new evidence on how riots affect confidence in state institutions, offering evidence on the trust-based channels that underpin economic exchange, complementing Adhya (2026)'s findings. Finally, the two studies differ sharply in scope and period: while the internet-focused analysis relies on ACLED data for 2019–2023, our analysis covers the pre-Modi period from 1994–2013, using a newly extended version of the Varshney and Wilkinson (2006) riots dataset. As a result, the two papers are complementary: Adhya (2026) highlights an important amplification mechanism in the digital era, whereas we establish the baseline economic costs of communal riots and identify the channels through which episodic religious violence translates into real economic losses.

While a large body of work documents the economic costs of civil wars and prolonged insurgencies (Blattman and Miguel, 2010), far less is known about the magnitude of the economic losses generated by recurrent, localised, and episodic violence such as communal riots. Using night-time lights as a proxy for local economic output, we causally quantify the impact of exposure to Hindu–Muslim riots on district-level economic activity. Our estimates show that even short-lived episodes of communal violence generate economically significant and persistent output losses, providing a benchmark for understanding the economic consequences of riots in the absence of the post-2014 political dynamics. The pre-2014 period we study is also substantively important in its own right. It covers the full arc of the BJP's rise as an electoral force, the aftermath of the Babri Masjid demolition in 1992, and the 2002 Gujarat riots, the most severe episode of communal violence in independent India's post-partition history, in which at least a thousand people died and at least 98,000 were displaced (Jha, 2014).

Second, this paper contributes to the literature on communal violence in India by shifting the focus from the determinants and electoral consequences of riots to their economic effects and, importantly, the underlying mechanisms. Building on the Friday–festival identification strategy of Iyer and Shrivastava (2018), we provide district-level causal evidence on how communal riots affect local economic activity and explicitly examine the channels through which these effects operate. We trace the impact of riot exposure across multiple margins, including production, labour inputs, financial intermediation, and local market infrastructure. In addition, drawing on quasi-random variation in the timing of interviews in the India Human Development Survey (IHDS), we provide the first causal estimates of the effect of riot exposure on institutional confidence. The results show that communal riots significantly erode confidence in state institutions, revealing an institutional trust channel through which episodic religious violence translates into persistent economic losses.

Third, we contribute new data by extending the Varshney–Wilkinson riots dataset (Varshney and Wilkinson, 2006). Our extension substantially expands the temporal coverage of the original event-level data, enabling the construction of a consistent district–year panel of communal violence spanning 1994–2013. This extended dataset provides a new foundation for future research on the economic and institutional consequences of communal violence in India.

3 Data

3.1 Riots

Our main explanatory variable captures the occurrence of Hindu–Muslim riots. We start from the Varshney and Wilkinson (2006) dataset covering 1950–1995. They collected the data using the Mumbai edition of the Times of India. Data for 1995-2006 come from Iyer and Shrivastava (2018), who extended the original Varshney and Wilkinson (2006) dataset. We further extend the series to 2014, following this established protocol. Specifically, we collect riot reports from digitised newspaper archives accessed via Factiva and ProQuest. We use structured keyword prompts to identify potentially relevant articles, which are then manually reviewed and coded to extract event dates and locations. Further detail on the data collection is in Appendix G.1. For the main analysis, we restrict the panel to 1994–2013 to match the availability of our nightlights variable.[4] Each event report includes the location and the date(s) where available. We harmonise all events to consistent district boundaries from 16 states.[5]

Our main explanatory variable is a continuous, distance-weighted measure of riot exposure, defined in Section 4, which aggregates nearby riot events with weights that decay smoothly with distance. This approach captures variation in exposure intensity across space and time while allowing us to include year fixed effects, a crucial feature when working with nightlights, which are subject to strong common time shocks (Henderson et al., 2011).

In the full riots panel for 1993–2014, there are 9,416 district–year observations, of which 292 (3.10%) record at least one riot. Figure 1 plots the annual number of religious riot incidents for 1993–2013, the period covered by the main nightlights outcome sample. The series peaks in 2002, corresponding to the Gujarat riots; all results are robust to excluding this year.

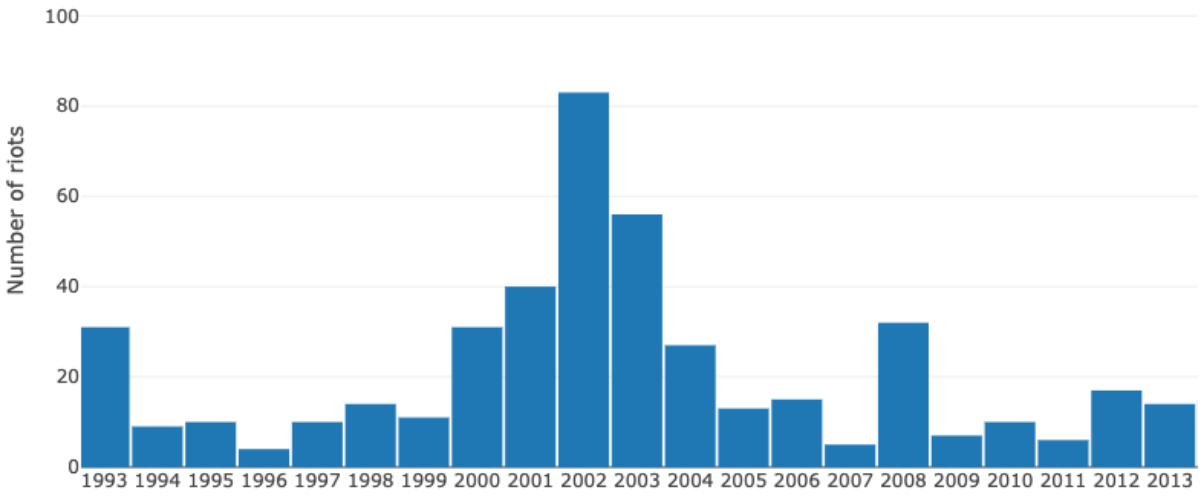


Figure 1: Number of religious riot incidents between 1993 and 2013

Following Iyer and Shrivastava (2018), we geo-code riot incidents to their most precise reported location whenever possible.[6] This procedure allows us to construct our main explanatory

[4]For the mechanism analysis, we rely on various subsets of the complete time span, depending on data availability.

[5]A single riot reported in the newspaper can mask substantial heterogeneity in duration. Some incidents last only a single day, while others persist for several days. Although a subset of entries records the exact duration, many reports describe the event only imprecisely, for example as lasting 'several days.' Hence, we abstract from duration altogether and treat each reported riot as a single event, regardless of its length.

[6]When the exact location is unavailable, we assign the centroid of the finest geographic unit reported.

variable and to assign each incident consistently to a district.

3.2 Nightlights

To measure the economic impact of religious riots at a granular level over time, we use nightlight data, as these data are available at the district level in yearly intervals.

The nightlights data come from the Operational Linescan System (OLS) on the U.S. Department of Defense's DMSP weather satellites; the series is processed and archived by NOAA and widely distributed via NASA platforms.[7] Annual global composites are available from 1992 onwards Elvidge et al. (1997, 2001), mapped to a geo-referenced 30-arc-second grid ($\approx$ 1 km at the equator). The data are cleaned and mapped to Indian districts by The Socioeconomic High-resolution Rural-Urban Geographic Platform for India (SHRUG). We use the natural logarithm of 1+ mean calibrated nightlight intensity within each district polygon by year.[8][9]

Nightlights are widely used as a proxy for subnational development, especially where income data are weak, following early work by Elvidge et al. (1997) and subsequent applications by Henderson et al. (2011); Hodler and Raschky (2014); Michalopoulos and Papaioannou (2013, 2014); Storeygard (2016). Evidence from Doll et al. (2006) shows that light density at night robustly tracks economic activity.[10][11] In India, Dhillon et al. (2016) links the 2004 National Election Study at the constituency level and reports correlations of nightlights with wealth around 0.6, and with income and education around 0.4–0.45. Overall, the literature (see Pinkovskiy and Sala-i Martin (2016)) finds a strong within-country relationship between GDP levels and nightlight intensity and growth rates. Practically, nightlight pixels aggregate cleanly to districts, and the annual frequency supports finer temporal analysis. Conceptually, increases in nightlights reflect greater local economic activity via channels such as electrification and rising incomes. Asher et al. (2021) also confirms that nightlights highly correlate with employment, population, per capita consumption, and electrification at very granular geographic levels.

In the nightlights data, each pixel's annual average brightness is recorded on a 6-bit scale (0-63), reflecting relative luminosity rather than calibrated radiance, as the DMSP–OLS sensors lack on-board radiance calibration. Consequently, year-to-year variation in observed brightness may reflect changes in sensor gain or satellite platform characteristics in addition to true changes in economic activity on the ground.

We use nightlights as a measure of economic activity for several reasons. Although we could, in principle, use the Census of India to construct various measures of public goods, it is only available for 1991, 2001, and 2011, and therefore does not allow us to examine outcomes annually. Similarly large surveys, such as the National Sample Survey, the India Human Development Survey, and the Economic Census of Firms could be used to proxy for overall economic activity,

[7] Each night—typically between 8:30 and 10:00 p.m. local time—the satellites image the entire globe from about 830 km altitude. The sensors detect concentrations of outdoor lighting, fires, and gas flares at a native resolution of $\approx$ 0.56 km, with a commonly used smoothed product at $\approx$2.7 km. Annual composites are then built over the calendar year, excluding scenes obscured by clouds or contaminated by aurora/solar glare, and filtering out ephemeral sources (e.g., fires, temporary lighting) and noise.

[8] The results are robust to different functional forms of the outcome variable, including the inverse hyperbolic sine (IHS) transformation; see Table 15.

[9] Nightlight values are calibrated to ensure consistency over time; see Asher et al. (2021).

[10] Henderson et al. (2011) show that nightlights capture short-run shocks (for example, Indonesia's 1997–98 Asian Financial Crisis and the 1993–94 Rwandan genocide), implying that nightlights are a useful proxy for economic activity at temporal and spatial scales where conventional data are scarce or unreliable. Prior work also finds that the DMSP-OLS sensor reliably detects electrified villages in developing countries, making nightlights a practical proxy for electricity access (Doll et al., 2006; Min et al., 2013; Baskaran et al., 2015). Recent applications use nightlights to study urban growth in sub-Saharan Africa (Storeygard, 2016), production in blockaded West Bank towns (Abrahams, 2022), and urban form in China (Baum-Snow et al., 2017) and India (Harari, 2020).

[11] Michalopoulos and Papaioannou (2014) shows that nightlight density correlates strongly with proxies for public goods, such as access to electricity, sewage systems, piped water, education, and broader development. Min (2015) likewise finds a strong association between electricity access, public goods provision, and nightlight intensity in low-income countries.

but are not available annually. Official GDP data from the Reserve Bank of India (RBI) are only available at the state level. District-level GDP data from ICRISAT are a potential alternative, but they do not cover Jammu and Kashmir. Moreover, the quality and construction of these data have not been widely evaluated in peer-reviewed work. For these reasons, we rely on nightlights as our main proxy for economic activity and use the ICRISAT district GDP data only as a robustness check.

Following Henderson et al. (2011) and Chen and Nordhaus (2011), we address these comparability issues by including year fixed effects, which absorb sensor-specific artefacts and other shocks common across districts in a given year. We further rely on the stable-lights product, which filters out ephemeral light sources and transient noise, and we verify robustness by restricting the sample to periods in which a single satellite was used. Together, these steps ensure that identification relies on within-district changes in relative luminosity over time, rather than on cross-year differences driven by measurement artefacts.

Table 1: Summary statistics for main and auxiliary analysis samples

Variable	Obs.	Mean	SD	Min	Max	Years
Log night lights	8,560	1.647	0.650	0.228	3.542	1994–2013
Lagged all riots (100 km)	8,560	0.483	1.450	0.000	37.408	1994–2013
Lagged riot occurred (0/1)	8,560	0.032	0.176	0.000	1.000	1994–2013
Lagged festival (0/1)	8,560	0.509	0.500	0.000	1.000	1994–2013
Lagged instrument (100 km)	8,560	5.194	5.665	0.000	20.101	1994–2013
Log bank branches	6,366	4.754	0.739	2.197	7.533	1995–2014
Number of principal markets	4,899	6.388	4.425	0	35	1990–2014
Number of sub-markets	4,288	9.729	10.181	0	66	1990–2014
Log agricultural credit	5,366	14.979	1.343	9.216	18.814	2002–2014
Log total credit	5,366	16.293	1.449	10.631	23.298	2002–2014
Log primary-sector GDDP	1,792	11.380	0.761	6.727	13.484	1993–2013
Log secondary-sector GDDP	1,792	10.866	1.012	7.408	14.309	1993–2013
Log tertiary-sector GDDP	1,792	11.786	0.952	8.717	14.826	1993–2013
Log total GDDP	2,031	12.407	1.052	7.413	15.307	1993–2013
Total output (Census subsample)	3,092	19.033	1.593	13.212	21.976	2000–2007
Capital (Census subsample)	3,092	17.921	1.987	11.305	21.863	2000–2007
Labour (Census subsample)	3,092	5.348	0.909	2.197	7.380	2000–2007
Entry share (Census subsample)	2,716	0.0017	0.0041	0	0.0296	2000–2007
Exit share (Census subsample)	2,504	0.0037	0.0073	0	0.0417	2000–2007
BJP (0/1)	8,560	0.177	0.381	0.000	1.000	1994–2013
Muslim share	8,560	0.138	0.168	0.001	0.993	1994–2013
Male share	8,560	0.517	0.015	0.458	0.599	1994–2013

Notes: Summary statistics are computed on the estimation sample relevant for each variable; sample sizes therefore differ across rows. Logged variables use the transformation $\log(1+x)$. Riot exposure measures aggregate distance-decayed events using the baseline 100 km Gaussian bandwidth. The instrument is the corresponding festival-based exposure measure. Variables are reported to three decimal places where appropriate. Muslim and male population shares are interpolated and extrapolated between census years.

3.3 Manufacturing data

Our firm-level data come from the Annual Survey of Industries (ASI), India's principal source of information on registered manufacturing establishments. We use the longitudinal ASI covering the financial years 2000–01 to 2007–08. The ASI's sampling frame is based on state-maintained

registers of factories and includes establishments employing at least 10 workers with power or 20 workers without power; as a result, it does not capture informal manufacturing activity.[12] Establishments employing 100 workers or more generally form the census sector and are surveyed annually, while the remaining units constitute the sample sector, which is surveyed according to the ASI's rotating sampling design.

Although the ASI does not cover Arunachal Pradesh, Mizoram, Sikkim, or the island of Lakshadweep, this limitation is not relevant for our analysis, as none of these states are included in the 16-state sample used throughout the paper. The ASI reports detailed firm-level information on production, assets, employment, inputs, and financial outcomes.

The ASI sampling frame covers all registered (formal) manufacturing firms in India. Large firms are classified as part of the Census sector and are surveyed every year, while smaller firms belong to the Sample sector and are surveyed intermittently. The data provides sampling weights, making the weighted sample representative at the state-industry level (Martin et al., 2017), but our analysis is at the district level. Hence, we will focus on the census subsample in our analysis.

3.4 ICRISAT data

To study the mechanisms underlying our main results, we use district-level data from the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT), drawing on its District Level Database (DLD). The DLD aggregates annual district-level information from multiple administrative sources, including the Directorate of Economics and Statistics of the Ministry of Agriculture, State Directorates of Economics and Statistics (DES), the Reserve Bank of India (RBI), and the National Bank for Agriculture and Rural Development (NABARD). The database provides measures of banking infrastructure, market access, agricultural and total credit, and district-level GDP. Coverage spans from the mid-1990s through the 2010s, closely overlapping with our riots panel, though data availability varies by variable and source.

Our main analysis is conducted on a balanced set of 16 states, listed in Table 2. The ICRISAT DLD, by contrast, covers a broader set of 20 states but does not include Jammu and Kashmir, which is part of our main analysis.[13]

Banking infrastructure data are compiled by ICRISAT from State Directorates of Economics and Statistics and are available for the period 1995–2014. Market access measures, assembled using data from State Directorates of Agricultural Marketing, Public Works Departments, and State DES, span 1990–2014 but contain missing observations across districts and years; results based on market counts should therefore be interpreted with caution. Finally, district-level credit data are compiled by ICRISAT from the Reserve Bank of India and are available for the years 2002–2014. Variable-specific coverage is discussed in greater detail in the corresponding sections.

3.5 IHDS data

To study how riots affect confidence in institutions, we use individual-level data from the India Human Development Survey (IHDS), a nationally representative household survey conducted jointly by the University of Maryland and the National Council of Applied Economic Research. We draw on both waves of the survey—IHDS-I (2004–2005) and IHDS-II (2011–2012)—which provide detailed information on household demographics, socioeconomic characteristics, and individuals' confidence in a wide range of political, legal, and service-delivery institutions. The IHDS has been widely used in applied economic research on India (e.g., Afridi et al. (2018)).

[12] More details on the ASI sampling design and survey methodology are available on the website of the Ministry of Statistics and Programme Implementation, Government of India.

[13] The main results are robust to excluding Jammu and Kashmir.

A feature of the IHDS for our purposes is that interviews within a given district and survey wave were conducted over an extended fieldwork period rather than on a single date. This staggered interview timing allows us to use variation in whether respondents were interviewed before or after a riot occurred in their district, while holding fixed district-level characteristics and survey-wave conditions.[14]

Our analysis focuses on confidence in institutions that are directly relevant for economic coordination and governance, including politicians, the police, courts, state government, the military, local panchayats, and the news media. These variables are measured using ordered categorical responses that capture respondents' stated confidence in each institution. We complement these outcomes with a set of placebo variables, such as confidence in schools and health providers, as well as household asset ownership, that are unlikely to respond immediately to riot exposure and are used to assess the validity of the identification strategy.

To maintain consistency across survey waves and to avoid the reassignment of households over time and space, for the household level analysis, we use the district defined in the 2011 Census[15]. We further restrict attention to district–survey-wave cells in which exactly one riot occurred during the IHDS interview period. This restriction ensures a well-defined treatment event and allows for clean comparisons of respondents interviewed before and after the same riot within the same district and survey wave. Applying this restriction yields a final sample of 1,088 individual observations across 8 districts. Summary statistics for the main outcomes and covariates are reported in Appendix F.

4 Empirical Strategy

We begin by outlining a simple, benchmark specification that relates local economic activity to exposure to communal riots. Let $y_{i,t}$ denote a proxy for economic activity in district i in year t, measured using annual night-time light intensity. Let $Riots_{i,t-1}$ capture riot exposure in district i during the previous year. In its most naïve form, the relationship of interest can be written as:

$$y_{i,t} = \alpha + \beta Riots_{i,t-1} + \epsilon_{it}. \tag{4.0.1}$$

This specification asks whether districts that experience greater riot exposure subsequently exhibit lower levels of economic activity. We use lagged riot exposure for two reasons. First, economic losses from violence typically materialise with a delay, as disruptions to production, labour supply, and trade unfold over time rather than instantaneously. Second, contemporaneous night-time lights may be mechanically affected by the conflict itself: fires, emergency lighting, and security operations can temporarily increase radiance, masking underlying economic conditions. Lagging riot exposure mitigates these short-run, non-economic lighting effects and yields a cleaner estimate of how violence propagates into subsequent economic activity.[16]

While equation 4.0.1 provides a natural starting point, its causal interpretation is problematic. Time-varying unobserved factors, such as political tensions or local economic shocks, may simultaneously affect both riot incidence and economic outcomes. As a result, ordinary least squares estimates of β may be biased.

To address the endogeneity of riot exposure, we build on the instrumental-variable strategy developed by Iyer and Shrivastava (2018). Newspaper accounts indicate that many communal riots occur around highly visible religious processions and public gatherings on significant holy days. Fridays draw especially large Muslim congregations for weekly prayers, creating dense crowds around mosques. When a major Hindu festival coincides with a Friday, competition over public space can intensify, increasing the likelihood of communal tension and, in some cases,

[14] See Appendix G for details.

[15] As collected from the merged IHDS data available on their website

[16] Conflict can temporarily raise night-time radiance (e.g., fires or emergency lighting), so contemporaneous regressions are biased upward, attenuating the estimated negative effect.

violent incidents. Because Hindu festival dates follow a lunar calendar, their timing shifts across years and varies across regions, generating plausibly exogenous variation in riot risk.

Guided by this logic, we first construct an indicator that captures the coincidence of Fridays with major Hindu festivals. For each state, we follow Iyer and Shrivastava (2018) and focus on the five most important Hindu festivals (see Table 2). National festivals such as Dussehra and Diwali are celebrated widely, while more region-specific festivals (e.g., Holi in Madhya Pradesh, Ganesh Chaturthi in Maharashtra) are assigned to the relevant states. For each district–year, this baseline instrument equals one if at least one of the selected festivals in the corresponding state falls on a Friday in that year, and zero otherwise.

	State	**Festivals**				
No.		I	II	III	IV	V
1	Andhra Pradesh	Ramnavami	Durga ashtami	Navami	Dushehra	Diwali
2	Assam	Janmashtami	Durga ashtami	Navami	Dushehra	Diwali
3	Bihar	Holi	Ramnavami	Navami	Dushehra	Diwali
4	Gujarat	Holi	Ramnavami	Navami	Dushehra	Diwali
5	Haryana	Shivratri	Holi	Janmashtami	Dushehra	Diwali
6	Jammu and Kashmir	Shivratri	Ramnavami	Janmashtami	Dushehra	Diwali
7	Karnataka	Shivratri	Ganesh Chaturthi	Navami	Dushehra	Diwali
8	Kerala	Shivratri	Janmashtami	Navami	Dushehra	Diwali
9	Madhya Pradesh	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
10	Maharashtra	Ramnavami	Ganesh Chaturthi	Navami	Dushehra	Diwali
11	Orissa	Holi	Durga ashtami	Navami	Dushehra	Diwali
12	Punjab	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
13	Rajasthan	Holi	Ramnavami	Janmashtami	Dushehra	Diwali
14	Tamil Nadu	Janmashtami	Ganesh Chaturthi	Navami	Dushehra	Diwali
15	Uttar Pradesh	Ramnavami	Janmashtami	Navami	Dushehra	Diwali
16	West Bengal	Holi	Durga ashtami	Navami	Dushehra	Diwali

Table 2: State-wise list of major festivals.

While this construction captures plausibly exogenous shifts in riot risk, our empirical setting motivates two extensions. First, riots may generate spillovers across space: violence in nearby districts can affect local economic activity through channels such as reduced trust, disrupted trade, or heightened uncertainty. Second, the inclusion of year fixed effects, which is standard in analyses using night-time lights to absorb common macroeconomic shocks and satellite-specific artefacts (Henderson et al., 2011), requires sufficient cross-sectional variation in the instrument.

To accommodate these features, we extend the baseline approach by constructing a spatially weighted measure of riot exposure. Specifically, we assume that multiple riots exert cumulative effects on economic activity and that the influence of each riot decays smoothly with distance. We define:

$$RiotExposure'_{i,t} = \sum_{j=1}^{n(t)} \phi\big(d_{ijt}\big),$$

where d_{ijt} denotes the distance between district i's centroid and riot j occurring in year t, $n(t)$ is the total number of riots in year t, and $\phi(\cdot)$ is a Gaussian distance-decay function. This measure allows riots in nearby districts to affect local outcomes while placing greater weight on events occurring closer to district i.[17]

We instrument this spatially weighted riot exposure using an analogous distance-weighted measure of Friday–festival coincidences, which we define formally below. Because the instrument varies across districts within a given year through spatial proximity, this approach allows us to include year fixed effects without absorbing all identifying variation.

[17] Iyer and Shrivastava (2018) refer to this object as $AllRiots'_{i,t}$.

Finally, as riot incidence is measured using newspaper reports, under-reporting may bias our results. We address this concern explicitly in Section 5.1.

Hence, our main specification is the regression:

$$y_{it} \;=\; \alpha_i \;+\; \beta\,\text{RiotExposure}'_{i,t-1} \;+\; \gamma_t \;+\; \varepsilon_{it}, \tag{4.0.2}$$

where $y_{it} \equiv \ln\!\big(1+\text{mean calibrated night light}_{it}\big)$ for district i in year t; α_i and γ_t are the district and year fixed effects. Our exposure measure $\text{RiotExposure}'_{it}$ aggregates all riots in year t using a Gaussian distance–decay kernel,

$$\text{RiotExposure}'_{it} \;=\; \sum_{j=1}^{n(t)} \phi\big(d_{ijt}\big), \qquad \phi(d) \;=\; \exp\!\big(-d^2/2\sigma^2\big),$$

with d_{ijt} the great-circle distance from district i's centroid to riot j's location in year t, and σ the standard deviation parameter of our decay function. We report robustness to alternative standard deviation parameters, which modulate the spatial reach of a riot's influence on the economy.

We adopt the district delineations of Martin et al. (2017) (1998 district boundaries from ASI) to keep the night-lights and firm analyses on the same geographic units. This partition yields 541 districts. Using India's total land area of 3,287,263 km^2, the implied mean district area is 6,076 km^2,[18] which is equivalent to a circle with a radius $\approx$ 44 km. Guided by this scale, we set the baseline Gaussian distance–decay bandwidth to 100 km and increase it in 50 km steps, as reported in Table 3.

To build intuition for our independent variable, note that the effect of far-away riots is very small. For example, using the coordinates in our data, the great-circle distance from Patna to the Lucknow district centroid is approximately 443 km. With our baseline bandwidth $\sigma = 100$, the marginal contribution of such a riot to Lucknow's exposure is approximately 5.6×10^{-5}.[19] By comparison, a riot 150 km away from a district centroid is weighted by approximately 0.325;[20] hence, the kernel decays rapidly with distance.

Our distance-decayed riot exposure variable therefore has a clear and interpretable scale. With a Gaussian kernel and $\sigma = 100$ km, a riot occurring at the district centroid contributes exactly 1. A riot 40 km away contributes 0.92, a riot 60 km away contributes 0.84, and a riot 100 km away contributes 0.61. Thus, a "one-unit" change corresponds very closely to a single riot occurring within the district or to several riots occurring nearby whose weighted contributions sum to one. For example, three riots located 20, 40, and 60 km from the centroid generate an exposure value of 2.74, while two distant riots at 100 km and 140 km generate an exposure value of 0.98.

To address the endogeneity of riot exposure, we instrument $\text{RiotExposure}'_{i,t-1}$ with a distance-weighted analogue of the festival coincidence used by Iyer and Shrivastava (2018):

$$Z_{i,t-1} \;=\; \sum_{k} \phi\big(d_{ik}\big)\,\text{Festival}_{k,t-1}.$$

Here $\text{Festival}_{k,t-1} = 1$ if, in district k during the year $t-1$, a listed important Hindu festival falls on a Friday; d_{ik} is the distance between districts i and k. Intuitively, when nearby districts experience the Friday–festival coincidence, the risk of riots in and around district i rises. Because the instrument varies across districts within a given year through spatial proximity, this construction allows us to include year fixed effects without absorbing the relevant identifying variation.

[18] Computed as 3,287,263/541.

[19] This calculation uses the Patna riot coordinate in our data, $(25.611224, 85.130692)$, and the Lucknow district centroid, $(26.813461, 80.901548)$. The Gaussian kernel gives $\exp[-443^2/(2\cdot100^2)] \approx 5.6\times10^{-5}$.

[20] $\exp[-150^2/(2\cdot100^2)] \approx 0.325$.

4.1 Identification Assumptions

4.1.1 Identification and the role of fixed effects

Because $Z_{i,t-1}$ varies *within* a year across districts via distances d_{ik}, we can include year fixed effects γ_t without saturating the instrument. Year fixed effects are crucial with night lights: they absorb global shocks and sensor-year idiosyncrasies, as well as nationwide macro shocks to the economy. District fixed effects α_i absorb all time-invariant local confounders (baseline brightness, settlement patterns, long-run religiosity, etc.). Under the exclusion restriction that Friday-festival timing affects y_{it} only through riots, β identifies the causal effect of riot exposure on nighttime luminosity.

Our baseline specification uses lagged riot exposure. Any direct effects of Friday–festival coincidences on economic activity, such as temporary market closures or changes in lighting, are short-lived and confined to the festival year. These transitory shocks cannot mechanically generate declines in nightlight luminosity, firm outcomes, or credit.

4.1.2 Exclusion Restriction

Our instrument $Z_{i,t-1}$ uses quasi-random coincidences in the ritual calendar: years in which a major Hindu festival falls on a Friday in (nearby) districts. A potential concern is that the Friday festival coincidence could affect local economic activity directly rather than only through riots. For example, contestation over public space on such days might transiently alter commerce or mobility independent of any riot incident.

Any direct effect of the Friday–festival coincidence on y_{it} at the annual district scale is negligible for three reasons. First, DMSP–OLS stable-lights are annual composites that explicitly filter short lived sources; a localised one-day spike in illumination or brief disruptions to commerce will be averaged out in a district of roughly 6,000 km^2. Second, year fixed effects absorb any nationwide festival timing patterns, while district fixed effects account for time-invariant local propensities to celebrate or illuminate. Third, the likely direct sign runs against our result: festivals tend to raise light output (extra lighting, gatherings), which would push the reduced form toward higher nightlights and thus attenuate, rather than generate, the negative 2SLS effect we attribute to riots. In that sense, our estimates are conservative with respect to this threat.

4.1.3 Sample and inference

Nightlights are available for 1994–2013. Our main analysis therefore focuses on this period, and we assess robustness to alternative samples and subsample restrictions in Section 8. Other outcome variables used to study mechanisms are available for different time periods, as outlined in Table 1. Standard errors are clustered at the district level. As a robustness check, we compute spatial HAC standard errors (Conley, 1999, 2008) that allow disturbances to be correlated across nearby districts and over time for different cut-offs and lags (using a Bartlett kernel). Although these standard errors are slightly larger, inference remains unchanged.

5 Results

As mentioned before, the independent variable is defined by the Gaussian distance decay kernel, which depends on the variable σ, indicating the bandwidth within which nearby riots impact a district. We present the AIC and BIC in order to choose the model in Table 3. As is clear from the table, we will use the 100 km SD as it has the lowest AIC and BIC, indicating the best model fit. Furthermore, we want the model to allow riots within a given location to have a higher impact on the economic outcomes. Overall, it is evident that with a higher σ, the coefficient of interest decreases. This is intuitive: placing more weight on distant districts, whose riots have

little effect on local outcomes, attenuates the coefficient. However, note that all the coefficients are significant and are of comparable magnitude.

Table 3: Regression results for different standard deviations for $\phi(\cdot)$

SD in km	Coefficient	SE	KP F	Root MSE	AIC	BIC
100	**-0.105*** **	**0.023**	**56.345**	**0.1982**	**-3433.985**	**-3405.766**
150	-0.075***	0.015	61.503	0.2067	-2712.530	-2684.311
200	-0.064***	0.011	66.050	0.2181	-1790.614	-1762.395
250	-0.059***	0.010	74.085	0.2308	-824.340	-796.120
300	-0.056***	0.009	86.649	0.2438	113.526	141.745

Notes: The dependent variable is ln(1 + mean calibrated night light), winsorised at the 1st and 99th percentiles. Each row reports a separate 2SLS specification using a different Gaussian distance-decay bandwidth. The coefficient is on lagged riot exposure, instrumented by the corresponding lagged festival-based exposure measure, both constructed with the standard deviation shown in the first column. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level. KP F is the Kleibergen–Paap first-stage F statistic. The bold row marks the bandwidth selected by AIC and BIC. Significance levels: * 10%, ** 5%, *** 1%.

For the rest of the results for the full sample regressions, we will stick to the specification with SD=100 unless stated otherwise.

Table 4 reports the baseline OLS and instrumental-variable results linking riot exposure to nightlights. Column (1-2) presents the naive OLS regression of $\ln(1+\text{mean calibrated nightlights}_{it})$ on lagged riot exposure, $\text{RiotExposure}_{\text{sd100},i,t-1}$. Column (3-4) reports the reduced-form regression of $\ln(1+\text{mean calibrated nightlights}_{it})$ on the instrument $Z_{i,t-1}$. Column (5-6) presents the 2SLS estimates of the effect of riot exposure on nightlights using $Z_{i,t-1}$ as an instrument. Column (7-8) reports the corresponding first-stage regression of $\text{RiotExposure}_{\text{sd100},i,t-1}$ on $Z_{i,t-1}$.

The coefficient in the OLS regression is very small and positive. However, the 2SLS regression results are negative, indicating that the naive regression suffered from positive omitted variable bias. This could be because increasing urbanisation is likely positively correlated with both the dependent and independent variables, leading to the positive omitted variable bias.

The Kleibergen–Paap F-statistic is 56.35 in our preferred specification (59.16 without controls), well above conventional thresholds, indicating no weak-instrument concerns. The first stage is positive and significant at the 1% level, confirming that the instrument increases riot incidence.

Table 4: Structural, reduced form, IV, and first-stage results

	(1) OLS	(2) OLS	(3) Reduced	(4) Reduced	(5) IV	(6) IV	(7) 1st stage	(8) 1st stage
Dependent var.	$y_{i,t} = \ln(1 + \text{Mean calibrated night light}_{i,t})$						$\text{RiotExposure}_{\text{sd100},i,t-1}$	
$\text{RiotExposure}_{\text{sd100},i,t-1}$	0.0023***	0.0034***			-0.1237***	-0.1051***		
	(0.0008)	(0.0008)			(0.0233)	(0.0229)		
$\text{Z}_{\text{sd100},i,t-1}$			-0.0040***	-0.0032***			0.0327***	0.0308***
			(0.0006)	(0.0006)			(0.0043)	(0.0041)
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Districts	428	428	428	428	428	428	428	428
Observations	8,560	8,560	8,560	8,560	8,560	8,560	8,560	8,560
KP rk Wald F					59.16	56.35		

Notes: The outcome is ln(1+mean calibrated night light), winsorised at the 1st and 99th percentiles. Columns (1)–(2) report OLS estimates, columns (3)–(4) report reduced-form estimates, columns (5)–(6) report IV estimates, and columns (7)–(8) report the corresponding first-stage regressions. Riot exposure and the festival-based instrument are both constructed using the baseline Gaussian bandwidth of 100 km and enter with a one-year lag. Controls are BJP incumbency, Muslim population share, and male population share. All specifications include district and year fixed effects. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: * 10%, ** 5%, *** 1%.

5.1 Adjusting for reporting error

Our riot measures are constructed from newspaper reports, raising the concern that not all riots are observed. A formal treatment of under-reporting in newspaper based riot data is provided by Iyer and Shrivastava (2018), who derives the bias induced by imperfect reporting under a set of assumptions about the reporting process. In particular, assuming (i) no over-reporting, (ii) reporting probabilities independent of the instrument, and (iii) riot effects independent of whether the riot is reported, Iyer and Shrivastava (2018) shows that under-reporting leads to the magnitude of the IV estimates being biased upward by a factor inversely proportional to the reporting rate.

Let p denote the probability that at least one riot is reported, conditional on a riot having occurred. As shown in Iyer and Shrivastava (2018), in a regression without additional covariates, the IV estimate obtained using reported riots, $\hat{\beta}_{IV}$, is related to the true IV coefficient, $\hat{\beta}^{*}_{IV}$, as

$$\hat{\beta}_{IV} \;=\; \frac{1}{p}\,\hat{\beta}^{*}_{IV}.$$

Existing validation exercises suggest reporting rates of roughly p = 0.77–0.78 for Hindu–Muslim riots in India when comparing Times of India, Mumbai edition newspaper accounts to administrative records and eyewitness data (Brass, 2003; Jha, 2014). Importantly, if media coverage amplifies economic responses to riots, the coefficient estimate after correcting for under-reporting would be a lower bound for the coefficient. Using p=0.77, our main specification results change from 0.59% to 0.45% (using Bickenbach et al. (2016)'s elasticity of log GDP with respect to log night-lights of 0.056 for Indian districts).

A related concern is measurement error arising from our digitised search procedure used to identify riots in the Times of India archive. To assess the accuracy of this approach, we conduct a manual validation exercise using the Varshney–Wilkinson (VW) dataset.

We randomly sampled 100 riot incidents from the VW dataset occurring in 1995 and earlier (the time period of the original VW dataset), and manually verified whether these incidents were detected by our search algorithm in the Times of India on the dates corresponding to the reported riot occurrences. Two incidents could not be located in the Times of India at all on the relevant dates, possibly due to errors in the reported dates in the VW dataset (or issues

with the digital archive). We exclude these two cases from the validation sample.[21]

Among the remaining 98 incidents, 7 are detected by the algorithm but are non-religious riots, and 7 religious riots are missed by the algorithm. This implies a total classification error rate of 0.0769. Overall, this exercise suggests that misclassification due to the digitised search procedure is limited in magnitude and unlikely to impact the results.

To further alleviate concerns that digitisation or reporting errors drive our results, we re-estimate our main specification, restricting the sample to 2005 and earlier. This period corresponds to the riot dataset constructed by Iyer and Shrivastava (2018), which is based on manual coding from physical newspaper copies rather than digitised search.[22]

Using this restricted sample, we obtain a coefficient of -0.070, which implies a decline in GDP of approximately 0.30%, bearing in mind the under-reporting. This estimate is quantitatively close to our baseline result, suggesting that measurement error arising from the digitised search procedure is unlikely to be driving our main findings.

6 Mechanisms

Riots may affect economic activity through multiple channels. First, violence can directly disrupt firm-level production by damaging assets, preventing workers from attending work, and increasing uncertainty, potentially affecting firms' scale of operation as well as entry and exit decisions. Second, riots may weaken the local institutional environment by reducing the presence of banks and markets that facilitate trade, credit, and exchange. Third, heightened social tension may lead to a contraction or reallocation of credit, as violence can erode trust and increase perceived lending risk.

Other mechanisms may also be at play but are more difficult to examine empirically. Riot exposure may induce out-migration of workers or entrepreneurs, reducing local economic activity, but annual district-level migration data are not available.[23] Similarly, riots may affect economic outcomes through changes in social trust and intergroup relations; while trust is not directly observable at the district-year level, prior work shows that riot exposure reduces lending to affected groups within Indian banks, consistent with a deterioration in trust and confidence (Fisman et al., 2020). Finally, riots may influence public investment decisions, either through increased reconstruction spending or through neglect of riot-prone areas, but detailed annual district-level data on public goods provision are generally not available annually.

In this section, we study the mechanisms for which district-level yearly data are available. We first examine firm-level production responses in formal manufacturing, focusing on adjustments in labour, capital, output, and firm turnover. We then study changes in local economic institutions (bank branches and markets), and finally examine the composition and allocation of district-level credit.

6.1 Manufacturing Firms

We begin by examining firm-level responses in formal manufacturing. Although formal manufacturing accounts for a relatively small share of employment in India, the Annual Survey of Industries (ASI) is the only data source that allows us to separately study adjustments in labour, capital, output, and firm turnover at an annual frequency at the district level. For this reason, the ASI is well suited to studying production-side mechanisms. Furthermore, any effects

[21] If one were to assume that these were missed due to issues with the archive, then the classification error is still less than 10% at 0.0968

[22] Iyer and Shrivastava (2018) collects data until July 2006.

[23] Qualitative evidence from affected industrial areas supports this mechanism. Reporting from Muzaffarnagar following the 2013 riots documents that migrant Muslim workers left the area and did not return due to fear of renewed targeting (Indian Express, 2013).

observed in the formal sector are likely to represent a lower bound, as the informal sector is plausibly even more adversely affected by religious riots.

We study firm responses along three margins: levels, growth, and extensive-margin adjustment. Level regressions capture changes in the scale of activity following riot exposure, while growth regressions capture changes in the rate of adjustment of activity over time. Studying both objects allows us to distinguish between changes in levels and changes in dynamic responses. We then examine entry and exit to assess whether riots reshape the local firm distribution.

To examine this channel, we estimate 2SLS specifications using ASI plant-level data. Table 5 reports the effects of lagged riot exposure on district-average plant capital, labour, total output, and on district-level entry and exit, for the census part of the ASI dataset.

We begin with level regressions, $\ln(1 + x_{it})$. In the census sample, we find a statistically significant decline in labour one year after riot exposure: the coefficient implies a reduction of approximately 8.3 percent. The corresponding effects on capital and total output are negative but imprecisely estimated and not statistically significant at conventional levels (the p-value for total output is 0.110). Overall, the level results suggest limited effects on capital and output, alongside a marked contraction in labour input.

We next study growth regressions, $\ln(1+x_{it}) - \ln(1+x_{it-1})$. The census subsample exhibits sharp short-run contractions following riot exposure. Labour growth declines by about 17.7 percent and output growth by about 13.0 percent one year after a riot. Capital growth is also negative, though less precisely estimated. Taken together, these results indicate that riots primarily operate by reducing firms' labour input. This mechanism is consistent with Ksoll et al. (2023), who show that electoral violence lowers firm performance partly through worker absenteeism.

To examine whether riots reshape the firm distribution, we study district-level entry and exit within the census subsample. Exit is measured using a cohort-based index that tracks the survival of plants observed in 2000. Cohorts are defined by plant age in 2000, and for each cohort we compare the decline in plant mass within a district to the corresponding national cohort-specific exit probability. District-level exit shares are constructed by aggregating over cohorts using district-specific initial cohort weights, following Chatterjee et al. (2025). Entry is measured as the district's share of national entrants, using plant age information and sampling multipliers, again following Chatterjee et al. (2025).[24] Further details on the construction of the entry and exit measures are provided in Appendices E.1 and E.1.1.

At the one-year horizon, we find no evidence that riots significantly affect either entry or exit. Point estimates are close to zero and statistically insignificant, indicating little short-run adjustment along the extensive margin. This absence of immediate entry and exit responses is consistent with the frictions emphasised in the literature, which argues that entry and exit in Indian formal manufacturing are slow-moving (Chatterjee et al., 2025).

Overall, the census-based evidence suggests that riots primarily affect economic activity through contractions in labour input rather than through capital adjustment or firm turnover. These results should be interpreted as mechanisms operating within the formal manufacturing sector.

6.2 Market institutions

Table 6 presents the IV estimates for the effect of riot exposure on the local presence of banks. Using our baseline specification with one-year lagged exposure, we find a large and statistically strong contraction in the number of bank branches. A one-unit increase in Gaussian-weighted

[24] Because plant identifiers are difficult to track, as intermittent survey participation can make continuing plants appear to exit and re-enter, we define entrants as plants whose reported age is one, rather than relying on firm identifiers.

Table 5: Effect of lagged riot exposure on capital, labor, and total output

	Census subsample							
	(1) Capital	(2) Labor	(3) Total output	(4) Entry	(5) Exit	(6) ΔCapital	(7) ΔLabor	(8) ΔTotal output
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.087	-0.087***	-0.083	0.000	0.000	-0.097	-0.195***	-0.139**
	(0.065)	(0.032)	(0.052)	(0.000)	(0.000)	(0.080)	(0.051)	(0.060)
Observations	3088	3088	3088	2712	2499	2666	2666	2666
Districts	398	398	398	398	368	395	395	395
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Unit FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
KP rk Wald F	34.9	34.9	34.9	28.3	23.2	27.6	27.6	27.6

Notes: Each column reports a separate IV regression for the fully enumerated ASI subsample. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth, and the instrument is the corresponding lagged festival-based exposure measure. Outcome variables are winsorised at the 1st and 99th percentiles; columns (6)–(8) use first differences of the winsorised outcomes. All specifications include unit and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the unit level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

riot intensity corresponds to an estimated 21.4 percent decline in the number of banks. This indicates that banks respond swiftly to violent unrest and begin scaling back operations within the following year.

This mechanism is economically important. Banks provide working-capital finance, liquidity, savings instruments, and credit to local firms and households, and a large literature shows that local banking presence plays a central role in shaping real economic outcomes. In the Indian context, Burgess and Pande (2005) demonstrate that bank branch expansion causally reduced poverty and increased economic activity, while more generally Guiso et al. (2004) and Beck et al. (2000) show that local financial development facilitates investment, production, and growth. A contraction in banking presence can therefore have immediate and economically significant effects on local economic activity.

Furthermore, these findings align with recent work by Fisman et al. (2020), who show that riot exposure increases intergroup animosity and induces discriminatory lending behaviour within Indian banks. Our evidence is consistent with a broader deterioration in trust and confidence in the local financial environment: if banks perceive riot-prone districts as riskier or socially fractured, they may reduce branch activity, adopt more conservative lending practices, or halt expansion.

We observe similar dynamics for local market structure. The number of sub-markets declines sharply following riot exposure, with the one-year lag coefficient around -1.648. Interpreted in levels, this implies that an additional unit of riot intensity corresponds to roughly 1.6 fewer sub-markets operating in the district by the following year. In contrast, principal markets show a more muted response: the coefficient is negative but smaller in magnitude and less precisely estimated. The evidence suggests that smaller market entities react quickly to worsening security conditions, whereas larger markets may adjust more slowly or require more severe or persistent instability to exit.

Taken together, the banking and market results highlight an important mechanism through which riots depress local economic performance: a contraction in the physical and institutional infrastructure that supports commerce. Riots reduce the presence of banks and erode the density of local markets, implying fewer trading opportunities, weaker credit availability, and higher operating risks for firms and households. These institutional and structural responses complement the firm-level patterns documented earlier, the short-run contractions in labour input, and help explain the persistent decline observed in aggregate night-lights and credit flows.

6.3 Credit

Table 6 presents the IV estimates using our preferred one-year lag specification. Column 5 shows that lagged riot exposure leads to a significant contraction in agricultural credit. Interpreted in levels, the coefficient of -0.119 corresponds to an estimated 11.2% decline in agricultural credit in the year following increased riot activity. This pattern is consistent with the findings of Fisman et al. (2020), who document that riots increase intergroup animosity and induce more conservative and discriminatory lending behaviour within Indian banks. Given that agricultural loans are typically risky Binswanger and Khandker (1995), the decline in credit to this sector indicates a tightening of lending standards in response to increased social tension. Although communal riots are largely an urban phenomenon, banks in an affected district appear to pull back from their riskiest lending, and agricultural loans are among the first to be cut. As a result, agricultural credit can fall even in rural areas.

Column 4 shows the results for total district-level credit. The coefficient is small and statistically indistinguishable from zero, implying no detectable aggregate reduction in overall lending beyond the sharp contraction observed in agriculture. This pattern suggests that banks partly reallocate lending toward less risky sectors. The evidence, therefore, points to targeted retrenchment rather than broad-based credit withdrawal.

Taken together with the results on banking presence and market structure, a coherent pattern emerges. Riot-exposed districts experience a reduction in the number of operating bank branches, a decline in sub-markets, and a substantial contraction in agricultural lending. These segments of the local economy rely heavily on stable relationships, repeated interactions, and predictable enforcement. By increasing perceived risk and uncertainty, riots are associated with a retrenchment of the financial and commercial infrastructure that supports local trade, investment, and liquidity provision. This institutional contraction provides a plausible explanation for the persistent reductions in nightlights luminosity observed in districts more exposed to spatial riot intensity.[25]

Table 6: Effect of lagged riot exposure on banking, markets, and credit

	(1) Banks	**(2)** Principal markets	**(3)** Sub-markets	**(4)** Total credit	**(5)** Agricultural credit
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.241***	-0.118	-1.648**	0.091	-0.119***
	(0.072)	(0.157)	(0.807)	(0.058)	(0.044)
District FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Districts	414	356	325	414	414
Observations	6,186	4,589	4,007	5,366	5,366
KP rk Wald F	16.67	11.30	10.54	17.43	17.43

Notes: Each column reports a separate IV regression. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. Outcomes are winsorised at the 1st and 99th percentiles. Banking, principal markets, total credit, and agricultural credit are measured as $\ln(1+x)$; sub-markets are measured in levels. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

6.4 Trust in institutions

One of the key mechanisms through which riots may affect economic and social outcomes is the erosion of trust between economic agents. In the Indian context, however, there is no large-scale household survey that directly measures intergroup trust, such as trust between Hindus and

[25] Data availability differs across outcomes, so the credit, market, and banking regressions are estimated on different district-year samples.

Muslims. The World Values Survey includes questions on generalised trust, but its geographic aggregation is above the district level and therefore unsuitable. As an alternative, we use data from the India Human Development Survey (IHDS) waves 1 and 2, which provide rich household-level information on individuals' confidence in a wide range of institutions, including the police, courts, government, media, schools, and health providers. While these measures do not capture intergroup trust directly, institutional trust is closely related to it.

To estimate the causal impact of riots on institutional trust, we use quasi-random variation in the timing of IHDS interviews relative to riot events, following a growing literature that uses random survey timing as a source of identification (e.g., Guo and An (2022); Depetris-Chauvin et al. (2020)).

In our data, there are nine district–survey-wave cells in which at least one riot occurred during the period over which the IHDS interviews were conducted in that district. One district–wave experienced two riots during the interview window; we drop this case to keep the analysis consistent, with a well-defined treatment. We further verify that, for the remaining eight district–survey waves, there were no riots in the 60 days prior to the first interview, ensuring a clean pre-period.

This setting allows us to compare respondents interviewed just before a riot in their district to respondents interviewed just after the same riot, within the same district and survey wave.

We estimate the following specification:

$$y_{idt} = \beta Post_{dt} + \delta X_i + \lambda_{dt} + \varepsilon_{idt}$$

where y_{idt} denotes the outcome of interest for individual i living in district d, interviewed in survey wave t. Outcomes include measures of confidence in various institutions, such as politicians' ability to fulfil promises, the police's ability to enforce the law, courts' ability to deliver justice, and confidence in public and private service providers. The main explanatory variable, Post_{dt}, is an indicator equal to one if individual i was interviewed on or after the date of the riot in district d, and zero if the individual was interviewed before the riot. The vector X_i includes predetermined individual and household characteristics, such as indicators for being Hindu or Muslim, urban residence, and the highest level of education attained by an adult in the household. The fixed effect λ_{dt} denotes *district* $\times$ *survey-wave fixed effects*, which absorb all district-specific characteristics as well as wave-specific district shocks. Because we restrict the sample to district–wave cells in which a riot occurred during the interview period, including these fixed effects is equivalent to including *event-specific fixed effects*. Identification therefore comes from comparisons of individuals interviewed before and after the *same riot* in the same district. The coefficient β captures the effect of riot exposure on individual trust outcomes. Identification relies on the assumption that, in the absence of a riot, outcomes would have evolved smoothly with interview timing within a given district and survey wave. Under this assumption, variation in interview timing relative to the riot date provides plausibly exogenous exposure to violence. Standard errors are clustered at the primary sampling unit (PSU) level—villages in rural areas and urban neighbourhoods in cities—allowing for arbitrary correlation among respondents within the same local sampling unit.[26]

To focus on short-run effects and strengthen the plausibility of the identifying assumption, we restrict the analysis to respondents interviewed within ± 50 days of the riot date. We also present robustness checks using a wider window of ± 350 days to assess whether the estimated effects persist beyond the immediate aftermath of the riot as our main results were the effects of lagged riots on economic outcomes.

Our identification strategy relies on the plausibly exogenous timing of a riot relative to the timing of IHDS surveys. Since it uses scenarios where the timeline of the survey in one region

[26]Note that we do not cluster at the district level due to the potential issues associated with having a small number of clusters (Cameron and Miller, 2015).

overlapped with the date of the riot in the same region, it provides a unique advantage in understanding individuals' responses to district-level riots. We probe deeper into this short-run impact of riots by narrowing the time window, focusing on the sample of respondents interviewed within 50 days around the local riot (41.4% of observations) and further on respondents interviewed within 350 days (99.9% of observations) to assess whether the effects are indeed persistent.

A potential concern is that the presence of riots may disrupt the implementation of the IHDS survey, leading to systematic differences between respondents interviewed before and after the event. While it is implausible that riots and survey fieldwork are coordinated, riots could, in principle, affect access to certain neighbourhoods or households.

We address this concern in several ways. First, the IHDS employs random sampling at every stage of survey design. IHDS-1 is a nationally representative survey of 41,554 households conducted in 2004–2005, covering 384 of the 593 districts in India as defined in the 2001 Census. The rural sample was drawn using stratified random sampling, while the urban sample was selected using probability proportional to population. Because these sampling decisions were presumably made well before the riots in our sample, it is unlikely that low-casualty riot events occurring during the interview period affected the underlying sampling design. Second, our identification relies on comparisons within district–survey waves, greatly limiting the scope for compositional changes to drive the results. Any time-invariant district characteristics or wave-specific district shocks are absorbed by the fixed effects. Third, we conduct placebo tests using outcomes that are unlikely to respond immediately to riots, such as confidence in schools or ownership of different types of assets. Finding no discontinuities in these placebo outcomes increases confidence that our results are not driven by changes in survey composition or implementation.

Table 7: Impact of Riots on Institutional Confidence

	Politicians	Military	Police	State govt.	News media	Panchayats
Panel A: Narrow window (±50 days)						
Post-riot	-0.143**	-0.150***	-0.257**	-0.186**	-0.096	-0.118*
	(0.069)	(0.047)	(0.103)	(0.087)	(0.108)	(0.067)
Observations	447	444	445	445	418	443
Clusters	31	31	31	31	31	31
Controls	Yes	Yes	Yes	Yes	Yes	Yes
District × Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Wide window (±350 days)						
Post-riot	-0.174**	-0.098**	-0.208**	-0.114	-0.097	-0.077
	(0.070)	(0.047)	(0.095)	(0.084)	(0.098)	(0.073)
Observations	1,079	1,079	1,080	1,077	1,025	1,078
Clusters	73	73	73	73	73	73
Controls	Yes	Yes	Yes	Yes	Yes	Yes
District × Wave FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports a separate regression for the listed institutional confidence outcome. The coefficient is for being interviewed after a riot in the district, estimated within the event window shown in each panel. Coefficients are multiplied by −1, so negative estimates indicate a decline in confidence after a riot. The original confidence variables are coded from 1, "a great deal of confidence," to 3, "hardly any confidence at all." For ease of interpretation, coefficients are multiplied by −1, so negative entries indicate lower confidence after a riot. Controls are Hindu, Muslim, urban residence, and household adult education. All specifications include district-by-survey-wave fixed effects. Standard errors are clustered at the PSU level and reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Placebo Outcomes: Confidence in Education and Health Services

	Gov. schools	Priv. schools	Gov. hospitals	Priv. hospitals
Panel A: Narrow window (±50 days)				
Post-riot	-0.050	0.001	-0.054	-0.037
	(0.114)	(0.052)	(0.192)	(0.099)
Observations	303	303	303	303
Clusters	22	22	22	22
Controls	Yes	Yes	Yes	Yes
District × Wave FE	Yes	Yes	Yes	Yes
Panel B: Wide window (±350 days)				
Post-riot	-0.068	0.003	-0.071	-0.062
	(0.110)	(0.049)	(0.182)	(0.098)
Observations	461	461	461	461
Clusters	34	34	34	34
Controls	Yes	Yes	Yes	Yes
District × Wave FE	Yes	Yes	Yes	Yes

Notes: Each column reports a separate regression for the listed placebo confidence outcome. The coefficient is for being interviewed after a riot in the district, estimated within the event window shown in each panel. The original confidence variables are coded from 1, "a great deal of confidence," to 3, "hardly any confidence at all." For ease of interpretation, coefficients are multiplied by -1, so negative entries indicate lower confidence after a riot. Controls are Hindu, Muslim, urban residence, and household adult education. All specifications include district-by-survey-wave fixed effects. Standard errors are clustered at the PSU level and reported in parentheses. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Placebo Tests: Household Assets and Expenditures

	Livestock	Mobile phone	House	Bicycle	Any vehicle	Motorcycle	School books
Post-riot	-0.143*	0.047	0.061	0.023	0.102	0.075	946.6
	(0.082)	(0.029)	(0.042)	(0.068)	(0.065)	(0.049)	(592.7)
Observations	449	304	447	448	448	448	448
Clusters	31	22	31	31	31	31	31
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District × Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bandwidth	±50 days	±50 days	±50 days	±50 days	±50 days	±50 days	±50 days

Notes: Each column reports a separate regression for the listed household asset or expenditure outcome. The coefficient is for being interviewed after a riot in the district. The sample is restricted to respondents interviewed within ±50 days of a riot. Controls are Hindu, Muslim, urban residence, and household adult education. All specifications include district-by-survey-wave fixed effects. Standard errors are clustered at the PSU level and reported in parentheses. These outcomes are unlikely to respond immediately to riot exposure and therefore serve as placebo checks. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A key concern with identification strategies that rely on random survey timing is that respondents interviewed after an event may differ systematically from those interviewed before it. We begin by addressing this concern using outcomes that capture household assets and consumption (see Table 9). These variables are slow-moving and determined largely by long-run economic conditions, and therefore should not respond to short-run exposure to riots. Consistent with the assumption of random interview timing, we find no evidence of systematic changes in asset ownership or consumption related outcomes around riot dates, supporting the view that survey implementation and respondent composition are not disrupted by local violence.[27]

We then turn to our main outcomes of interest: institutional confidence. These measures capture respondents' trust in key political, security, and governance institutions, including politicians, the military, the police, state government, news media, and local panchayats. Unlike economic outcomes, institutional confidence may respond rapidly to salient political shocks such as communal violence. Panel A of Table 7 reports estimates using a narrow window of ± 50 days around the riot date. We find statistically significant declines in confidence in politicians, the military, the police, state government, and panchayats. The magnitude of these effects is economically significant, corresponding to shifts of roughly 0.15–0.26 standard units in the confidence scale (this corresponds to a change of 5%-13% compared to the respective means). These institutions are either directly responsible for preventing violence, maintaining law and order, or governing at the local level, making them especially salient during episodes of unrest; these factors are extremely relevant to economic activity. By contrast, confidence in the news media does not respond significantly to riot exposure, suggesting that the erosion of trust is not uniform across all institutions.[28]

Panel B extends the event window to ± 350 days. While some effects attenuate with the wider window, several remain statistically and economically significant. In particular, confidence in politicians, the military, and the police continues to decline even over this longer horizon, indicating persistent erosion of trust in core political and security institutions. In contrast, the effects on confidence in state governments and panchayats become statistically insignificant, suggesting that these responses are more short-lived and closely tied to immediate exposure rather than long-run belief updating. This distinction is important, as persistent losses of trust in central political and security institutions may help explain why the economic effects of communal riots endure well beyond the immediate aftermath of violence.

To further rule out the possibility that these results are driven by a general increase in pessimism or by changes in survey implementation, we conduct placebo tests using confidence in schools and hospitals (Table 8). These services are less directly connected to riot dynamics and are unlikely to be immediately affected by episodes of communal violence. If post-riot respondents were simply more negative or if interviewing conditions changed after riots, we would expect to observe discontinuities in these outcomes as well. Reassuringly, we find no statistically significant effects on confidence in either government or private schools, or in public or private hospitals, in both the narrow and wide windows. The absence of effects in these placebo outcomes strengthens the interpretation that the observed declines in institutional confidence reflect a genuine response to riot exposure rather than survey artefacts or compositional changes.

Overall, the pattern of results is consistent with riots generating a marked and selective decline in institutional trust rather than a generalised shift in survey responses. The lack of effects on placebo outcomes strengthens the credibility of the identification strategy, while the persistence of effects for certain institutions suggests that riots may have lasting consequences for citizens' perceptions of the state's capacity and legitimacy.

[27] Owning livestock is marginally significant at the 10% level.

[28] Trust in banks and courts is not affected, and the results are available upon request.

7 A simple model: trust in the state, participation, and firm output

This section provides a simple partial-equilibrium framework that rationalises the main mechanism suggested by the empirical results. The model is not intended as a full structural account. Rather, it shows how a riot-induced decline in perceived state protection can reduce participation in public economic activity, thereby lowering labour input, market thickness, and local production.

Consider a district populated by a continuum of agents indexed by h. Agents include workers, entrepreneurs, and traders who may engage in public economic activity, such as commuting to work, opening a shop, or travelling to trade. Let $\theta \in [0,1]$ denote perceived *state protection*, capturing safety in public space, police responsiveness, and the credibility of public order. We interpret riots as reducing θ.

Unified participation margin. An agent chooses whether to participate in market activity. Participation yields a monetary return w, a wage for workers or expected market income for entrepreneurs and traders, plus a non-pecuniary shifter $c(\theta)$ capturing perceived safety and protection. Assume

$$c'(\theta) > 0, \tag{7.0.1}$$

so that higher perceived state protection increases the net attractiveness of participation. Each agent has an idiosyncratic outside option or disutility u_h, for example reliance on family support, home production, or avoidance of public space. The participation rule is

$$\text{Participate if} \qquad w + c(\theta) \geq u_h. \tag{7.0.2}$$

Let u_h be distributed according to a cumulative distribution function F, with density $f > 0$ over the relevant range. The share of agents who participate is therefore

$$P(w,\theta) = \Pr\{u_h \leq w + c(\theta)\} = F\big(w + c(\theta)\big). \tag{7.0.3}$$

Differentiating with respect to θ gives

$$\frac{\partial P(w,\theta)}{\partial \theta} = f\big(w + c(\theta)\big)c'(\theta) > 0. \tag{7.0.4}$$

A riot-induced decline in θ therefore reduces participation in public economic activity.

Labour supply and market thickness. We map participation into both the effective labour available to firms and the thickness of local markets. Holding w fixed in the short run, write

$$L(\theta) = \lambda_L P(w,\theta), \qquad \lambda_L > 0, \tag{7.0.5}$$

and

$$M(\theta) = \lambda_M P(w,\theta), \qquad \lambda_M > 0, \tag{7.0.6}$$

where $L(\theta)$ denotes effective labour supply and $M(\theta)$ denotes market thickness, or the intensity of public commercial interactions. Since $P_\theta(w,\theta) > 0$, it follows that

$$L'(\theta) > 0 \qquad \text{and} \qquad M'(\theta) > 0. \tag{7.0.7}$$

The interpretation is that when perceived state protection is low, fewer people are willing to travel, transact, work, or be visible in public economic spaces. As a result, both effective labour supply and local market activity decline.

Firms. A representative firm produces using capital and effective labour:

$$y(\theta) = AK^{\beta}\big(L(\theta)\big)^{\alpha}, \qquad \alpha \in (0,1), \quad \beta \in (0,1), \quad \alpha + \beta < 1. \tag{7.0.8}$$

Capital K is held fixed in the short run. By the chain rule,

$$\frac{dy(\theta)}{d\theta} = AK^{\beta}\alpha\big(L(\theta)\big)^{\alpha-1}\frac{dL(\theta)}{d\theta} > 0. \tag{7.0.9}$$

Thus, a decline in perceived state protection reduces output through lower participation and lower effective labour availability.

Discussion and link to the empirics. This simple framework rationalises our main empirical patterns. A negative shock to perceived state protection (a decline in θ following riots) reduces participation in public economic activity, $P(w,\theta)$. Because both effective labour supply and market thickness are increasing functions of participation—$L(\theta) = \lambda P(w,\theta)$ and $M(\theta) = (1-\lambda)P(w,\theta)$—the model implies simultaneous declines in (i) employment and production and (ii) markets.[29] The same mechanism can also account for contractions in banks. In particular, if banks are viewed as entrepreneurial endeavours—that is, as market participants within $M(\theta)$ that operate in public space and rely on local enforcement and safety—then a decline in θ leads to a fall in the number of banks. The same logic also helps interpret the contraction in agricultural credit. A decline in θ is likely to raise the perceived security risk of agricultural activity more than that of some other productive activities (e.g., a factory operating on enclosed premises). Banks may therefore view agricultural borrowers as disproportionately exposed to theft, looting, and disruptions to input and output markets, which increases expected default and enforcement costs. Consequently, agricultural lending tightens more strongly than credit to less security-sensitive sectors. Overall, the model highlights a single underlying channel: riots reduce confidence in the state's ability to provide protection, which withdraws workers and entrepreneurs from public economic space. This lowers labour input and market interactions, depresses output and profits, and reduces the number of banks (and reduces lending to riskier sectors such as agriculture). In this sense, the estimated declines in production, market activity, banking presence, and sectoral lending are mutually consistent manifestations of a common deterioration in perceived state protection.

Interpersonal vs. state-protection trust. An alternative interpretation of the trust channel is that riots primarily reduce *interpersonal* trust across groups, triggering boycotts and a breakdown of exchange that disproportionately affects minority communities. While such effects may well be present, we are unable to test them directly given the data constraints discussed above. Nevertheless, our evidence suggests that interpersonal channels are unlikely to be the dominant driver of the aggregate estimates. If minority-targeted exclusion were the main mechanism, the economic impact of riots should be mechanically smaller in districts with very low Muslim presence: when the Muslim share is below 10%, even an extreme scenario in which all Muslim agents are boycotted would affect only a limited fraction of local transactions and labour input, making it difficult to generate a large decline in aggregate outcomes.

Consistent with this argument, the 2SLS effect remains large and statistically significant when restricting the sample to districts with Muslim share being less than 10%. In this subsample, the coefficient on lagged riot exposure is -0.160 (see Table 11), implying a sizable contraction that is, if anything, larger than the full-sample estimate despite the small minority share.

Taken together, these results indicate that even if interpersonal trust and boycott-type mechanisms contribute, they are unlikely to account for the estimated aggregate effects on

[29] Security concerns may be more relevant for sub-markets, which is why the coefficient for that outcome variable is larger.

their own. Instead, the findings are more naturally consistent with a broader deterioration in perceived state protection that reduces participation and market interactions among the population at large.

8 Robustness

8.1 Alternative economic data

ICRISAT provides GDP data at the district level for 1993 to 2013 based on State DES.[30]

Table 10: Effect of Lagged Riot Exposure on Sectoral GDDP

	Primary sector		Secondary sector		Tertiary sector		Total GDDP	
	$\ln(1+x)$	$\Delta\ln(1+x)$	$\ln(1+x)$	$\Delta\ln(1+x)$	$\ln(1+x)$	$\Delta\ln(1+x)$	$\ln(1+x)$	$\Delta\ln(1+x)$
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.049	-0.161^{***}	-0.104^{***}	-0.034^{*}	0.003	-0.024^{**}	-0.031^{*}	-0.086^{***}
	(0.038)	(0.057)	(0.025)	(0.020)	(0.012)	(0.011)	(0.018)	(0.026)
Observations	1769	1549	1769	1549	1769	1549	2007	1699
Districts	223	203	223	203	223	203	282	233
KP rk Wald F	29.74	25.66	29.74	25.66	29.74	25.66	32.22	27.43

Notes: Each column reports a separate IV regression. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. Columns labelled $\ln(1+x)$ use outcome levels; columns labelled $\Delta\ln(1+x)$ use first differences. Outcomes are winsorised at the 1st and 99th percentiles. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Riot exposure depresses real economic activity across all major sectors, with the strongest and most robust effects appearing in the primary sector and in total district GDP. In the growth specification, a one-unit increase in lagged riot exposure reduces primary-sector growth by roughly 14.8% and total GDP growth by approximately 8.2%. The secondary and tertiary sectors also contract, though with more modest but still statistically significant declines.

The level regressions reinforce this overall pattern. The secondary sector exhibits the largest contraction: a one-unit increase in riot exposure reduces its real GDP by approximately 9.9 percent. The primary sector displays smaller, statistically insignificant effects, though the point estimate remains negative. Total GDP is significantly lower in both growth and level specifications, reinforcing our nightlight results.

8.1.1 Dropping states and years

Punjab and Jammu and Kashmir differ from the rest of the sample as they are not Hindu-majority states, and they, along with Assam, have experienced separatism. We exclude all three states as a robustness check. The coefficient is largely unchanged, as shown in Table 11.

The demolition of the Babri Mosque in Ayodhya in 1992 was a major event, and the riots that occurred in its aftermath in 1992-1993 were significant. Furthermore, the 2002 Gujarat riots were a major incident that had a widespread impact on the country. To check that our results are not being driven mainly by these major events,we drop the years 1994 and 2003 (separately and together) and run the regressions.[31]

We also try winsorising at the 5-95 percentile for the outcome variable to demonstrate the robustness that top coding and outliers are not driving our results. Also, we show that using the full sample does not change the results. Table 12 confirms that coefficients barely change with using the full sample.

[30] All values are expressed in constant 1999 rupees. Note that Jammu and Kashmir is not covered. Furthermore, the data quality has not been discussed extensively in peer-reviewed papers; hence, we use the data as a robustness check.

[31] Our nightlights data starts in 1994.

Table 11: Robustness to Dropping Selected States and Years

	(1) Drop Assam, Punjab & J&K	(2) Drop 1994	(3) Drop 2003	(4) Drop states & years	(5) Muslim share $< 10\%$
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.111^{***}	-0.080^{***}	-0.199^{***}	-0.112^{***}	-0.160^{***}
	(0.021)	(0.021)	(0.057)	(0.024)	(0.051)
District FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Other controls	Yes	Yes	Yes	Yes	Yes
KP F	61.90	44.06	18.70	39.76	17.91
Districts	376	428	428	376	243
Observations	7,520	8,132	8,132	6,768	4,567

Notes: Each column reports a separate IV regression of $\ln(1 + \text{mean calibrated night light})$ on lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. The outcome is winsorised at the 1st and 99th percentiles. Column (1) drops Assam, Punjab, and Jammu & Kashmir; column (2) drops 1994; column (3) drops 2003; column (4) drops Assam, Punjab, Jammu & Kashmir, 1994, and 2002; column (5) restricts the sample to districts with Muslim population share below 10%. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 12: Robustness to Alternative Winsorisation of the Dependent Variable

	No winsorisation	Winsorised 1–99%	Winsorised 5–95%
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.105***	-0.105***	-0.104***
	(0.023)	(0.023)	(0.022)
District FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Districts	428	428	428
Observations	8,560	8,560	8,560
KP F	56.35	56.35	56.35

Notes: Each column reports a separate IV regression using a different version of the dependent variable, $\ln(1 + \text{mean calibrated night light})$. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: * 10%, ** 5%, *** 1%.

8.1.2 Conley Standard errors

Because there is limited cross-sectional variation in the instrument and shocks may be spatially correlated across districts, we also re-estimate our specifications using Conley (spatial HAC) standard errors with clustering radii of 50, 100, and 150 km, and we take into account time lags of 1 and 5 years. Table 13 reports the corresponding results. All specifications remain statistically significant at the 5% level.

Table 13: Spatial-HAC IV Estimates

	Lag length = 1			Lag length = 5		
	(1) 50 km	(2) 100 km	(3) 150 km	(4) 50 km	(5) 100 km	(6) 150 km
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.1051***	-0.1051***	-0.1051**	-0.1051***	-0.1051***	-0.1051**
	(0.0234)	(0.0328)	(0.0423)	(0.0232)	(0.0326)	(0.0421)
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
KP F	55.53	26.48	14.91	51.46	25.51	14.60
Districts	428	428	428	428	428	428
Observations	8,560	8,560	8,560	8,560	8,560	8,560

Notes: Each column reports an IV regression with spatial-HAC standard errors. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. Columns vary the spatial cutoff used in the covariance correction and the number of temporal lags. The outcome is $\ln(1 + \text{mean calibrated night light})$, winsorised at the 1st and 99th percentiles. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: * 10%, ** 5%, *** 1%.

8.1.3 Alternative outcome variables

We show that the results are similar when we use alternative outcome variables in absolute values: maximum value of nightlight, mean calibrated nightlight, and total calibrated nightlight (controlling for the number of cells used by the satellite).

We also show that the results are negative and statistically significant at the 1% level when using the following transformations: $ln(1 + x)$ and $ln(0.01 + x)$, where x is one of the variables from the levels regressions.

Table 14: Effect of Lagged Riot Exposure on Nightlight Outcomes in Levels

Dependent variable	(1) Max light	(2) Mean light	(3) Total light
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-3.467***	-0.595***	-1,630.876
	(0.662)	(0.113)	(1,064.676)
Total cells	–	–	46.390
			(42.802)
District FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
KP F	56.345	56.345	56.325
Districts	428	428	428
Observations	8,560	8,560	8,560

Notes: Each column reports a separate IV regression. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. Outcomes are winsorised at the 1st and 99th percentiles. Column (3) additionally controls for the number of nightlight cells. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: * 10%, ** 5%, *** 1%.

Table 15: Effect of Lagged Riot Exposure under Alternative Nightlight Transformations

Dependent variable	(1) ln(1 + Total)	(2) ln(1 + Max)	(3) ln(0.01 + Mean)	(4) ln(0.01 + Max)	(5) ln(0.01 + Total)	(6) IHS(Total)	(7) IHS(Mean)	(8) IHS(Max)
$\text{RiotExposure}_{\text{sd100},i,t-1}$	-0.124***	-0.083***	-0.135***	-0.084***	-0.134***	-0.134***	-0.131***	-0.084***
	(0.042)	(0.018)	(0.037)	(0.019)	(0.040)	(0.040)	(0.030)	(0.019)
ln(1 + cells) included	Yes	No	No	No	No	No	No	No
ln(0.01 + cells) included	No	No	No	No	Yes	No	No	No
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
KP F	56.3	56.3	56.3	56.3	56.3	56.3	56.3	56.3
Districts	428	428	428	428	428	428	428	428
Observations	8,560	8,560	8,560	8,560	8,560	8,560	8,560	8,560

Notes: Each column reports a separate IV regression using an alternative transformation of the nightlight outcome. The endogenous variable is lagged riot exposure, constructed with the baseline 100 km Gaussian bandwidth. All transformed outcomes are winsorised at the 1st and 99th percentiles. Columns (1)–(5) use logarithmic transformations with alternative offsets; columns (6)–(8) use the inverse hyperbolic sine transformation. Cell-count controls are included where indicated. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. Significance levels: * 10%, ** 5%, *** 1%.

8.1.4 Contemporary independent variable

The sign and significance are consistent with riots harming the economy when using contemporary riot variables (Table 16). The magnitude is smaller, which is consistent with the discussion earlier.

Table 16: Contemporaneous Effect of Riot Exposure across Gaussian Bandwidths

	SD = 100 km	SD = 150 km	SD = 200 km	SD = 250 km	SD = 300 km
$\text{RiotExposure}_{\text{sd},i,t}$	-0.0758***	-0.0514***	-0.0417***	-0.0364***	-0.0331***
	(0.0261)	(0.0142)	(0.0098)	(0.0076)	(0.0064)
Controls	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Districts	428	428	428	428	428
Observations	8,560	8,560	8,560	8,560	8,560
KP F	36.569	47.339	56.328	66.161	77.495
AIC	**-5781.409**	-5641.214	-5325.178	-5070.154	-4884.749
BIC	**-5753.189**	-5612.994	-5296.959	-5041.935	-4856.530

Notes: Each column reports a separate IV regression using contemporaneous riot exposure constructed with the Gaussian bandwidth shown in the column heading. The outcome is $\ln(1 + \text{mean calibrated night light})$, winsorised at the 1st and 99th percentiles. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. AIC and BIC are reported for model comparison; bold entries mark the lowest value in each row. Significance levels: * 10%, ** 5%, *** 1%.

8.1.5 Panel unit root

Although our nightlights panel spans up to 20 yearly observations (1994–2013), we prudently test for unit roots. We use the Im–Pesaran–Shin (IPS) test Im et al. (2003), which is designed for panels with fixed T and $N \to \infty$ and allows heterogeneous autoregressive parameters across units. Following Levin et al. (2002), we also remove cross-sectional means to mitigate cross sectional dependence and allow for time trends. Results are robust to omitting the demeaning step and/or the trend. Across all specifications in Table 17, we reject the null that all panels contain a unit root.[32]

[32] In the IPS framework, the alternative is that a non-zero fraction of panels are stationary as $N \to \infty$.

Table 17: Im–Pesaran–Shin panel unit root tests for y_{it} (1994–2013)

	(1) Trend + demean	(2) Trend	(3) No trend + demean	(4) No trend
Panels (N)	428	428	428	428
Periods (T)	20	20	20	20
Panel means included	Yes	Yes	Yes	Yes
Time trend included	Yes	Yes	No	No
Cross-sectional means removed	Yes	No	Yes	No
ADF lags	0	0	0	0
t-bar	-3.3394	-3.3093	-2.1143	-1.9617
$\tilde{t}$-bar	-2.5804	-2.5820	-1.8371	-1.7469
$Z(\tilde{t}\text{-bar})$	-31.8735	-31.9162	-11.8886	-9.4618
p-value for Z	0.0000	0.0000	0.0000	0.0000
Fixed-N exact critical values:				
1%	-2.380	-2.380	-1.730	-1.730
5%	-2.320	-2.320	-1.670	-1.670
10%	-2.280	-2.280	-1.640	-1.640

Notes: IPS tests with panel-specific AR parameters. "Demean" indicates cross-sectional means removed. All four specifications reject the null that all panels contain a unit root at conventional levels.

8.1.6 Growth regressions

We check how the results hold for growth instead of levels. Figure 2 shows that the negative effect of riot exposure on economic growth is concentrated in the 2000s. When estimating the IV specification on rolling time windows, the coefficients are consistently negative and statistically significant for all samples ending before 2010, with point estimates between –0.30 and –0.05. After 2010, the magnitudes attenuate and lose statistical significance. This pattern suggests that riot shocks had larger growth effects in the early part of the sample, while the later years show weaker responses, potentially reflecting changes in economic structure, measurement, or the declining frequency and intensity of riots.

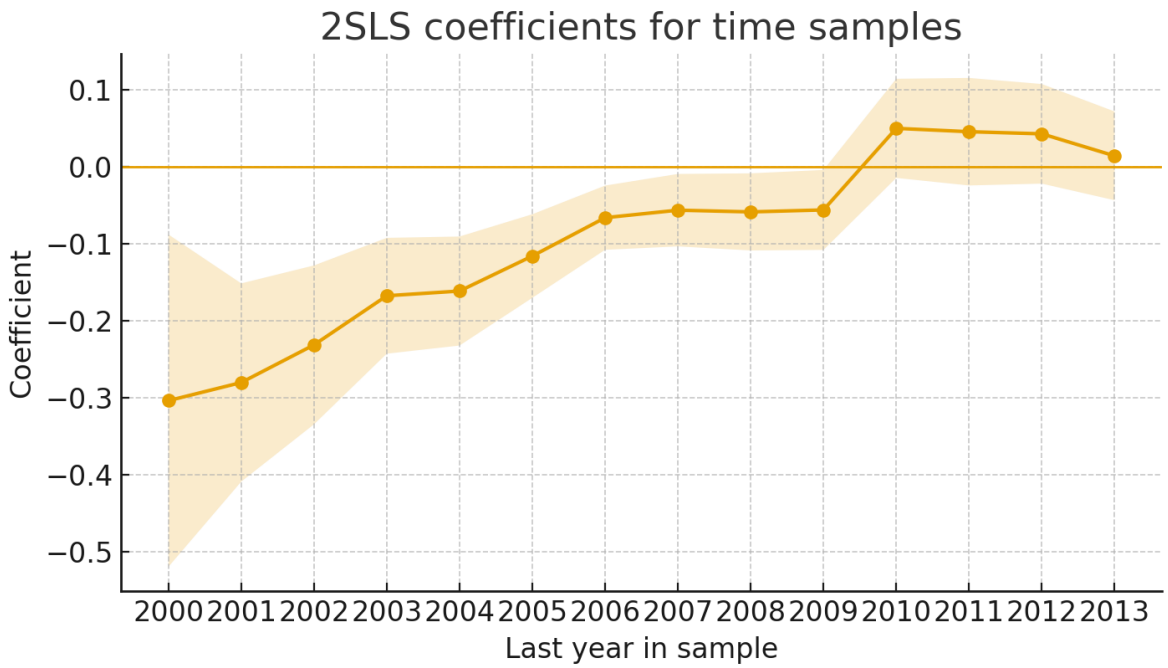


Figure 2: IV coefficient on lagged riot exposure across rolling time windows (with controls; Winsorised 1–99%).

Figure 1 shows that the intensity of religious riots declines substantially after the early 2000s, with particularly high activity around 1999–2004 and markedly lower levels thereafter. This pattern helps explain why the IV effects on nightlight growth are large and significant when the sample is restricted to years before 2010, but attenuate in later periods. When riot shocks are less frequent and less intense, their year-to-year impact on local growth becomes correspondingly harder to detect. The level regressions nevertheless capture the cumulative effects of earlier intense periods, whereas the growth regressions depend on contemporaneous variation. The fall in riot activity is therefore fully consistent with the observed attenuation of growth effects in the later sample years.

9 Conclusion

This paper provides the first district-level estimates of the economic consequences of communal riots in India. Using a newly extended dataset on Hindu–Muslim riots from 1994 to 2013, combined with nightlights, manufacturing microdata, banking and credit statistics, and market infrastructure records, we show that riot exposure generates sizable and persistent local economic losses. Using quasi-random variation in the coincidence of major Hindu festivals with Fridays, we construct a spatially weighted instrument that captures exogenous shifts in riot risk. Across all specifications, a one-unit increase in lagged riot exposure reduces nightlight luminosity by roughly 0.105 log points, implying a GDP decline of around 0.45 percent, after taking into account the under-reporting of riots.

Our results reveal a consistent set of mechanisms through which riots depress economic activity. Manufacturing firms contract immediately in terms of labour and output. Financial infrastructure responds strongly: districts exposed to greater riot intensity experience a substantial withdrawal of bank branches, sharp reductions in agricultural lending, and a sharp contraction in local market structure. Complementing these findings, we show that riots also lead to declines in confidence in state institutions, such as politicians, the police, and local gov-

ernments, suggesting an erosion of the institutional trust environment that supports economic coordination.

More broadly, the findings highlight that communal riots impose economic costs not only through direct disruptions to production, but by weakening the institutional and relational foundations of local economies. By documenting how violence propagates spatially and affects multiple layers of economic activity (firms, markets, finance, and institutional trust), this paper provides a unified account of how episodic religious conflict can generate persistent economic losses. The results highlight the importance of institutional confidence as a transmission channel linking social violence to long-run economic outcomes, and suggest that the economic consequences of communal riots extend well beyond their immediate physical and temporal footprint.

References

Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American Economic Review*, 93(1):113–132, 2003.

Alex S. Abrahams. Hard traveling: Unemployment and road infrastructure in the shadow of political conflict. *Political Science Research and Methods*, 10(3):545–566, 2022. doi: 10.1017/psrm.2021.8. URL https://doi.org/10.1017/psrm.2021.8.

Kaustabh Adhya. Conflict, economic activities and the internet: Disentangling the (world wide) web. Manuscript, January 2026.

Farzana Afridi, Taryn Dinkelman, and Kanika Mahajan. Why are fewer married women joining the work force in india? a decomposition analysis over two decades. *American Economic Review*, 108(1):1–38, 2018. doi: 10.1257/aer.20130472.

Sam Asher, Tobias Lunt, Ryu Matsuura, and Paul Novosad. Development research at high geographic resolution: an analysis of night-lights, firms, and poverty in india using the shrug open data platform. *The World Bank Economic Review*, 35(4), 2021.

Thushyanthan Baskaran, Brian Min, and Yogesh Uppal. Election cycles and electricity provision: Evidence from a quasi-experiment with indian special elections. *Journal of Public Economics*, 126:64–73, June 2015. doi: 10.1016/j.jpubeco.2015.03.011. URL https://doi.org/10.1016/j.jpubeco.2015.03.011.

Nathaniel Baum-Snow, Loren Brandt, J. Vernon Henderson, Matthew A. Turner, and Qinghua Zhang. Roads, railroads, and decentralization of chinese cities. *The Review of Economics and Statistics*, 99(3):435–448, July 2017. doi: 10.1162/REST_a_00660. URL https://doi.org/10.1162/REST_a_00660.

Thorsten Beck, Ross Levine, and Norman Loayza. Finance and the sources of growth. *Journal of Financial Economics*, 58(1–2):261–300, 2000.

Frank Bickenbach, Eckhardt Bode, Peter Nunnenkamp, and Mareike Söder. Night lights and regional GDP. *Review of World Economics*, 152(2):425–447, 2016. doi: 10.1007/s10290-016-0246-0.

Hans P. Binswanger and Shahidur R. Khandker. The impact of formal finance on the rural economy of india. *Journal of Development Studies*, 32(2):234–262, 1995. doi: 10.1080/00220389508422413.

Christopher Blattman and Edward Miguel. Civil war. *Journal of Economic Literature*, 48 (1):3–57, 2010. ISSN 0022-0515. doi: 10.1257/jel.48.1.3. URL https://www.aeaweb.org/articles?id=10.1257/jel.48.1.3.

Anjali Thomas Bohlken and Ernest John Sergenti. Economic growth and ethnic violence: An empirical investigation of hindu–muslim riots in india. *Journal of Peace Research*, 47(5):589–600, 2010. URL https://www.jstor.org/stable/20798928.

Paul R. Brass. *The Production of Hindu–Muslim Violence in Contemporary India*. University of Washington Press, Seattle, 2003.

Robin Burgess and Rohini Pande. Do rural banks matter? evidence from the indian social banking experiment. *American Economic Review*, 95(3):780–795, 2005.

Business Standard. Riots paralyse industry in muzaffarnagar. *Business Standard*, September 2013. URL https://www.business-standard.com/article/current-affairs/riots-paralyse-industry-in-muzaffarnagar-113091000841_1.html. Accessed: 2026-01-17.

A. Colin Cameron and Douglas L. Miller. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372, March 2015. doi: 10.3368/jhr.50.2.317.

Shoumitro Chatterjee, Kala Krishna, Kalyani Padmakumar, and Yingyan Zhao. No country for dying firms: Evidence from india. Unpublished manuscript, April 2025.

Xi Chen and William D. Nordhaus. Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594, 2011. doi: 10.1073/pnas.1017031108. URL https://doi.org/10.1073/pnas.1017031108.

William J. Collins and Robert A. Margo. The labor market effects of the 1960s riots. NBER Working Paper 10243, National Bureau of Economic Research, 2004a. URL https://www.nber.org/papers/w10243.

William J. Collins and Robert A. Margo. The economic aftermath of the 1960s riots in american cities: Evidence from property values. Working Paper 0410, Vanderbilt University Department of Economics, 2004b. URL https://ideas.repec.org/p/van/wpaper/0410.html.

William J. Collins and Robert A. Margo. The economic aftermath of the 1960s riots in american cities: Evidence from property values. *The Journal of Economic History*, 67(4):849–883, December 2007. doi: 10.1017/S0022050707000423. URL https://doi.org/10.1017/S0022050707000423.

Timothy G. Conley. Gmm estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45, 1999.

Timothy G. Conley. Spatial econometrics. In Steven N. Durlauf and Lawrence E. Blume, editors, *Microeconometrics*, pages 303–313. Palgrave Macmillan, London, 2008.

Emilio Depetris-Chauvin, Ruben Durante, and Filipe Campante. Building nations through shared experiences: Evidence from african football. *American Economic Review*, 110(5):1572–1602, 2020. doi: 10.1257/aer.20180906.

Amrita Dhillon, Pramila Krishnan, Manasa Patnam, and Carlo Perroni. Electoral accountability and the natural resource curse: Theory and evidence from india. CEPR Discussion Paper DP11377, Centre for Economic Policy Research (CEPR), July 2016. URL https://ssrn.com/abstract=2807788. Available at SSRN: https://ssrn.com/abstract=2807788.

Denise DiPasquale and Edward L. Glaeser. The Los Angeles riot and the economics of urban unrest. *Journal of Urban Economics*, 43(1):52–78, 1998.

C. N. H. Doll, J.-P. Muller, and J. G. Morley. Mapping regional economic activity from nighttime light satellite imagery. *Ecological Economics*, 57:75–92, 2006. doi: 10.1016/j.ecolecon.2005.03.007. URL https://doi.org/10.1016/j.ecolecon.2005.03.007.

William Easterly and Ross Levine. Africa's growth tragedy: Policies and ethnic divisions. *Quarterly Journal of Economics*, 112(4):1203–1250, 1997.

Christopher D. Elvidge, Kimberly E. Baugh, Erik A. Kihn, Herbert W. Kroehl, Ethan R. Davis, and C. W. Davis. Mapping city lights with nighttime data from the DMSP operational linescan system. *Photogrammetric Engineering & Remote Sensing*, 63(6):727–734, 1997.

Christopher D. Elvidge, Kimberly E. Baugh, J. B. Dietz, T. Bland, Paul C. Sutton, and Herbert W. Kroehl. Night-time lights of the world: 1994–1995. *ISPRS Journal of Photogrammetry and Remote Sensing*, 56(2):81–99, 2001.

James D. Fearon. Ethnic mobilization and ethnic violence. In Barry R. Weingast and Donald A. Wittman, editors, *The Oxford Handbook of Political Economy*, pages 852–868. Oxford University Press, Oxford, 2008.

Erica Field, Matthew Levinson, Rohini Pande, and Sujata Visaria. Segregation, rent control, and riots: The economics of religious conflict in an Indian city. *American Economic Review: Papers and Proceedings*, 98(2):505–510, 2008.

Raymond Fisman, Arkodipta Sarkar, Janis Skrastins, and Vikrant Vig. Experience of communal conflicts and intergroup lending. *Journal of Political Economy*, 128(9):3346–3375, 2020. doi: 10.1086/708856.

Devaki Ghose and Divya Pandey. The effects of communal violence on women's marital outcomes. *Demography*, 62(3):1087–1108, 2025. doi: 10.1215/00703370-11995654. URL https://doi.org/10.1215/00703370-11995654.

Luigi Guiso, Paola Sapienza, and Luigi Zingales. Does local financial development matter? *Quarterly Journal of Economics*, 119(3):929–969, 2004.

Shiqi Guo and Jiafu An. Does terrorism make people pessimistic? evidence from a natural experiment. *Journal of Development Economics*, 155(C):None, None 2022. doi: 10.1016/j.jdeveco.2022.102817. URL https://ideas.repec.org/a/eee/deveco/v155y2022ics0304387822000013.html.

Mariaflavia Harari. Cities in bad shape: Urban geometry in india. *American Economic Review*, 110(8):2377–2421, August 2020. doi: 10.1257/aer.20171673. URL https://doi.org/10.1257/aer.20171673.

J. Vernon Henderson, Adam Storeygard, and David N. Weil. A Bright Idea for Mesuring Economic Growth. *American Economic Review*, 2011.

Roland Hodler and Paul A. Raschky. Regional favoritism. *The Quarterly Journal of Economics*, 129(2):995–1033, 2014. doi: 10.1093/qje/qju004. URL https://doi.org/10.1093/qje/qju004.

Donald L. Horowitz. *Ethnic Groups in Conflict*. University of California Press, Berkeley, 1985.

Kyung So Im, M. Hashem Pesaran, and Yongcheol Shin. Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, 115(1):53–74, 2003. doi: 10.1016/S0304-4076(03)00092-7. URL https://doi.org/10.1016/S0304-4076(03)00092-7.

Indian Express. The economy of fear. *The Indian Express*, September 2013. URL https://indianexpress.com/article/india/india-others/the-economy-of-fear/. Accessed: January 2026.

Sriya Iyer. *The Economics of Religion in India*. The Belknap Press of Harvard University Press, Cambridge, MA, 2018. ISBN 9780674979642.

Sriya Iyer. The economics of religious conflicts. In Carla Bagnoli and Vincenzo Pacillo, editors, *Religious Authorities and Practices of Conflict Resolution*, chapter 6. Routledge, London, 2025. doi: 10.4324/9781003599487-8. URL https://doi.org/10.4324/9781003599487-8.

Sriya Iyer and Anand Shrivastava. Religious riots and electoral politics in india. *Journal of Development Economics*, 131:104–122, 2018. doi: 10.1016/j.jdeveco.2017.11.003. URL https://doi.org/10.1016/j.jdeveco.2017.11.003.

Christophe Jaffrelot. *Religion, caste, and politics in India*. Primus Books, 2010.

Saumitra Jha. "unfinished business": Historic complementarities, political competition, and ethnic violence in gujarat. *Journal of Economic Behavior & Organization*, 104:18–36, 2014.

Christopher Ksoll, Rocco Macchiavello, and Ameet Morjaria. Electoral violence and supply chain disruptions in kenya's floriculture industry. *The Review of Economics and Statistics*, 105(6):1335–1351, 2023. doi: 10.1162/rest_a_01185. URL https://doi.org/10.1162/rest_a_01185.

Rafael La Porta, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny. The quality of government. *Journal of Law, Economics, and Organization*, 15(1):222–279, 1999.

Andrew Levin, Chien-Fu Lin, and Chia-Shang James Chu. Unit root tests in panel data: Asymptotic and finite-sample properties. *Journal of Econometrics*, 108(1):1–24, 2002. doi: 10.1016/S0304-4076(01)00098-7. URL https://doi.org/10.1016/S0304-4076(01)00098-7.

Leslie A. Martin, Shanthi Nataraj, and Ann E. Harrison. In with the big, out with the small: Removing small-scale reservations in india. *American Economic Review*, 107(2):354–386, February 2017. doi: 10.1257/aer.20141335. URL https://doi.org/10.1257/aer.20141335.

Paolo Mauro. Corruption and growth. *Quarterly Journal of Economics*, 110(3):681–712, 1995.

Stelios Michalopoulos and Elias Papaioannou. Pre-colonial ethnic institutions and contemporary african development. *Econometrica*, 81(1):113–152, 2013. doi: 10.3982/ECTA9613. URL https://doi.org/10.3982/ECTA9613.

Stelios Michalopoulos and Elias Papaioannou. National institutions and subnational development in africa. *The Quarterly Journal of Economics*, 129(1):151–213, 2014. doi: 10.1093/qje/qjt029. URL https://doi.org/10.1093/qje/qjt029.

Brian Min. *Power and the Vote: Elections and Electricity in the Developing World*. Cambridge University Press, Cambridge, 2015.

Brian Min, Kwawu Mensan Gaba, Ousmane Fall Sarr, and Alassane Agalassou. Detection of rural electrification in africa using dmsp-ols night lights imagery. *International Journal of Remote Sensing*, 34(22):8118–8141, 2013. doi: 10.1080/01431161.2013.833358. URL https://doi.org/10.1080/01431161.2013.833358.

Anirban Mitra and Debraj Ray. Implications of an economic theory of conflict: Hindu-Muslim violence in India. *Journal of Political Economy*, 122(4):719–765, 2014.

Nathan Nunn and Leonard Wantchekon. The slave trade and the origins of mistrust in africa. *American Economic Review*, 101(7):3221–3252, 2011. ISSN 0002-8282. doi: 10.1257/aer.101.7.3221. URL https://www.aeaweb.org/articles?id=10.1257/aer.101.7.3221.

Maxim L. Pinkovskiy and Xavier Sala-i Martin. Lights, camera, income! illuminating the national accounts–household surveys debate. *The Quarterly Journal of Economics*, 131(2):579–631, 2016. doi: 10.1093/qje/qjw003. URL https://doi.org/10.1093/qje/qjw003.

Dominic Rohner, Mathias Thoenig, and Fabrizio Zilibotti. War signals: A theory of trade, trust, and conflict. *The Review of Economic Studies*, 80(3):1114–1147, 2013. doi: 10.1093/restud/rds049. URL https://www.jstor.org/stable/43551458.

Adam Storeygard. Farther on down the road: Transport costs, trade, and urban growth in sub-saharan africa. *The Review of Economic Studies*, 83(3):1263–1295, 2016. doi: 10.1093/restud/rdw020. URL https://doi.org/10.1093/restud/rdw020.

Sander van Noort and Tanushree Goyal. Is ethnic violence self-perpetuating? quasi-experimental evidence from hindu–muslim riots in india. *Journal of Conflict Resolution*, 0(0):1–29, 2025. doi: 10.1177/00220027251383563.

Ashutosh Varshney. *Ethnic Conflict and Civic Life: Hindus and Muslims in India*. Yale University Press, New Haven, 2002. URL https://www.jstor.org/stable/j.ctt1nq5hn.

Ashutosh Varshney and Steven Wilkinson. Varshney–wilkinson dataset on hindu–muslim violence in india, 1950–1995, version 2. https://doi.org/10.3886/ICPSR04342.v1, February 2006. URL https://doi.org/10.3886/ICPSR04342.v1. Inter-university Consortium for Political and Social Research [distributor].

David Wilkinson. Focus group research. In Clive Seale, editor, *Qualitative Research: Theory, Method and Practice*, pages 177–199. SAGE Publications, London, 2004.

Chris Wilson. The ethnic origins of religious conflict in North Maluku province, Indonesia, 1999–2000. *Indonesia*, 79:69–91, 2005.

A Appendix: Rings

To assess whether our baseline results are driven by any particular spatial band, we re-estimate the main specification using ring-specific measures of riot exposure that restrict attention to riots occurring within predefined distance intervals. We estimate:

$$y_{it} \;=\; \alpha_i \;+\; \beta\, \mathrm{R}^{\mathrm{r}}_{i,t-1} \;+\; \gamma_t \;+\; \varepsilon_{it}, \tag{A.0.1}$$

where $y_{it} \equiv \ln\!\big(1 + \text{mean calibrated night light}_{it}\big)$ for district i in year t; α_i and γ_t denote district and year fixed effects. The exposure measures R^{r}_{it} ($r \in \{\text{near}, \text{mid}, \text{far}\}$) are ring–specific sums of distance–decayed riot intensities, defined as

$$\mathrm{R}^{(a,b)}_{it} \;=\; \sum_{j \in \mathcal{J}_t} \phi_{(a,b)}\big(d_{ij}\big), \qquad \phi_{(a,b)}(d) \;=\; \mathbf{1}\{a \le d < b\}\, \exp\!\left[-\frac{(d-a)^2}{2\sigma^2}\right],$$

where d_{ij} is the great–circle distance (in km) from district i's centroid to riot j's location in year t, and $[a, b)$ denotes the inner and outer bounds of the ring. The truncated Gaussian kernel $\phi_{(a,b)}(d)$ ensures that only riots within a given ring affect $\mathrm{R}^{(a,b)}_{it}$: weight is highest at the inner edge ($d = a$) and decays smoothly at the outer boundary ($d = b$).

To address potential endogeneity of riot occurrence, each exposure is instrumented with an analogous ring–specific measure based on exogenous festival timing shocks,

$$\mathrm{IV}^{(a,b)}_{it} \;=\; \sum_{k} \phi_{(a,b)}\big(\mathrm{dist}(x_i, x_k)\big)\, \mathrm{Festival}_{kt},$$

where $\mathrm{Festival}_{kt}$ equals the indicator of major religious festivals falling on Fridays in district k and year t. This instrument propagates the exogenous spatial–temporal variation in riot risk through the same truncated Gaussian kernel, ensuring that the IV mirrors the spatial structure of the endogenous exposure.

In this framework, the "near", "mid", and "far" rings capture respectively the local, intermediate, and distant effect of riots on economic activity.

Bellow are different specifications (lagged/not lagged, using ring specific IV); We use rings of 50 km in (1)-(3) and rings of 100 km in (4)-(6).

Across all specifications, the estimated effects of riot exposure within the near, mid, and far rings are consistently negative and of similar magnitude. In particular, we find no evidence of positive spillovers from riots occurring in more distant districts, supporting the interpretation that riot exposure in surrounding areas does not offset local economic losses and reinforcing the spatial robustness of our main specification.[33]

[33] The 0–50 km ring suffers from weak instruments in some specifications; however, in the preferred lagged specification the Anderson–Rubin p-values remain below conventional 5% thresholds, indicating significance even under weak identification.

Table 18: Effect of Riot Exposure by Distance Band

	(1)	(2)	(3)	(4)	(5)	(6)
Distance band (km)	0–50	50–100	100–150	0–100	100–200	200–300
Panel A: Lagged riots						
Coefficient	-0.3806	-0.1904***	-0.1626***	-0.1526***	-0.1056***	-0.1254***
	(0.3166)	(0.0551)	(0.0350)	(0.0454)	(0.0224)	(0.0279)
KP F	2.04	29.70	49.40	29.53	46.16	32.72
AR p-value	0.041	0.000	0.000	0.000	0.000	0.000
Panel B: Contemporaneous riots						
Coefficient	-0.0791	-0.1345*	-0.1136***	-0.0929	-0.0689***	-0.0812***
	(0.3469)	(0.0701)	(0.0338)	(0.0583)	(0.0202)	(0.0219)
KP F	1.10	15.62	38.82	17.24	36.92	23.64
AR p-value	0.815	0.010	0.000	0.062	0.000	0.000
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,560	8,560	8,560	8,560	8,560	8,560
Districts	428	428	428	428	428	428

Notes: Each column reports a separate IV regression using riot exposure from the distance band shown in the column heading. Panel A uses lagged riot exposure; Panel B uses contemporaneous riot exposure. The outcome is $\ln(1 + \text{mean calibrated night light})$, winsorised at the 1st and 99th percentiles. All specifications include district and year fixed effects and control for BJP incumbency, Muslim population share, and male population share. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. AR p-values are from Anderson–Rubin tests. Significance levels: * 10%, ** 5%, *** 1%.

B Alternative explanatory variables

B.1 Binary variable

We also estimate a specification using a binary measure equal to 1 if at least one riot occurs in district i in year t:

$$y_{it} \;=\; \alpha_i \;+\; \beta\,\text{riots happened}_{i,t-1} \;+\; \gamma' X_{it} + \gamma_t \;+\; \varepsilon_{it}, \tag{B.1.1}$$

While the coefficient is negative and the Anderson–Rubin test indicates significance, the first-stage F-statistic is very low, implying that the point estimate is unreliable under weak instruments. Using a specification without year fixed effects but with a linear time trend yields a negative and significant coefficient with an F-statistic above 10 but below the Stock–Yogo threshold of 16.38. Because this specification omits year fixed effects, which are empirically important in our setting, the results should be interpreted with caution.

More fundamentally, this binary specification ignores important spatial spillovers: riots in nearby districts are treated as irrelevant even though they plausibly affect local economies, leading to likely attenuation bias.

B.1.1 Distance to riot

We next estimate:

$$y_{it} \;=\; \alpha_i \;+\; \beta\,\text{distance to closest riot}'_{i,t-1} \;+\; \gamma' X_{it} + \gamma_t \;+\; \varepsilon_{it}, \tag{B.1.2}$$

where the exposure variable is the centroid-to-nearest-riot distance (in km).

The results (reported in Table 19, Columns 3 and 4) show that greater distance from the nearest riot is associated with higher night-light intensity, consistent with economic disruption being stronger closer to violence. The specification with year fixed effects suffers from weak instruments, yet the Anderson–Rubin test remains significant at the 1% level. The specification with a linear time trend has strong first-stage strength and yields a positive and highly significant coefficient, but it once again suffers from the omission of year fixed effects.

B.1.2 Closest riot exposure

We also estimate a specification using the kernel-weighted exposure to the closest riot:

$$y_{it} \;=\; \alpha_i \;+\; \beta\, \text{Riots}'_{i,t-1} \;+\; \gamma' X_{it} + \gamma_t \;+\; \varepsilon_{it}, \tag{B.1.3}$$

with d_{ijt} the distance to the nearest riot.

As reported in Table 19 (columns 5–6), the specification with year fixed effects again exhibits weak instruments but yields a significant AR p-value. The version with a linear trend has strong instrument strength but omits year fixed effects. Both specifications also ignore the extensive margin—multiple riots versus a single distant one—which limits their economic interpretability.

B.1.3 Number of riots

Finally, we examine the intensive margin:

$$y_{it} \;=\; \alpha_i \;+\; \beta\, \text{n_riots}'_{i,t-1} \;+\; \gamma' X_{it} + \gamma_t \;+\; \varepsilon_{it}, \tag{B.1.4}$$

where n_riot_{it} counts the number of riots in district i in year t.

As shown in Table 19 (column 7), the specification with year fixed effects again suffers from weak identification but remains significant according to the Anderson–Rubin test. The specification with a linear time trend (column 6) shows strong first-stage strength but is undermined by the omission of year fixed effects. Moreover, these specifications ignore the spatial distribution of riots: a district with many small nearby riots may be more economically affected than one with a single riot far from its centroid, a distinction that the count variable cannot capture.

Table 19: Summary of IV Estimates Across Key Specifications

	(1) Riot lagged	(2) Riot lagged	(3) Distance	(4) Distance	(5) Lagged intensity	(6) Lagged intensity	(7) Lagged events
Coefficient	−2.723	−3.181**	0.001844*	0.000457***	−1.169*	−0.605***	−0.695**
Standard error	(2.043)	(1.056)	(0.001030)	(0.000051)	(0.633)	(0.095)	(0.312)
KP F	2.19	10.61	3.86	214.10	5.29	77.54	9.13
AR p-value	0.0002	0.0000	0.0002	0.0000	0.0002	0.0000	0.0002
Observations	8,560	8,560	8,560	8,560	8,560	8,560	8,560
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes
Time trend	No	Yes	No	Yes	No	Yes	No

Notes: Each column reports a separate IV regression using an alternative measure of riot exposure. The row labelled "Coefficient" reports the estimated coefficient on the endogenous riot measure named in the column heading. The excluded instrument is the corresponding lagged festival-based shock. All specifications control for BJP incumbency, Muslim population share, and male population share. District and year fixed effects, or district fixed effects with a linear time trend, are included as indicated. Standard errors are clustered at the district level and reported in parentheses. KP F is the Kleibergen–Paap first-stage F statistic. AR p-values are from Anderson–Rubin tests. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Across all alternative specifications, the estimated effects remain consistently negative. However, each alternative measure has substantive limitations: several suffer from weak instruments, others rely on specifications that omit year fixed effects, and many fail to capture the spatial spillovers that are central to understanding how violence propagates across local economies.

Taken together, these issues highlight the value of our preferred approach, which embeds a continuous spatial distance decay structure and leverages strong instrumental variation to identify the impact of riots more reliably, whilst allowing for district fixed effects.

C Back-of-the-envelope: aggregate GDP effects

Let i index districts and t years. Our preferred estimates imply that a one-unit increase in lagged riot exposure reduces district GDP by approximately 0.45%. We treat this as a semi-elasticity in logs and write the implied effect on district GDP as

$$\Delta \ln(GDP_{it}) \approx -0.0045 \cdot \text{RiotExposure}_{i,t-1}. \tag{C.0.1}$$

The coefficient 0.0045 is obtained by translating our estimated effect on nightlights into GDP units using an external elasticity of GDP with respect to nightlights.

For convenience, we define the implied district-year *log GDP loss magnitude* as

$$\ell_{it} \equiv 0.0045 \cdot \text{RiotExposure}_{i,t-1} \tag{C.0.2}$$

so that the predicted change in district GDP satisfies $\Delta \ln(GDP_{it}) \approx -\ell_{it}$.

To aggregate local losses to the national level, we weight districts by their share of national economic activity. Because annual district GDP is unavailable, we proxy economic size using winsorised mean calibrated nightlight intensity in levels, denoted NL_{it}. We define year-specific weights

$$\omega_{it} \equiv \frac{NL_{it}}{\sum_j NL_{jt}}, \qquad \sum_i \omega_{it} = 1 \ \ \forall t, \tag{C.0.3}$$

and compute the implied national log GDP loss magnitude in year t as the weighted average

$$L_t \equiv \sum_i \omega_{it}\, \ell_{it}. \tag{C.0.4}$$

Exponentiating the implied log loss allows us to recover percentage changes in GDP. We report the *percentage GDP loss* in year t as

$$\%\Delta GDP_t^{loss} \equiv 100 \left(e^{-L_t} - 1\right), \tag{C.0.5}$$

which is negative. Equivalently, we can express the same calculation as the *counterfactual gain* in GDP in a no-riots scenario:

$$\%\Delta GDP_t^{no\,riots} \equiv 100 \left(e^{L_t} - 1\right), \tag{C.0.6}$$

which is positive. The two expressions differ only in the choice of baseline (actual versus counterfactual); when L_t is small, they are approximately equal in magnitude, since $e^{L_t} - 1 \approx L_t$ and $e^{-L_t} - 1 \approx -L_t$.

Assumptions and interpretation. This back-of-the-envelope aggregation relies on: (i) district GDP shares being proportional to the nightlights-based weights ω_{it} in (C.0.3); (ii) applying the same semi-elasticity uniformly across districts and years; and (iii) a partial-equilibrium interpretation that abstracts from general-equilibrium adjustments, such as reallocation of activity across districts or price responses.[34]

Finally, our riot exposure measure incorporates spatial spillovers by construction: riots in one location contribute to $\text{RiotExposure}_{i,t-1}$ in multiple nearby districts.[35]

[34] For example, if violence shifts production or investment from more exposed districts to less exposed districts, national losses could be smaller than the sum of local losses; conversely, if spillovers operate through trade, credit, or migration networks, national losses could be larger than what is captured by within-district outcomes alone.

[35] Because exposure is distance-weighted, a single riot can increase predicted losses in several districts. Aggregating ℓ_{it} across districts therefore counts both direct and spillover exposure and may attribute the influence of the same riot to multiple districts. For this reason, L_t should be interpreted as an upper-bound measure of the aggregate loss implied by total local exposure, rather than a precise estimate of national GDP losses.

Year-by-year aggregate effects. Table 20 reports the year-by-year back-of-the-envelope aggregate GDP losses computed using the procedure described above.

Table 20: Back-of-the-envelope aggregate GDP effects by year

Year	GDP loss (%)	GDP gain in no-riots counterfactual (%)
1994	-0.39	0.39
1995	-0.11	0.11
1996	-0.12	0.12
1997	-0.08	0.08
1998	-0.07	0.07
1999	-0.18	0.18
2000	-0.15	0.15
2001	-0.33	0.33
2002	-0.41	0.42
2003	-0.77	0.78
2004	-0.58	0.58
2005	-0.26	0.26
2006	-0.13	0.13
2007	-0.21	0.21
2008	-0.06	0.06
2009	-0.46	0.46
2010	-0.06	0.06
2011	-0.13	0.13
2012	-0.10	0.10
2013	-0.15	0.15

Notes: The table reports year-by-year back-of-the-envelope aggregate GDP effects obtained by aggregating district-level log GDP loss magnitudes using winsorised nightlights-based economic size weights. GDP loss is computed as $100(e^{-L_t} - 1)$ and GDP gain as $100(e^{L_t} - 1)$, where $L_t = \sum_i \omega_{it}\ell_{it}$ and $\ell_{it} = 0.0045 \cdot \text{RiotExposure}_{i,t-1}$.

D Back-of-the-envelope: State–Year GDP Losses from Riots

This appendix describes the construction of state–year GDP loss measures used to benchmark the magnitude of riot-induced economic disruptions and to contextualize our estimates relative to existing work on conflict and economic activity.

D.1 Construction of State–Year GDP Loss Measures

Let i index districts, s states, and t years. Our baseline instrumental-variables estimates imply that a one-unit increase in lagged riot exposure reduces log GDP by 0.0045. Because district-level GDP is not observed annually, we proxy economic activity using night-time lights and aggregate district-level losses to the state–year level using nightlights-based weights.

For each district–year, we define the implied log GDP loss as

$$\Delta \ln GDP_{i,t} = 0.0045 \times \text{RiotExposure}_{i,t-1}. \tag{D.1.1}$$

To aggregate these losses to the state–year level, we construct a nightlights-weighted average exposure:

$$L_{s,t} = \sum_{i \in s} w_{i,s,t} \cdot \text{RiotExposure}_{i,t-1}, \qquad w_{i,s,t} = \frac{\text{Nightlights}_{i,t}}{\sum_{j \in s} \text{Nightlights}_{j,t}}. \tag{D.1.2}$$

This weighting scheme assigns greater importance to economically larger districts within each state-year, ensuring that the resulting object reflects exposure of economic activity.

The implied state–year log GDP loss is then given by

$$\Delta \ln GDP_{s,t} = 0.0045 \times L_{s,t}. \tag{D.1.3}$$

We convert log losses into percentage terms using the exact transformation

$$\Delta GDP_{s,t}(\%) = 100 \times \left(e^{-\Delta \ln GDP_{s,t}} - 1 \right). \tag{D.1.4}$$

D.2 Distribution of Implied State–Year GDP Losses

Table 21 reports summary statistics for the distribution of implied GDP losses across all state–years in the sample. The distribution is highly skewed. While the median state–year loss is modest, the upper tail features large implied declines in economic activity.

In particular, the median implied loss is approximately 0.06%, while the mean loss is about 0.23%, reflecting the influence of a small number of highly exposed state–years. A substantial mass of state–years exhibits losses close to zero, consistent with the episodic and spatially concentrated nature of communal violence.

D.3 Most and Least Affected State–Years

Table 21 also reports the ten state–years with the largest implied GDP losses and the ten least affected state–years.

The least affected state–years exhibit losses of zero, reflecting the fact that lagged riot exposure is zero for these observations given the spatial decay in riot exposure. These cases provide a useful placebo benchmark, confirming that the aggregation procedure does not mechanically generate economic losses in the absence of violence.

The single most exposed state–year in the sample exhibits an implied GDP loss of approximately 8.87%, while the second most affected state–year experiences a loss of about 5.58%. The remaining third through tenth state–years in the upper tail exhibit implied losses of approximately 1.33%, substantially smaller than the most extreme episodes but still economically meaningful.

Table 21: Distribution and Extremes of Implied State–Year GDP Losses

Panel A: Distribution of implied state–year GDP losses (%)	
Mean	-0.2264
SD	0.6458
p1	-1.3079
p5	-0.8586
p10	-0.5618
p25	-0.2567
Median (p50)	-0.0621
p75	-0.0010
p90	-1.11×10^{-8}
p95	-2.22×10^{-13}
p99	0

Panel B: Ten state–years with the largest implied GDP losses			
State (pc11)	Year	Exposure	GDP loss (%)
24	2003	20.63927	-8.8694
24	2004	12.75955	-5.5801
10	2002	2.98534	-1.3344
1	2005	2.92553	-1.3079
6	2009	2.88559	-1.2901
24	2002	2.56349	-1.1469
10	2004	2.56235	-1.1464
24	2005	2.52623	-1.1304
27	2009	2.45348	-1.0980
18	2013	2.43649	-1.0904

Panel C: Ten least affected state–years (near-zero losses)			
State (pc11)	Year	Exposure	GDP loss (%)
32	2006	3.48×10^{-28}	0
18	1997	1.12×10^{-38}	0
19	2012	2.71×10^{-19}	0
33	2006	9.51×10^{-34}	0
19	1997	2.81×10^{-26}	0
32	2012	7.16×10^{-15}	0
18	2012	1.79×10^{-27}	0
32	2000	5.05×10^{-24}	0
33	2000	1.49×10^{-15}	0
33	2012	7.38×10^{-21}	0

Notes: Panel A reports the distribution of implied state–year GDP losses computed from the nightlight-weighted average of lagged riot exposure. Panels B and C report the ten most and least affected state–years. Implied GDP losses are calculated using a semi-elasticity of $\ln(GDP)$ with respect to riot exposure equal to 0.0045 and are expressed as percent differences relative to a no-riots counterfactual. Nightlight weights use winsorised mean calibrated nightlight intensity. State–years with zero losses correspond to observations with no lagged riot exposure.

D.4 Interpretation and Benchmarking

The most extreme state–year loss in the sample is approximately 8.87%, comparable in magnitude to the roughly 10% GDP level shortfall attributed to sustained terrorist activity in the Basque Country by Abadie and Gardeazabal (2003). The second most affected state–year also exhibits a large implied loss of approximately 5.58%, about half the magnitude of the Basque Country terrorism effect. In contrast, the third through tenth state–years experience substantially smaller losses, tightly clustered between 1.09% and 1.33%. This pattern is consistent with the nature of the underlying violence: the Basque Country experienced sustained and organized terrorist activity, whereas communal riots in India are episodic, localized, and short-lived.

E Firm entry and exit

E.1 Cohort-based exit index

Let $t_0 = 2000$ denote the base year. We define cohorts by plant age in the base year. For each plant i that is observed in year t_0, let

$$c(i) = \text{age of plant } i \text{ in year } t_0$$

be its cohort. We restrict attention to plants that are alive in t_0 and track these cohorts over time.

For each cohort c, district d, and year $t \geq t_0$, define:

$$\begin{aligned} M_c^t &= \text{total number (mass) of plants in cohort } c \text{ in year } t \quad \text{(all-India)}, \\ M_{cd}^t &= \text{number (mass) of plants in cohort } c \text{ in district } d \text{ in year } t. \end{aligned}$$

In particular, $M_c^{t_0}$ and $M_{cd}^{t_0}$ denote the cohort sizes in the base year.

For a given cohort c and district d, the decline in the number of plants between t_0 and t is

$$M_{cd}^{t_0} - M_{cd}^t$$

at the district level, and

$$M_c^{t_0} - M_c^t$$

at the all-India level. We interpret these as the number of exits from cohort c in district d and nationally, respectively.

We define the cohort- and district-specific relative exit share for $t > t_0$ as

$$\delta_{cd}(t) \;=\; \frac{M_{cd}^{t_0} - M_{cd}^t}{M_c^{t_0} - M_c^t}\,, \tag{E.1.1}$$

whenever the denominator is non-zero. Intuitively, $\delta_{cd}(t)$ measures the fraction of all exits in cohort c between t_0 and t that occur in district d.

Not all cohorts are equally important in a given district. We therefore weight cohorts by their initial importance in district d in the base year. Define

$$w_{cd} \;=\; \frac{M_{cd}^{t_0}}{\sum_{c'} M_{c'd}^{t_0}}\,, \tag{E.1.2}$$

so that w_{cd} is the share of cohort c in total base-year employment (or mass of plants) in district d, and $\sum_c w_{cd} = 1$ for each district d.

Our cohort-based exit index at the district–year level is then given by

$$\text{ExitIndex}_{dt} \;=\; \sum_c w_{cd}\,\delta_{cd}(t), \qquad t > t_0. \tag{E.1.3}$$

This index is a weighted average of cohort-specific relative exit shares, with weights reflecting the initial cohort composition of the district. Larger values of ExitIndex_{dt} indicate that district d experiences more exits, relative to the national pattern, among the plants that were active there in the base year.

E.1.1 Entry of new plants

We construct a measure of entry based on plant age. A plant is classified as an entrant in year t if the ASI reports its age in year t as exactly 1 year.

Let Entrants_{dt} denote the number (mass) of plants in district d and year t whose age equals 1. Formally,

$$\text{Entrants}_{dt} \;=\; \sum_{i \in d,t} \mathbf{1}\{\text{age}_{it} = 1\}$$

We next aggregate these entrant counts across all districts to obtain the total number of entrants in India in year t:

$$\text{Entrants}_{t} \;=\; \sum_{d'} \text{Entrants}_{d't}$$

We define district–year entry shares by normalising entrant counts by these national totals. Thus, the share of all new plants in India that locate in district d in year t is

$$\text{EntryShare}_{dt} \;=\; \frac{\text{Entrants}_{dt}}{\text{Entrants}_{t}}, \tag{E.1.4}$$

These entry-share variables measure the fraction of all entrants in India in year t that are located in district d. A larger value of EntryShare_{dt} therefore indicates that a district attracts a disproportionately large share of new industrial activity in that year.

F IHDS summary statistics

Table 22: Summary statistics: IHDS confidence and household variables

Variable	Observations	Mean	Std. dev.	Min	Max
Confidence in politicians	1,081	2.438	0.661	1	3
Confidence in police	1,081	1.306	0.659	1	3
Confidence in courts	1,082	2.029	0.745	1	3
Confidence in state government	1,079	2.032	0.771	1	3
Confidence in military	1,027	1.609	0.631	1	3
Confidence in panchayat	1,080	1.932	0.732	1	3
Confidence in news media	461	1.479	0.693	1	3
Confidence in private schools	461	1.254	0.496	1	3
Confidence in government schools	461	1.521	0.664	1	3
Confidence in hospitals	461	1.275	0.523	1	3
Household owns house	1,074	0.929	0.257	0	1
Household owns livestock	1,088	.389	.488	0	1
Any mobile phone owned	461	0.935	0.247	0	1
Any vehicle owned	1,087	0.695	0.461	0	1
Motorcycle owned	1,087	0.290	0.454	0	1
Cycle owned	1,087	0.535	0.499	0	1
School books: Total value Rs	1,086	1606.7	3138.6	0	40000
Post-riot interview indicator (D)	1,088	0.494	0.500	0	1
Hindu (0/1)	1,088	0.750	0.433	0	1
Muslim (0/1)	1,088	0.230	0.421	0	1
Urban residence (0/1)	1,088	0.447	0.497	0	1
Highest adult education in HH	1,087	8.603	5.110	0	15

Notes: Confidence variables are measured on a three-point scale, with higher values indicating lower confidence. All variables are drawn from the India Human Development Survey (IHDS). The sample is restricted to district and survey-wave cells in which a single riot occurred during the IHDS interview period. Highest adult education in HH capped at 15.

G IHDS Data cleaning

We combine the India Human Development Survey (IHDS) with our riot database by harmonising geographic identifiers and constructing interview timing variables at the district and survey-wave level.

First, we assign each riot to a 2011 Census district. Second, we prepare the IHDS data for timing-based analysis. Using the IHDS interview date variable, we construct an interview date for each respondent and summarise the fieldwork period by district and survey wave. Specifically, for each district and wave, we compute the earliest and latest interview dates and the number of interviews conducted, which defines the district-specific interview window. Third, we link riots to IHDS interview windows. We parse riot dates from the riot dataset and, for each district and survey wave, identify whether a riot occurred during the corresponding IHDS interview window. We also compute whether any riots occurred 60 days before the first interview date in that district and wave. Finally, we restrict the sample to district-by-wave cells with a well-defined single event. We retain only district–survey-wave cells in which exactly one riot occurred during the interview window and drop cells with any riots occurring in the 60 days prior to the start of the interview period. This restriction ensures a clean pre-period and avoids ambiguity from multiple events occurring within the same interview window. The summary statistics for the riots incidents and the intervew dates are summarised in Table 23.

Table 23: Riots used in the IHDS confidence analysis

Riot date	District	State	Interview start	Interview end	Days before	Days after	Window (days)	Obs.
2004-07-21	Muzaffarnagar	Uttar Pradesh	2004-01-17	2005-01-18	186	181	367	79
2004-11-01	Kachchh	Gujarat	2004-01-12	2005-10-04	294	337	631	90
2005-02-10	Vadodara	Gujarat	2004-05-02	2005-05-19	284	98	382	224
2005-02-20	Birbhum	West Bengal	2005-02-03	2005-06-27	17	127	144	153
2005-07-05	Faizabad	Uttar Pradesh	2004-04-28	2005-09-28	433	85	518	80
2012-06-24	Srinagar	Jammu and Kashmir	2012-06-06	2012-07-12	18	18	36	113
2012-07-28	Dakshina Kannada	Karnataka	2012-06-15	2012-08-22	43	25	68	192
2012-08-11	Mumbai	Maharashtra	2012-06-04	2012-12-08	68	119	187	158

Notes: The table reports the eight district–survey-wave cells in which exactly one riot occurred during the IHDS interview period. Interview start and end dates correspond to the first and last household interviews conducted in the district and survey wave. Days before (after) measure the number of days between the riot date and the start (end) of the interview window. Obs. reports the number of households in the district–survey-wave cell; one Kachchh observation has a missing interview date and is therefore excluded from the regression sample.

G.1 Riots data collection and cleaning

We follow the approach of Varshney and Wilkinson (VW) in collecting data on religious riots from the *Times of India*, Mumbai edition. For the period 2006–2013, we access digitised newspaper archives via ProQuest, while for post 2013 years we use Factiva. Together, these sources provide comprehensive coverage of the newspaper archive over our sample period. We use data up to 2014 in line with the data availability of The Defense Meteorological Program (DMSP) Operational Line-Scan System (OLS) nightlight data.

To identify riot-related articles, we apply the following keyword-based search query. For Factiva, the query is:

```
( riot* OR pogrom* OR massacre* OR lynch* OR arson OR curfew* OR
((communal OR sectarian OR religio* OR interfaith OR "minority community")
NEAR/5 (violenc* OR clash* OR tension* OR unrest OR disturb* OR attack*))
OR
((Hindu* NEAR/3 Muslim*) OR (Sikh* NEAR/3 Muslim*) OR (Sikh* NEAR/3 Hindu*)
OR (Hindu* NEAR/3 Christian*)) OR
("stone pelting" OR "mob lynching" OR "communal clash" OR "religious riot"
```

```
OR "communal violence") )
AND NOT ( hd=("Classified Ad") OR cricket OR football OR tournament OR
movie OR music OR festival OR review OR match OR "box office" OR horoscope
OR astrology OR recipe OR award )
```

All articles returned by the search algorithm are manually reviewed by research assistants. During this hand-coding stage, we drop events that are clearly unrelated to Hindu–Muslim violence, including incidents that are unambiguously terrorism-related or involve non-religious forms of violence. We retain only those events that can be classified as Hindu–Muslim religious riots based on the article content.

Could small, unreported altercations on Friday–festival days cumulatively depress annual luminosity without appearing as "riots"? That would require a district-wide, systematic impact substantial enough to survive annual averaging yet persistently escape journalistic and administrative recording—an implausible scenario given the scrutiny of Indian public order. By contrast, the links between festivals and riots Jaffrelot (2010), and between riots and broader economic outcomes such as lending in India Fisman et al. (2020), property value in the USA (Collins and Margo, 2007), labour market outcome in the USA (Collins and Margo, 2004a) are widely documented.

G.2 District definitions

Districts are defined using the ASI 1998 district classification, following Martin et al. (2017), to ensure consistent linkage of firm-level ASI microdata over time. All observations are mapped to this baseline district coding, and external datasets are harmonised accordingly.

G.3 Nightlights data cleaning

Night-time lights data are obtained from the SHRUG database at the 2011 Census district level. SHRUG aggregates pixel-level nightlights to districts using averages weighted by the number of grid cells, which reflects the geographic area of each district and ensures consistency across district boundaries. In years where nightlights data are available from two different satellite sources, we use the simple average of the two measures. When mapping from 2011 Census districts to our analysis districts, we aggregate nightlights using averages weighted by the number of grid cells.

H ASI Data Construction and Cleaning

Our ASI cleaning, aggregation, and sample construction procedures closely follow Martin et al. (2017) in their treatment of the ASI.

H.1 Data and sampling structure

We use an establishment-level panel from the Annual Survey of Industries (ASI) covering the period 2000–01 through 2007–08. The ASI sampling frame covers all registered (formal) manufacturing firms in India. Large firms are classified as part of the *Census* sector and are surveyed every year, while smaller firms belong to the *Sample* sector and are surveyed intermittently. The ASI provides firm-level sampling multipliers that allow the construction of representative aggregates at the state-by-industry level.

H.2 Dropping rules and sample restrictions

We apply the following exclusions to construct the analysis sample:

1. **Closed firms.** We drop firm–year observations classified as closed based on ASI closure indicators and labour–output consistency checks.
2. **Negative capital.** We drop observations with negative values of fixed capital stock.
3. **Zero labour and total output.** We drop observations where both labour and total output are zero.
4. **Zero capital and total output.** We drop observations where both capital and total output are zero.

In addition, firms that have zero labour in all observed years are dropped at the firm level. Observations with missing district identifiers are also excluded. After these restrictions, each firm–year appears at most once in the final panel.

H.3 Construction of core variables

Total output. Total output is constructed from Block J of the ASI and related revenue components. Our baseline measure, `TotalOutput`, includes ex-factory value of output together with additional revenue-side components (such as services income, electricity sold, own construction, and net goods flows).

Labour. Labour input is measured as the total number of persons engaged at the establishment level.

Capital. Capital is defined as the closing stock of fixed capital.

H.4 Additional cleaning and harmonisation

After stacking yearly panels, we apply further harmonisation steps. Year definitions are aligned to a consistent calendar convention. Firm age is defined as year minus the reported initial year of production. Initial production years equal to zero, below 1500, or greater than or equal to the observation year are treated as missing. Firms with inconsistent or implausible initial production years ($\leq$1600 or $\geq$2011) are dropped.

H.5 Price deflation

All nominal monetary variables are converted to real terms using year- and industry-specific deflators. To ensure consistency with ASI accounting conventions and survey timing, all deflators are aligned to the calendar year definition used in the ASI data.

Output deflator. Output and revenue variables are deflated using wholesale price indices (WPI) matched at the three-digit NIC industry level. We construct a concordance between three-digit NIC industries and WPI commodity groups and compute industry-specific price indices normalised to unity in the base year 1998. Real output measures are then obtained by dividing nominal output by the corresponding industry-year WPI index. This procedure is applied to total output and services income.

Wage deflator. Labour compensation variables are deflated using the Consumer Price Index (CPI). Nominal total emoluments are divided by the CPI (normalised to 1998 = 100) to obtain real wage payments. We additionally construct real average wages as the ratio of real total emoluments to employment.

Capital deflators. Capital stocks are deflated using two alternative price indices. Our baseline capital deflator is based on the wholesale price index for machinery, normalised to unity in 1998, and is applied to fixed capital and plant and machinery stocks. As a robustness check, we also construct an alternative capital deflator using the Penn World Table investment price index, similarly normalised to the 1998 level. Both deflators are merged by calendar year.

Implementation details. All deflators are merged at the firm-year level. Observations without valid matches to the required deflator series are dropped to maintain internal consistency. Nominal variables are retained alongside their deflated counterparts to facilitate robustness checks.

H.6 Aggregation to the district level

The final dataset used for estimation is constructed at the district–year level. Firm–level outcomes are aggregated to districts using ASI sampling multipliers. Specifically, firm–year variables are weighted by the ASI multiplier and summed within district–year cells. District-level averages are then obtained by dividing weighted totals by the weighted number of firms. All regressions are conducted using these district–year aggregates.