

# Religious Institutions and Economic Development in India<sup>\*</sup>

Kishen Shastry<sup>†</sup>

University of Cambridge

February 2026

## Abstract

Hindu monastic institutions (*mutts*) in Karnataka run schools, colleges and hostels, and their heads carry weight in state politics. We study how these institutions shaped local responses to India's economic and political reforms of the early 1990s. We build a new geocoded dataset of 253 mutts and link it to subdistrict panels from the Economic Census and Population Census and to assembly election results. Using a difference-in-differences design with continuous pre-1990 mutt exposure, we find that more exposed subdistricts added more non-farm jobs after the reforms: 15–20 percent more in our preferred estimates, mostly in manufacturing and private firms. Government employment barely changed. The employment gains are absent in 1998 and appear by 2005. More exposed subdistricts also gained more secondary schools and colleges per capita, and more exposed constituencies had narrower victory margins after 1994, with no detectable change in turnout or the number of parties. These patterns are consistent with mutts both running educational institutions and acting as political intermediaries.

**JEL codes:** Z12, O15, O17, D72, I25, P16

---

<sup>\*</sup>This paper is the first chapter of my PhD thesis at the University of Cambridge, examined in February 2026. I am deeply grateful to Sriya Iyer for her guidance and support. I thank my examiners, Benjamin Arold and Mohamed Saleh, for their detailed comments, which I am incorporating in a revised version, and Toke Aidt and Anand Shrivastava for helpful comments. I also thank participants at the Advanced Graduate Workshop at Azim Premji University and the SCORE 3 Early Career Panel (Social Consequences of Religion Initiative, University of Cambridge) for valuable feedback. All errors are my own.

<sup>†</sup>Email: ksk49@cam.ac.uk.

# 1 Introduction

Religious institutions have long played both secular and sacred roles in society, especially in developing regions where state capacity is limited. In India, Hindu monastic establishments known as *mutts* (or *mathas*) are a prime example. Religious organisations are important spiritual centres and also provide education, healthcare, and other social services to local communities (Iyer, 2018). This paper examines the economic and political consequences of variation in the influence of these Hindu monastic institutions across communities in the state of Karnataka<sup>1</sup>. We ask whether communities with greater historical exposure to mutts experience higher economic development and different political dynamics than those with lower exposure. In particular, we investigate whether higher mutt density is associated with gains in non-farm employment, improvements in educational attainment, and changes in electoral competitiveness. Understanding the impact of religious organisations like mutts matters for development policy, because they often function as alternative providers of public goods, and for political economy, because they can mobilise voters and shape local power structures.

This study contributes to the literature on religion and development by providing micro-level evidence on a specific type of religious institution. A large body of research has explored the broad relationship between religion and economic outcomes across countries (e.g. Barro and McCleary, 2003; Guiso et al., 2003; McCleary and Barro, 2006), but there is relatively less quantitative work that identifies the local impacts of particular religious organisations. Political economy research has documented how religious authorities and networks influence politics, from the role of churches in democratisation (Woodberry, 2012; Trejo, 2009) to religious endorsements in elections, yet much of this evidence is historical or qualitative. By focusing on Hindu mutts in the state of Karnataka, the analysis offers new insights in a contemporary developing country context. Mutts are uniquely influential in this region. They traditionally serve specific caste communities and have built extensive educational and charitable infrastructures since the late twentieth century (Ikegame, 2010, 2013, 2011). Mutt leaders, or *swamis*, are often courted by politicians for electoral support, which suggests that these institutions can affect political competition on the ground. To the best of current knowledge, this paper is the first to measure rigorously the socioeconomic impact of Hindu mutts and to situate the findings within the broader economics of religion.

The empirical results in this work indicate a substantial influence of mutts on economic

---

<sup>1</sup>As of 2025, the state of Karnataka has a population of over 70 million. It is a roughly US \$360 billion economy, accounting for about 8-9 percent of India's GDP, ranking among the country's three largest state economies, and functioning as a major export-oriented hub for information technology and other high-value services.

outcomes. We assemble novel data on the location of 253 mutts in Karnataka and combine it with subdistrict and constituency level panel data. To summarise the spatial footprint of mutts, we construct a pre-1990 measure of mutt exposure that counts weighted influence mutts within a 20 kilometre radius of each subdistrict centroid. Using the Socioeconomic High Resolution Rural Urban Geographic dataset (Asher et al., 2021), we then construct a panel that combines four rounds of the Economic Census for 1990, 1998, 2005, and 2013 with Population Census and Village and Town Directory information for 1991, 2001, and 2011.

The analysis shows that areas with historically active mutts have experienced faster structural transformation out of agriculture. In particular, the share of non-farm employment is significantly higher in mutt areas, reflecting greater take-up of skilled jobs after the reforms of the 1990s. We also find that mutt presence is associated with more secondary schools, senior secondary schools, and colleges per capita. This is consistent with mutts shaping local education capacity through two channels: (i) direct provision of mutt-run schools, colleges, and hostels and (ii) mobilize public support that expands overall higher-education infrastructure in mutt-influenced areas (Giriprakash, 2018). Interestingly, constituencies influenced by mutts exhibit closer races: victory margins are smaller, and party competition is more intense, while overall turnout and the fragmentation of the party system (as measured by the effective number of parties) remain essentially unchanged.

Pre-1990 mutt exposure predicts where non-farm employment and higher education expand after liberalisation and decentralisation, and where electoral competition tightens after 1994. The sectoral patterns and the timing hint at a plausible mechanism in which mutts invest and channel investment in human capital and social welfare, and provide organisational structure for community political voice. These channels align with theories of religious institutions as providers of club goods and as agents of social mobilisation (Iannaccone, 1992; Berman, 2000; Carvalho, 2016).

The findings have implications for both scholarship and policy. They show that religious organisations can be important catalysts of development, complementing and in some cases strengthening the state’s efforts in education and employment generation. They also highlight a nuanced role of religion in politics. Rather than simply reinforcing traditional authority, mutts in this context appear to foster competition and accountability by empowering new groups of voters. These insights contribute to a growing recognition that the architecture of religion, including the size, governance, and objectives of religious institutions, matters for economic and political outcomes (Iyer, 2016; Becker et al., 2025). By documenting these effects in the context of Hindu monastic orders, the paper expands the predominantly Western and Abrahamic focus of existing economics of religion research.

The remainder of the paper is organised as follows. Section 2 reviews related literature on religion and economic development, religious institutions and political economy, and organisational models of religion. Section 3 provides background on the history of mutts in Karnataka, and Section 4 describes the data. Section 5 outlines the empirical strategy, and thereafter Section 6 presents the results. Finally, Section 7 concludes.

## 2 Related Literature

This paper contributes to an expanding economics of religion that links religious institutions to development outcomes. Survey pieces emphasize that religion affects economic behavior through beliefs, norms and organisational structures, and call for more work on specific institutions in non-Western settings (Iannaccone, 1998; Iyer, 2016; Carvalho et al., 2019; Becker et al., 2025). A large empirical literature studies the association between religiosity and aggregate growth (Barro and McCleary, 2003; McCleary and Barro, 2006; Guiso et al., 2006), but most of it treats religion as an individual characteristic rather than as an organised provider of services. The focus here on Hindu *mutts* in one Indian state provides a complementary perspective. Mutts effectively combine spiritual authority with control over land, schools, and hostels (Ikegame, 2010, 2011, 2012a, 2013; Sood, 2006). More broadly, religious organisations in India frequently provide education and other welfare services alongside religious functions, as documented in Iyer’s work on the institutional landscape (Iyer, 2018) and in studies of religious organisations as providers of schools and health services (Iyer et al., 2014; Borooah and Iyer, 2005). Related work on faith-based provision in other settings shows that churches and religious charities can substitute for or complement the state in welfare provision (Hungerman, 2005; Gruber and Hungerman, 2007; Hungerman, 2020; Auriol et al., 2020). Compared with this largely cross-sectional and organisation-level evidence, this paper uses subdistrict and constituency panels and a time-invariant measure of exposure to a particular class of Hindu monastic institutions to track how pre-reform mutt presence is associated with the subsequent evolution of non-farm employment, higher-education capacity and electoral competition over nearly three decades. In doing so it connects the economics of religion to institutional and development economics (North, 1990; Greif, 2006; Iyer, 2016; Casson et al., 2010).

A first body of work to which this paper relates examines the relationship between religion and human capital. Historical studies show that Protestant regions of Europe invested earlier in schooling and that missionary activity raised literacy and incomes in Africa and Latin America (Becker and Woessmann, 2009; Woodberry, 2012; Valencia Caicedo, 2019). Complementary evidence from the Middle East shows that state-led modernization

and the expansion of public schooling interacted with religious identity in shaping human capital accumulation in Egypt (Saleh, 2015, 2016). For India, Calvi (2022) documents that Protestant missions had persistent effects on girls' schooling, while Borooah and Iyer (2005) and Iyer et al. (2014) show that religious organisations continue to play a central role in education and health provision. Other contributions highlight that religious elites can oppose mass education when it threatens their power. Chaudhary and Rubin (2011) show that districts dominated by traditional Muslim religious institutions in colonial India invested less in public schooling, and Malik and Mirza (2022) show that shrine elites in Pakistani Punjab slowed literacy growth when they gained control over development funds. The results are closest in spirit to this work, but differ in three ways. First, we study Hindu mutts rather than Christian or Islamic missions and shrines. Second, the outcome set includes not only education but also non-farm employment and political competition, allowing a broader view of economic change. Third, the context is a democratic Indian state in the post-1990 reform period, rather than a colonial or military regime. In this typology, mutts in Karnataka resemble more the "missionary" model, investing heavily in schooling, hostels, and higher education for their communities, than the shrine-based elites who may block mass education when it undermines their authority. The analysis also complements work on education expansion driven by the state and secular providers, such as large-scale school construction and private schooling growth, by focusing on religious organisations as an additional channel for human capital formation.

Beyond religion-specific studies, the paper is related to the literature on structural transformation and the rural non-farm economy in India. Work on agricultural productivity and rural diversification emphasises how local endowments and reforms shape the expansion of non-farm activities and wage employment (e.g. Foster and Rosenzweig, 2004). Research on misallocation and the geography of manufacturing and services stresses that spatial differences in local conditions and institutions condition the response to national reforms (Hsieh and Klenow, 2009; Ghani et al., 2016). A growing body of evidence on infrastructure, roads, electrification, and connectivity, shows that these investments can shift employment out of agriculture and raise rural incomes (Asher and Novosad, 2017, 2020). In this literature, mutts can be viewed as part of the local institutional infrastructure: pre-existing religious organisations with land, buildings, and networks that help some subdistricts take better advantage of economy-wide liberalisation and higher-education policies, over and above the effects of physical infrastructure and baseline geography.

A second body of work studies religion, political competition, and public goods. Political economy research for India and other democracies shows that more competitive elections and stronger media lead governments to expand public goods and respond more effectively

to shocks (Besley and Burgess, 2002; Besley et al., 2004; Finan and Mazzocco, 2021; Asher and Novosad, 2017, 2020). At the same time, identity-based organisations mediate access to the state. Caste and ethnic networks can both improve informal insurance and distort labour allocation (Munshi and Rosenzweig, 2006, 2016; Munshi, 2019). Religious leaders and organisations often act as political brokers: Trejo (2009) shows that competition between Catholic and Evangelical churches altered the Catholic Church’s stance towards indigenous movements in Mexico. In Pakistan, Sufi shrines combine spiritual and electoral power and shape the distribution of education and public goods (Malik and Mirza, 2022); and in Indian Punjab, *deras* mobilize marginalized castes and lobby for public services (Batra and Prabhakar, 2024; Bajwa, 2017). In Karnataka, ethnographic work describes mutts as power centers that bargain over candidate selection, ministerial portfolios and welfare schemes (Ikegame, 2010, 2012a,b, 2013; Sood, 2006), with journalistic accounts reinforcing their electoral salience.<sup>2</sup> The analysis relates this qualitative evidence to work on democratic decentralisation and Panchayati Raj reforms, which documents how local institutions and quotas reshape political competition and public goods provision (Aziz, 2000; Johnson, 2003; Besley et al., 2004). By showing that pre-reform mutt exposure is associated with smaller victory margins in assembly contests after 1994, the paper provides evidence that Hindu religious institutions are linked to tighter electoral competition rather than only to the protection of incumbents. The results hint at a plausible mechanism for better governance, and that mutts act as local intermediaries that affect how competitive elections are on the ground.

Finally, the paper speaks to organisational models of religion. Early economic models treat religious participation as part of household time allocation and human capital investment (Azzi and Ehrenberg, 1975; Iannaccone, 1990). Building on Olson’s (1965) logic of collective action, club-good models view religious groups as providers of collective benefits that use costly participation rules to screen out free riders and sustain cooperation (Iannaccone, 1992; Stark, 1999). Subsequent work shows how religious strictness and shared identity support high-commitment activities such as education, mutual insurance and even political violence (Berman, 2000; Berman and Laitin, 2008; Berman, 2009; Henrich, 2009; Chen, 2010). Recent contributions emphasize identity-based organisations that invest in group identity to mobilize costly effort and sustain participation (Carvalho, 2013; Carvalho and Koyama, 2016; McBride, 2015, 2016, 2019), and related evidence shows that adverse shocks can raise religiosity and demand for religious activity (Bentzen, 2019).

Research on formal and informal institutions stresses that informal religious authority often coexists with, and shapes, formal state structures (Casson et al., 2010; Platteau,

---

<sup>2</sup>See, for instance, Giriprakash (2018); Poovanna (2023); Menon (2024).

2008; Hansen and Stepputat, 2001; Carvalho and Koyama, 2016). Mutts fit squarely within this perspective. They are decentralized, caste-linked organisations that build religious and social capital through long-term guru-disciple relations and caste-based solidarity, run large networks of schools and hostels, and negotiate directly with the state (Ikegame, 2011, 2012a; Sood, 2006). The evidence in this paper provides a quantitative example of how such identity-based religious organisations can influence the trajectory of local development and political competition in a large Hindu-majority democracy, a setting largely missing from existing empirical work.

The interaction between religious organisations and government policy also alters the incentives faced by religious leaders. Governments regulate religious institutions through subsidies, land grants, and oversight. In India, Hindu temples are subject to varying degrees of government control. However, mutts in Karnataka have received land and grants to build colleges and hostels, and they often bargain with governments over recognition, aided status, and the location of new or upgraded higher-education capacity, which can increase public investment in education in mutt-influenced areas (Ikegame, 2011)<sup>3</sup>. Once established, these mutts became influential actors in their own right and in some cases challenged political leaders. This fits models in which religious organisations professionalise and diversify their activities when given resources, yet retain legitimacy through their spiritual mission. The results in this paper, which document sizeable development effects for mutt-exposed areas and changes in electoral competition, reinforce the idea that understanding religion’s impact on economic and political outcomes requires attention to how religious institutions are organised, how they finance themselves and how they interact with the state (Iyer, 2016; Becker et al., 2025).

## 3 Institutional background

Hindu religious institutions in India take many organisational forms, from large pilgrimage temples to small caste shrines and guru-centred movements. In Karnataka, one particular form stands out: the *mutt* or *matha*, a Hindu monastic institution headed by a hereditary or appointed guru who presides over a network of land, temples, schools, hostels, and charities. Ethnographic work describes mutts as caste-based, community-oriented centres of religious, educational, and social life, anchored in a particular caste community and led by an ascetic who embodies that community in public life (Sood, 2006; Ikegame, 2010, 2012a). The largest mutts function as apex bodies for extensive networks of schools, colleges, hospitals,

---

<sup>3</sup>On these land grants and educational investments, see also contemporaneous reporting in Giriprakash (2018).

and other social service institutions, channelling resources from devotees, diaspora donors, and the state into their home communities (Sood, 2006). Rather than purely spiritual retreats, mutts in Karnataka have long acted as political and social intermediaries: gurus dispense ritual recognition and access to material support, arbitrate local disputes and lobby the state for resources (Ikegame, 2011, 2013; Sood, 2006). Sood characterises the resulting configuration as a regional “Matha State”, in which the diverse religious, welfare and electoral activities of mathas generate a shared horizon of power for caste communities in Karnataka’s public sphere (Sood, 2006), while Ikegame stresses their role as a “parallel state” that combines spiritual authority with control over land, education and welfare provision (Ikegame, 2013).

Mutts in Karnataka differ from the large temple complexes that fall under the statutory control of the Hindu Religious and Charitable Endowments departments in many states. Most mutts in Karnataka are registered as private religious trusts or educational societies. They control their own property, appoint their own successors, and run their own institutional networks, often through separately registered educational or charitable societies (Ikegame, 2012a; Sood, 2006; Iyer, 2018). This institutional autonomy allows gurus to act as independent administrative and juridical authorities for their sponsoring caste community, mediating between followers, other mutts, and the state administration (Sood, 2006). The result is a dense landscape of relatively autonomous Hindu organisations that overlap only partially with the formal state and party system. Journalistic syntheses emphasise that mutts are typically caste-linked spiritual centres with tens of thousands of followers each, providing a focal point for social and political mobilisation within specific communities.<sup>4</sup> Estimates reported in the popular press suggest that Lingayats, a prominent caste in Karnataka, alone have more than a thousand mutts, with historical counts of roughly 5,300 Lingayat mutts at Independence and around 1,100 still functioning today.<sup>5</sup> Mutts headed by Lingayat and Vokkaliga seers have been especially prominent in state-level politics since at least the 1960s (Ataulla, 2022), and Sood documents how their success has inspired a wave of new mutts among backward-class and Dalit communities seeking similar institutional vehicles for welfare and political representation (Sood, 2006).

The scale of individual mutts is also substantial. Siddaganga Mutt in Tumakuru, a major Lingayat institution, combines a traditional monastic centre with a large education network. Its educational trust runs on the order of 125-130 schools and colleges, ranging from primary schools to engineering, pharmacy, and nursing colleges, and associated blind

---

<sup>4</sup>For an overview in the popular press, see, for example, Poovanna (2023).

<sup>5</sup>These figures are drawn from reporting in ET Online (2023).

schools and special education centres.<sup>6</sup> Contemporary reporting and institutional materials note that the mutt’s institutions educate roughly 50,000 students across Karnataka and that its residential *gurukula* provides free food, education, and accommodation to around 10,000 poor children irrespective of caste or creed.<sup>7</sup> Sood’s ethnography of Siddaganga and other major mutts similarly highlights their student hostels, Sanskrit schools, and professional colleges as key instruments for expanding the life chances of their caste followers (Sood, 2006).

A second example is Suttur Math near Mysuru, whose education and welfare activities are organised through the JSS Mahavidyapeetha. Originating from a small hostel established in 1928, JSS Mahavidyapeetha has grown into a large educational movement. Recent statements report that it now oversees more than 300 institutions, from kindergartens to postgraduate and postdoctoral centres, catering to the educational needs of over 100,000 students in India and abroad. Its network includes a science and technology university, a health sciences university, engineering colleges, polytechnics, management schools, and a large free residential school at Suttur for children from low-income rural families. The mutt also sponsors hospitals, rural development programmes, and cultural activities,<sup>8</sup> making it a key provider of both public-good-like services and private educational opportunities in large regions. Sood (2006) notes that such Vokkaliga institutions became exemplars for other caste groups, who explicitly read their educational and charitable networks as evidence that having a mutt can transform a community’s economic and political position.

A third illustration comes from Sri Kshetra Dharmasthala in coastal Karnataka. Although centred on a major pilgrimage temple rather than a monastic order, its Shri Dharmasthala Manjunatheshwara Educational (SDME) Society operates in a similar way to mutt-linked educational trusts. The SDME Society reports that it manages 56 educational institutions from kindergarten to doctoral studies across multiple districts of Karnataka, spanning general, legal, technical, medical, and management education.<sup>9</sup> Together these institutions educate tens of thousands of students and employ several thousand staff, and they are explicitly framed as instruments for rural uplift and value-based education. Like the large mutts, Dharmasthala’s network blurs the lines between religious philanthropy, private provision, and publicly funded welfare schemes.

Contemporary reporting reinforces the picture of mutts and related religious trusts as

---

<sup>6</sup>These figures are reported on Siddaganga Mutt’s website (Sree Siddaganga Education Society, 2025).

<sup>7</sup>See, for example, Sree Siddaganga Education Society (2025); Poovanna (2023).

<sup>8</sup>These details are drawn from JSS Mahavidyapeetha’s institutional profile (JSS Mahavidyapeetha, 2022).

<sup>9</sup>These counts and descriptions are from the Shri Dharmasthala Manjunatheshwara Educational Society’s own materials (Shri Dharmasthala Manjunatheshwara Educational Society, 2018).

major non-state providers and political intermediaries. Journalistic accounts describe how Lingayat and Vokkaliga mutts in northern and southern Karnataka run extensive chains of schools, colleges, hostels, and hospitals, and how chief ministers and party leaders regularly visit seers to seek blessings and negotiate support during elections.<sup>10</sup> Mutts have mobilised followers around reservation policies, language issues, and temple management, and have at times confronted ruling parties over the distribution of welfare schemes or the selection of candidates (Ataulla, 2022).<sup>11</sup> Ethnographic work shows that these activities are not incidental: mutts routinely use mundane tools such as land records, legal correspondence, and alumni associations to negotiate land from the state, elicit donations, administer their networks of institutions, and intervene in electoral politics (Sood, 2006). These accounts portray mutts as intermediaries between followers and the state, as large employers and landlords in their own right and as important centres of political bargaining.

Despite their prominence in public life, mutts have received little attention in economics. Most quantitative work on religion and development in India has focused on formal temples or on broad measures of religious composition at the district or state level. In contrast, mutts are geographically local, historically persistent and organisationally independent of the state. In many parts of rural Karnataka, they are the only Hindu institutions that combine land, religious authority and an autonomous capacity to invest in education and public goods (Ikegame, 2012a, 2013; Sood, 2006; Iyer, 2018). Sood’s account of the “Matha State” emphasises that communities with established mutts appear to enjoy distinctive developmental and political advantages, and that other caste groups respond by creating their own mutts to access similar benefits (Sood, 2006). This paper focuses on older, influential mutts that predate the 1991 economic reforms and exploits their pre-existing spatial distribution to study how local religious institutions shape subsequent patterns of non-farm employment, higher education and electoral competition.

### 3.1 Post-1990 economic and political shocks

The period covered by this paper coincides with a set of large, economy-wide shocks that reshaped the opportunities facing rural Karnataka. Three are particularly relevant: national market-oriented reforms initiated in 1991, the rapid expansion and partial privatisation of higher education, and the political decentralisation associated with the 73rd and 74th Constitutional Amendments and the Karnataka Panchayat Raj Act of 1993.

In 1991, India faced a severe balance-of-payments crisis. The subsequent reform

---

<sup>10</sup>See, for instance, Giriprakash (2018); Poovanna (2023); Menon (2024).

<sup>11</sup>On recent episodes, see also Poovanna (2023); ET Online (2023).

package liberalised trade, dismantled most industrial licensing, eased restrictions on foreign investment and began to reform the financial sector (Ahluwalia, 2002). Tariffs were cut sharply, quantitative restrictions were phased out and industrial policies shifted from detailed controls over capacity and product choice to a more rules-based regime that treated private investment more favourably (Panagariya, 2004). These changes were national in scope, but they expanded demand for manufactured goods and services and relaxed long-standing constraints on firm entry and growth, especially in formally registered manufacturing. In a setting where mutts already controlled land, buildings and social networks, these reforms increased the returns to local organisational capacity to help followers start firms, access credit and plug into supply chains. The empirical analysis below, therefore, treats the 1990s liberalisation as a common shock that raises the payoff to having pre-existing religious organisations capable of brokering access to new non-farm opportunities.

The same period saw a rapid expansion of secondary and tertiary education. National policy encouraged the growth of private “self-financing” colleges, particularly in technical and professional fields, while state governments retained substantial discretion over regulation and aid. Karnataka moved early and aggressively in this direction. Especially during the 1990s, the state relied heavily on private trusts, religious institutions and minority organisations to supply colleges and professional institutes. Official reports and policy reviews describe Karnataka as a pioneer in allowing private colleges to be set up by education trusts, religious *mutts* and minority institutions, and document a sharp rise in the number of recognised colleges and professional institutes after 1990 (Mathew, 2016). These policies created a supply-side environment in which mutts that already ran schools and hostels could scale up into upper-secondary schools, degree colleges and engineering institutes, often as aided or recognised institutions under state schemes. In the empirical work, the post-1990 expansion of higher education is thus interpreted as a common policy shift that amplifies the role of mutts as providers and intermediaries of human capital.

A third set of shocks concerns political institutions. At the national level, the 73rd and 74th Constitutional Amendments (1992) gave constitutional status to rural and urban local bodies, mandated regular local elections, and required state governments to establish State Election Commissions and State Finance Commissions (Government of India, 1992a,b). The Karnataka Panchayat Raj Act of 1993 reorganised local government into a three-tier structure (Gram, Taluk and Zilla Panchayats) with elected representatives, reservations for Scheduled Castes, Scheduled Tribes and women, and formal responsibility for a broad set of local development functions (Government of Karnataka, 1993; Aziz, 2000). These reforms altered the channels through which state and central funds reached rural areas

and increased the salience of local institutions that could mobilise voters and interface with panchayat and assembly members. In Karnataka, mutts were well placed to play this role. Ethnographic work suggests that from the mid-1990s onwards, seers became more deeply involved in candidate selection, coalition-building and the allocation of welfare schemes (Ikegame, 2010, 2012a; Sood, 2006).

In what follows, we exploit the timing of these shocks. The Economic Census results compare 1990 with post-reform rounds in 1998, 2005 and 2013, the education results compare the 1991 Population Census with 2001 and 2011, and the elections analysis contrasts assembly contests before and after the first post-reform state election in 1994. The preferred interpretation is that macroeconomic liberalisation, higher-education policies and political decentralisation shifted the returns to local organisational capacity across Karnataka in broadly similar ways. Given the long historical roots and spatial stability of major mutts, their pre-1990 footprint can then be used to study how local religious institutions mediated the impact of these common shocks on structural change and political competition.

## 4 Data

The empirical analysis combines three ingredients: a new geocoded inventory of Hindu monastic institutions (mutts) in Karnataka, administrative data on employment, population, and local infrastructure from the Socioeconomic High-resolution Rural-Urban Geographic (SHRUG) database, which links four rounds of the Economic Census (1990, 1998, 2005, 2013) to three rounds of the Population Census (1991, 2001, 2011) and the associated Village and Town Directories for every village and town in India (Asher et al., 2021). Additionally, a panel of state assembly elections is drawn from the SHRUG elections module.<sup>12</sup> For Karnataka, SHRUG provides measures of non-farm employment from the Economic Censuses and demographic and public-goods variables from the Population Censuses and directories, which are aggregated to a common spatial frame based on subdistricts and assembly constituencies.

### 4.1 A new dataset on mutts

The first contribution of the paper is a new statewide dataset on large mutts in Karnataka. The starting point is a list of subdistrict names, combined with searches via the Google

---

<sup>12</sup>The SHRUG elections module is based on vote returns from the Trivedi Center for Political Data (Jensenius and Verniers, 2017).

Places API for potential mutts using a broad set of English and Kannada tokens, including generic terms such as “mutt”, “matha”, “mahasamsthana”, “peetha”, and “ashram”. For each query, all places in Karnataka that match these tokens are retrieved, together with their coordinates and a salience proxy (the number of Google ratings).

The raw search results contain many locations that are not monastic institutions, such as ordinary temples, schools, and so on. Automated filters are therefore combined with extensive manual checks. Only places tagged as Hindu temples or places of worship are retained. Other religions and obvious non-targets are dropped, and entries that appear under multiple spellings or tokens are de-duplicated. The remaining places are then inspected by hand, using names and associated web pages where available, to confirm that they are mutts rather than ordinary temples.

The final dataset contains 253 mutts. Historical sources, including institutional histories and local gazetteers, indicate that these are long-standing establishments that predate the 1990 Economic Census, often by centuries. The list includes all of the major mutts that figure prominently in accounts of Karnataka’s religious landscape, such as Sringeri Sharada Peetham, Udupi Krishna mutt and its associated Ashta mutts, Siddaganga mutt near Tumakuru, Adichunchanagiri mutt in Mandya district, and Suttur mutt near Mysuru. This gives confidence that the search procedure captures the most influential institutions for the purposes of this paper.

To translate point locations into a continuous measure of mutt exposure, the analysis works at the subdistrict level. For each subdistrict, a representative point (the centroid of the subdistrict polygon) is taken and distances are computed to every mutt in the dataset. Mutts primarily serve a local catchment of villages that lie within day-trip distance of their schools, hostels, and welfare activities, and these catchments often cut across administrative boundaries. To capture this, a 20 kilometre neighbourhood is defined around each subdistrict centroid and all mutts within this radius are considered. Each such mutt receives a distance weight given by a 20 kilometre log-kernel and a salience weight based on its number of Google ratings. The resulting distance-salience weights are then summed and divided by the subdistrict’s 1991 population to obtain a per-capita mutt exposure. Finally, mutt exposure is standardised to have mean zero and unit variance across subdistricts<sup>13</sup>. Throughout the paper, a one-unit increase in exposure should be


---

<sup>13</sup>The distribution of this kernel-based exposure is highly right-skewed, with a small number of subdistricts lying very close to large, high-salience mutts. In Section 6.4 as robustness checks, we also construct exposure measures based on (i) a salience-weighted sum of mutts whose points lie strictly inside the subdistrict boundary, without any spillover across subdistricts, and (ii) an unweighted raw count of mutts that treats all mutts as equally salient. The main results are similar when using these alternative exposure measures instead of the baseline salience-weighted kernel.

read as a one standard-deviation higher value of this pre-1990 mutt exposure.

Figure 1 visualises the spatial distribution and potential reach of mutts in Karnataka. Panel (a) plots mutt locations as points over subdistrict boundaries. Panel (b) overlays simple 20 kilometre circles around each mutt location, illustrating the local catchment areas that would arise if each mutt primarily served communities within day-trip distance. The dense overlap of circles in the central and southern parts of the state shows that many subdistricts lie within the notional catchment of multiple institutions, consistent with a geographically pervasive mutt presence. In the empirical analysis, this intuition is formalised through a 20 kilometre distance-weighted kernel around each subdistrict centroid.

Figure 1: Mutts in Karnataka and 20 kilometre circles



Notes: The figure shows the locations of 253 mutts in Karnataka and illustrative 20 kilometre circles around them. Panel (a) plots mutts as points over subdistrict boundaries. Panel (b) overlays a 20 kilometre radius circle around each mutt location, to convey the local catchment that a mutt might serve. The heavy overlap of circles in Panel (b) highlights that mutt influence is likely to be geographically widespread. In the empirical analysis, mutt exposure is constructed as a 20 kilometre log-kernel of mutt locations around each subdistrict centroid, weighted by a salience proxy (the number of Google ratings) and divided by the subdistrict's 1991 population. The circles in Panel (b), therefore, illustrate the geographic support of this kernel.

## 4.2 Economic outcomes, demographics, and infrastructure

The main economic outcomes and baseline covariates are drawn from the SHRUG database Asher et al. (2021), which links village and town level records from the Economic Census, Population Census, Village Directory, and Town Directory to consistent geographic identifiers. The starting point is the rural population census key for 1991. Villages that were classified as rural in the 1991 Population Census are retained, the Economic Census, Population Census Abstract, and Village and Town Directory modules are attached to these rural SHRUG identifiers, and the resulting data are aggregated to subdistricts using the standard subdistrict keys. When a 1991 rural village is reclassified as part of a town or municipality in 2001 or 2011, its census and directory information continues to be assigned to the same subdistrict.

Two implications are important for interpretation. First, the analysis focuses on rural Karnataka in the sense that the underlying communities were rural at the beginning of the period<sup>14</sup>. Second, because 1991 rural villages are followed forward even if they urbanise, mutt exposure is not mechanically confounded with contemporaneous rural–urban reclassification.

Economic outcomes are taken from four rounds of the Economic Census (1990, 1998, 2005, 2013). For each round, all non-farm workers across 1991 rural SHRUG villages within a subdistrict are summed. Per-capita measures are constructed by dividing each Economic Census count by the nearest Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplying by 1,000. The main outcome is total non-farm employment per 1,000 residents,<sup>15</sup> and separate series are also constructed for manufacturing, services, private, and government employment.

Baseline demographics are also used from the 1991 Population Census Abstract. Particularly, subdistrict-level shares of the population who are Scheduled Castes and Scheduled Tribes, the literacy rate, and the share of children aged 0–6 are computed. These 1991 characteristics are used both as covariates and as denominators for per-capita variables.

Similarly, local infrastructure variables like transport facilities, communications facilities

---


<sup>14</sup>Bengaluru city is excluded entirely and subdistricts that cannot be matched cleanly across sources are dropped.

<sup>15</sup>In the SHRUG, non-farm employment is constructed from the Economic Census, which is a complete enumeration of all non-crop-producing establishments in India, covering public and private firms in both manufacturing and services and in the formal and informal sectors. The raw data record workers at the place of work, and SHRUG aggregates these counts to the village or town level and drops agricultural and government administration establishments to ensure consistency over time. Employment per 1,000 residents therefore measures the density of non-farm jobs in the local economy, rather than a conventional labour-force employment rate based on household surveys.

such as post offices and telephone connections, different types of markets, and electricity connections are also used. These indicators provide a compact description of baseline connectivity and infrastructure and are used as additional controls in the Economic Census and higher-education specifications.

The same SHRUG backbone is used to construct education outcomes. From the 1991, 2001, and 2011 Village and Town Directories, counts of secondary schools, senior secondary schools, and colleges are extracted, combining government and private institutions. Aggregating these counts to subdistricts and dividing by the corresponding Census population yields a higher-education institutions index: the number of secondary schools, senior secondary schools, and colleges per 1,000 residents. Table 1 reports summary statistics for the mutt exposure measures, baseline demographics and infrastructure, and non-farm employment outcomes across the 668 subdistrict-year observations in the sample.

Figure 2: Non-farm employment in high- and low-exposure subdistricts



Notes: The figure plots mean non-farm employment per 1,000 residents for low and high exposure subdistricts in each Economic Census year (1990, 1998, 2005, and 2013). Exposure groups are defined using quartiles of the subdistrict-level mutt exposure in 1990: “low exposure (Q2)” denotes subdistricts in the second quartile of the exposure distribution, and “high exposure (Q4)” denotes subdistricts in the top quartile. Non-farm employment per 1,000 residents is constructed by dividing total Economic Census employment by the nearest Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplying by 1,000.

Figure 2 provides a simple descriptive illustration of how non-farm employment evolves in subdistricts with low and high mutt exposure. Up to the first post-reform Economic

Table 1: Summary statistics

|                                                              | Mean    | SD     | Min    | Max     |
|--------------------------------------------------------------|---------|--------|--------|---------|
| <b>Panel A. Mutt exposure (subdistrict, 1991)</b>            |         |        |        |         |
| Mutts within 20 km (count)                                   | 1.72    | 2.95   | 0.00   | 21.00   |
| Mutt exposure per 1,000 (20 km log kernel)                   | 8.73    | 71.28  | 0.00   | 899.84  |
| Exposure (standardised)                                      | 0.00    | 1.00   | −0.12  | 12.71   |
| Log mutt exposure                                            | 3.02    | 2.74   | 0.00   | 10.26   |
| <b>Panel B. 1991 population and baseline characteristics</b> |         |        |        |         |
| Population (1991)                                            | 164,713 | 65,734 | 28,093 | 395,890 |
| Scheduled Caste share (1991)                                 | 0.179   | 0.070  | 0.028  | 0.348   |
| Scheduled Tribe share (1991)                                 | 0.052   | 0.058  | 0.000  | 0.357   |
| Literacy rate (1991)                                         | 0.399   | 0.105  | 0.141  | 0.647   |
| Child population share, age 0-6 (1991)                       | 0.173   | 0.025  | 0.122  | 0.222   |
| Terrain ruggedness index (standardised)                      | 0.00    | 1.00   | −0.93  | 4.37    |
| Transport facilities per 1,000 (1991)                        | 0.609   | 0.261  | 0.214  | 1.743   |
| Communications facilities per 1,000 (1991)                   | 0.434   | 0.154  | 0.146  | 0.990   |
| Markets per 1,000 (1991)                                     | 0.040   | 0.022  | 0.000  | 0.092   |
| Electricity connections per 1,000 (1991)                     | 0.829   | 0.436  | 0.215  | 1.936   |
| <b>Panel C. Non-farm employment (per 1,000 residents)</b>    |         |        |        |         |
| Total non-farm employment <sup>a</sup>                       | 62.35   | 44.87  | 20.78  | 314.57  |
| Manufacturing employment                                     | 22.27   | 36.56  | 2.02   | 246.08  |
| Services employment                                          | 37.80   | 12.17  | 14.99  | 95.42   |
| Private non-farm employment                                  | 51.33   | 39.95  | 18.04  | 273.23  |
| Government non-farm employment                               | 10.02   | 5.99   | 0.48   | 47.91   |

Notes: All variables are measured at the subdistrict level.  $N = 668$  subdistrict-year observations. “Mutts within 20 km (count)” is the number of mutts whose point locations lie within 20 kilometres of the subdistrict representative point. “Mutt exposure per 1,000 (20 km log kernel)” is the salience-weighted log-kernel of mutt locations within 20 kilometres, divided by the 1991 Population Census total and multiplied by 1,000. “Exposure (standardised)” is this per-capita kernel-based mutt exposure standardised to have mean zero and unit variance across subdistricts. “Log mutt exposure” is the natural logarithm of one plus the *raw* kernel weight,  $\log(1 + \text{kernel weight})$ , where “kernel weight” refers to the salience-weighted 20 kilometre log-kernel of mutt locations around the subdistrict representative point prior to any population scaling. In the robustness regressions in Table 6, this log mutt exposure is standardised and interacted with the post-reform indicator.

<sup>a</sup> Economic Census employment counts are converted to per-capita terms by dividing by the closest available Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplying by 1,000. Panel C reports averages over the four Economic Census years; the underlying employment variables are winsorised at the 99th percentile of their cross-sectional distribution. Baseline characteristics are from the 1991 Population Census and associated Village Directory modules, aggregated from SHRUG to the subdistrict level.

Census in 1998, the paths for the two groups track each other closely, with both experiencing a common decline in non-farm employment per 1,000 residents between 1990 and 1998. The divergence emerges only thereafter: between 1998 and 2005 non-farm employment rises sharply in high-exposure subdistricts while it increases more modestly in low-exposure subdistricts, and this gap remains sizable in 2013. Even without any controls, this crude visual representation foreshadows the difference-in-differences results, which show that subdistricts with greater pre-1990 mutt exposure experience substantially larger post-reform gains in non-farm employment per capita.

### 4.3 Assembly elections and constituency-level exposure

To study political competition, the analysis uses the elections module that is distributed with SHRUG. This module compiles candidate-level vote totals and turnout from the Trivedi Center for Political Data (TCPD) Jensenius and Verniers (2017) and harmonises them across elections and delimitation regimes.<sup>16</sup> For each legislative assembly constituency and election between 1978 and 2018 the data report candidate-level vote shares and basic constituency identifiers.

From these data two main indicators of electoral competition are constructed. The first is the margin of victory, defined as the absolute difference in percentage points between the winner’s and the runner-up’s vote shares. The second is “competitiveness”, defined as 100 minus the winner’s vote share in percentage points, so that higher values correspond to closer races. For completeness, the effective number of parties (ENOP), a standard measure of vote fragmentation based on candidate-level vote shares, is also constructed, and turnout is used as an additional outcome. Results for the latter are reported in Appendix A.2.

Linking mutts to constituencies requires careful treatment of delimitation. First, mutt exposure is computed at the village level. Second, the SHRUG constituency key, which provides normalised fragment weights mapping each village to pre-2008 and post-2008 assembly constituencies, is used. For each delimitation regime, village-level exposure is multiplied by these fragment weights and summed across villages to obtain an average exposure for each constituency. This procedure yields a constituency-level measure of mutt exposure that is consistent across time, uses the same 20 kilometre kernel definition as the subdistrict-level exposure.

---

<sup>16</sup>See Jensenius and Verniers (2017) and Asher et al. (2021) for details on the underlying TCPD data and the construction of the SHRUG elections module. The SHRUG elections metadata are available at <https://docs.devdatalab.org/SHRUG-Metadata/Elections/elections-metadata/>.

## 5 Empirical strategy and identification

The empirical strategy exploits cross-sectional variation in pre-1990 mutt exposure across subdistricts and compares how outcomes evolve before and after the reforms. The design is a difference-in-differences with a continuous treatment intensity.

The analysis uses three related panels. For the Economic Census, the unit of observation is a subdistrict  $c$  observed in years  $t \in \{1990, 1998, 2005, 2013\}$ , with outcomes such as non-farm employment per 1,000 residents. For higher education, subdistrict-year data for Population Census years  $t \in \{1991, 2001, 2011\}$  are used, with outcomes based on the density of secondary schools, senior secondary schools, and colleges. For elections, subdistricts are aggregated to assembly constituencies  $\ell$  observed in state elections between 1978 and 2018, and outcomes such as the margin of victory and competitiveness are studied. All panels are restricted to rural Karnataka villages and aggregated to subdistrict (or assembly constituency) for analysis.

### 5.1 Exposure and timing

The key explanatory variable throughout is the time-invariant mutt exposure constructed in Section 4.1. For each subdistrict, this summarises the salience-weighted density of mutts within 20 kilometres per 1,000 residents in 1991. It is treated as predetermined with respect to all outcomes studied and is standardised to have mean zero and unit variance across subdistricts, so that a one-unit increase corresponds to a one-standard-deviation higher level of pre-1990 mutt exposure.

The shocks exploited in the analysis are the structural and political changes of the early 1990s. For the Economic Census, 1990 provides a pre-reform baseline and the 1998, 2005, and 2013 rounds capture medium-run adjustments to liberalisation and subsequent policy changes. For higher education, the 1991 Population Census is the pre-reform baseline and the 2001 and 2011 Censuses are post-reform. For elections, the 1994 state election is treated as the first post-reform contest, reflecting both national economic reforms and the implementation of the 73rd and 74th Constitutional Amendments.

### 5.2 Baseline difference-in-differences specification

In all applications, we compare pre and post-reform periods using a common difference-in-differences structure with a continuous treatment intensity. The basic question is whether places that were more exposed to mutts before 1990 experienced systematically larger changes in outcomes after the reforms than otherwise similar, less-exposed places.

For a subdistrict-level outcome  $Y_{ct}$  (such as non-farm employment per 1,000 residents or the higher-education institutions index), the baseline specification is

$$Y_{ct} = \alpha_c + \lambda_t + \beta (\text{Exposure}_c \times \text{Post}_t) + \theta'_t Z_c + u_{ct}, \quad (1)$$

where  $c$  indexes subdistricts and  $t$  indexes time.

Each term in (1) plays a distinct role. Subdistrict fixed effects,  $\alpha_c$ , absorbs all time-invariant differences across subdistricts: initial income levels, long-run geography, deep political history, pre-existing religious and caste institutions, and any other factors that do not change over the sample period. Including  $\alpha_c$  means that identification comes from how each subdistrict changes over time relative to its own baseline, rather than from differences in levels across places.

Year fixed effects,  $\lambda_t$ , captures shocks that are common to all subdistricts in a given year, such as statewide macroeconomic conditions, national liberalisation, or changes in data collection. They allow the average outcome to shift flexibly across Economic Census or Population Census rounds. As a result, the coefficient  $\beta$  is identified from relative changes in outcomes between high-and low-exposure subdistricts, over and above any state-wide trend.

The term  $\text{Post}_t$  is an indicator for post-reform years: 1998, 2005, and 2013 for the Economic Census and 2001 and 2011 for higher education (with 1990 or 1991 as the corresponding pre-reform baseline).  $\text{Exposure}_c$  is the time-invariant pre-1990 mutt exposure index constructed in Section 4.1 and standardised to have mean zero and unit variance across subdistricts. The interaction  $\text{Exposure}_c \times \text{Post}_t$  therefore implements a difference-in-differences with a continuous treatment: it asks whether, after 1990, subdistricts that were one standard deviation more exposed to mutts experienced larger changes in  $Y_{ct}$  than those that were less exposed.

The vector  $Z_c$  collects 1991 baseline characteristics: the shares of the population who are Scheduled Castes and Scheduled Tribes, the literacy rate, a standardised ruggedness index, and per-capita measures of transport facilities, postal and telephone infrastructure, markets, and electricity. We interact these with year dummies through  $\theta'_t Z_c$ , allowing each baseline characteristic to have its own time path. This relaxes the parallel trends assumption by permitting, for example, more literate or better connected subdistricts to grow faster (or slower) for reasons unrelated to mutts. In other words,  $\theta'_t Z_c$  soaks up systematic differential trends that are correlated with initial composition and infrastructure, so that  $\beta$  is identified from residual variation in changes that is not explained by these factors.

Standard errors are clustered at the subdistrict level to allow for arbitrary serial correlation and heteroskedasticity within subdistricts over time.

For the elections analysis we use an analogous specification at the constituency level:

$$E_{\ell t} = \alpha_{\ell} + \lambda_t + \gamma (\text{Exposure}_{\ell} \times \text{Post94}_t) + \psi'_t Z_{\ell} + \eta_{\ell t}, \quad (2)$$

where  $E_{\ell t}$  is an electoral outcome for constituency  $\ell$  in year  $t$  (such as the margin of victory or competitiveness),  $\alpha_{\ell}$  and  $\lambda_t$  are constituency and year fixed effects, and  $\text{Post94}_t$  is an indicator for elections from 1994 onwards. The vector  $Z_{\ell}$  contains 1991 constituency covariates (caste composition, literacy, ruggedness), again interacted with year dummies to allow these initial characteristics to have flexible time trends. Standard errors are clustered at the constituency level.

To see what identifies  $\beta$  in (1), it is useful to difference between a pre-reform year and the post-reform period within each subdistrict. With a single post indicator, the change in the outcome in subdistrict  $c$  can be written as

$$\Delta Y_c = (\bar{Y}_{c,\text{post}} - Y_{c,\text{pre}}) = \beta \text{Exposure}_c + \Delta \theta' Z_c + \Delta u_c,$$

where all time-invariant unobservables are absorbed by  $\alpha_c$  and common shocks are absorbed by the year effects. The coefficient  $\beta$  is therefore identified from *cross-subdistrict differences in the change in  $Y_{ct}$*  between the pre and post-reform periods as a function of pre-1990 mutt exposure, conditional on the baseline characteristics  $Z_c$  and their differential trends. Intuitively, the design asks whether subdistricts that happened to be more exposed to mutts before 1990 experienced systematically larger post-reform changes in outcomes than less exposed subdistricts with similar initial composition, geography, and infrastructure.

### 5.3 Identification assumptions and threats

The key identifying assumption is a continuous-treatment version of parallel trends: conditional on baseline covariates and their flexible time trends, subdistricts with higher and lower pre-1990 mutt exposure would have experienced the same average change in outcomes between the pre- and post-reform periods in the absence of mutts.

For the Economic Census panel in particular, there is only one true pre-reform round (1990), so we cannot estimate a full set of pre-trend coefficients. The dynamic specifications below, therefore, use 1998, the first post-reform round, as an early post benchmark and check that there is essentially no exposure gradient at that point. This is a limited diagnostic rather than a formal pre-trends test. A similar limitation applies to the

education panel, where 1991 is the only pre-reform Census: the 2001 and 2011 coefficients trace out post-reform dynamics relative to that single baseline rather than a full pre-trend.

As discussed above, the fixed effects absorb all time-invariant differences across subdistricts, and the year dummies capture shocks common to all subdistricts in a given year. Interactions between year dummies and the 1991 composition and infrastructure variables allow for differential trends associated with caste composition, literacy, physical geography, demographic structure, and baseline connectivity. Mutt exposure is built from mutts that pre-date the 1990 Economic Census, whereas the outcomes studied are measured in 1990 and later. Beyond these baseline safeguards, three classes of threats are particularly salient:

(i) *Placement of mutts*: Mutts could be systematically located in subdistricts that already had different long-run development potential, for instance, more rugged terrain, better transport links, or different caste composition. We address this in three ways. First, the rich control set  $Z_c$  flexibly absorbs any time-varying effects of 1991 caste composition, literacy, ruggedness, demographic structure, and infrastructure on later outcomes. Second, Appendix A.1 reports cross-sectional balance regressions showing that, although mutt exposure is correlated with some observables, these correlations are modest and explain less than five percent of the variation in exposure. Third, robustness checks in Section 6.4 show that the Economic Census results are stable to alternative constructions of mutt exposure and, to a lesser degree, to trimming the upper tail of the exposure distribution.

(ii) *Pre-existing trends*: High-exposure subdistricts might have been on different trajectories even in the absence of mutts. To probe this, we estimate dynamic versions of the preferred specification (1) in which the single post indicator is replaced with interactions between exposure and separate year dummies (Section 6.3). For non-farm employment, the exposure-year coefficient in 1998 is essentially zero, while large and precisely estimated differences emerge only later in 2005 and 2013. For higher-education institutions, there is no exposure gradient in 1991 by construction, a modest but significant gap by 2001, and a larger gap by 2011. These patterns are difficult to reconcile with simple pre-existing trend stories, and are plausibly consistent with mutt exposure shaping how subdistricts respond to reforms over time.

(iii) *Comparison outcomes*: As an additional check, we look at outcomes that are plausibly less directly shaped by mutts. In the Economic Census, the coefficient on exposure for government employment per capita is positive but considerably smaller than the effects for manufacturing and private employment, and the government employment share changes only modestly. Likewise, the estimated effect on services employment is positive but much smaller than for manufacturing and private-sector jobs. These patterns

suggest that the mutt gradient is concentrated in the parts of the economy where mutt-run schools, hostels and networks are most plausibly useful for starting and expanding enterprises, rather than reflecting a uniform shift in all outcomes.

Although we cannot fully rule out selection on unobservables, any alternative explanation would have to involve pre-1990 factors that are (i) strongly correlated with mutt exposure conditional on observed composition, ruggedness, and infrastructure, (ii) have little effect on outcomes in the pre- or early post-reform period, yet (iii) generate large relative gains in non-farm employment, higher-education capacity, and electoral competitiveness only from the mid-1990s onward. While such factors are possible in principle, the combination of historical evidence on mutts and the timing and sectoral pattern of the results makes mutts themselves a particularly plausible driver of these differential changes.

## 6 Results

This section documents how pre-1990 exposure to mutts is associated with subsequent patterns of structural change in Karnataka. We proceed in four steps.

First, we establish the main result: subdistricts with higher pre-1990 mutt exposure experience substantially larger increases in non-farm employment per capita between 1990 and the post-reform Economic Census rounds (1998, 2005, 2013). Second, we show that these gains are concentrated in manufacturing and private-sector jobs and emerge only after the 1990s reforms. Third, we present a set of robustness checks on the Economic Census results. Finally, we examine two mechanisms that can help explain the employment gains: the expansion of secondary and tertiary education and changes in local political competition.

Throughout, the unit of analysis is the subdistrict unless otherwise noted. Standard errors are clustered at the subdistrict (or constituency) level. All exposure measures are defined per 1,000 residents using the 1991 population and then standardised to have mean zero and unit variance so that coefficients can be interpreted as the effect of a one standard-deviation increase in exposure.

### 6.1 Mutts and non-farm employment

We begin by relating mutt exposure to the growth of non-farm employment per capita, using the four Economic Census rounds (1990, 1998, 2005, and 2013). Non-farm employment per 1,000 residents is constructed as described in Section 4.2 and denoted by  $Y_{ct}$ .

Using the standardised pre-1990 exposure index from Section 4.1 and the baseline

difference-in-differences structure in equation (1), the main estimating equation for non-farm employment can be written as

$$Y_{ct} = \alpha_c + \lambda_t + \beta (\text{Exposure}_c \times \text{Post}_t) + \theta'_t Z_c + u_{ct}, \quad (3)$$

where  $\alpha_c$  and  $\lambda_t$  are subdistrict and year fixed effects,  $Z_c$  collects 1991 baseline characteristics, and  $\theta_t$  is allowed to vary by year. Here  $\text{Post}_t$  is the post-reform indicator that equals one for the 1998, 2005, and 2013 Economic Census rounds and zero for 1990. The coefficient of interest,  $\beta$ , measures the change in non-farm employment per 1,000 residents between 1990 and the post-reform period that is associated with a one-standard-deviation increase in mutt exposure. Table 2 reports estimates from variants of (3) with progressively richer sets of controls.

Column (1) reports estimates without any time-varying controls beyond fixed effects. The coefficient on the interaction between the post period and mutt exposure is 11.7 jobs per 1,000 residents. Moving from a subdistrict at the mean level of exposure to one a standard deviation above the mean is thus associated with an increase of about 12 non-farm jobs per 1,000 residents over the post-1990 period relative to 1990, a 19 percent gain compared with the 1990 mean of 61.6 jobs per 1,000.

Column (2) adds flexible time trends in the 1991 Scheduled Caste share, Scheduled Tribe share, and literacy share; the estimated effect remains large and precisely estimated at 10.2 jobs per 1,000 residents. Column (3) further allows ruggedness and 1991 infrastructure variables to have their own year-specific slopes. In this full specification, which we take as the preferred one, the coefficient on the post-period exposure interaction is 10.0 jobs per 1,000 residents. Thus, even after allowing subdistricts with different caste composition, literacy, physical geography, and baseline infrastructure to follow different time trends, a one standard-deviation difference in pre-1990 mutt exposure is associated with roughly 10 extra non-farm jobs per 1,000 residents in the post-1990 period.

To make this magnitude concrete, note that the average subdistrict in the sample had a population of about 170,000 people in 1991. The preferred estimate in Column (3) therefore implies that moving from a subdistrict at the mean exposure to one standard deviation above the mean is associated with roughly  $10.0 \times 170 \approx 1,700$  additional non-farm jobs. These are large, economically meaningful differences in the context of the 1990s and 2000s Karnataka economy, and they arise after netting out both subdistrict fixed effects and common year shocks, as well as rich interactions of initial composition and infrastructure with time.

A natural concern is that the baseline exposure index uses contemporary Google

Table 2: Mutt exposure and non-farm employment per 1,000 residents

|                                                | (1)<br>Emp         | (2)<br>Emp         | (3)<br>Emp         |
|------------------------------------------------|--------------------|--------------------|--------------------|
| Post $\times$ exposure                         | 11.74***<br>(1.96) | 10.22***<br>(1.79) | 10.03***<br>(1.80) |
| Subdistrict fixed effects                      | Yes                | Yes                | Yes                |
| Year fixed effects                             | Yes                | Yes                | Yes                |
| 1991 SC/ST/literacy $\times$ year <sup>a</sup> | No                 | Yes                | Yes                |
| Ruggedness $\times$ year                       | No                 | No                 | Yes                |
| 1991 infrastructure $\times$ year <sup>b</sup> | No                 | No                 | Yes                |
| Observations                                   | 668                | 668                | 668                |
| Subdistricts                                   | 167                | 167                | 167                |
| $R^2$                                          | 0.52               | 0.54               | 0.55               |
| Mean of $Y_{c,1990}$                           | 61.6               | 61.6               | 61.6               |

Notes: The dependent variable is non-farm employment per 1,000 residents in a subdistrict, measured in Economic Census years 1990, 1998, 2005, and 2013. For each year, Economic Census employment counts are divided by the closest available Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplied by 1,000. The mean level of the dependent variable in 1990 is 61.6 jobs per 1,000 residents. The main regressor is the interaction between a post-reform indicator (1998, 2005, 2013) and mutt exposure measured using 1991 information. Exposure is constructed as a 20 kilometre log kernel of mutts around the subdistrict centroid, weighted by mutt salience and scaled by the 1991 population, then standardised to mean zero and unit variance. Column (1) includes subdistrict and year fixed effects only. Column (2) adds interactions between year dummies and the 1991 Scheduled Caste share, Scheduled Tribe share, and literacy share. Column (3) further adds interactions between year dummies and a standardised measure of terrain ruggedness and 1991 per-capita infrastructure variables (transport facilities, postal and telephone infrastructure, markets, and electricity). Standard errors, in parentheses, are clustered at the subdistrict level. *Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

<sup>a</sup> “1991 SC/ST/literacy  $\times$  year” denotes interactions between year dummies and the 1991 shares of Scheduled Castes, Scheduled Tribes, and literate population.

<sup>b</sup> “1991 infrastructure  $\times$  year” denotes interactions between year dummies and 1991 per-capita measures of transport, communications, markets, and electricity.

rating counts as salience weights, which are not themselves predetermined. Section 6.4.1 therefore replaces the salience-weighted kernel with a conceptually simpler exposure based on the unweighted count of mutts within 20 kilometres per 1,000 residents. The resulting coefficients are very similar in magnitude and remain statistically significant, indicating that the core Employment Census finding does not hinge on the use of Google rating weights.

## 6.2 How mutt exposure reshapes non-farm employment

The next step is to ask how the additional non-farm jobs in high-exposure subdistricts are distributed across sectors. Equation (3) is estimated separately for total non-farm employment and for employment in manufacturing, services, and the private sector, using the same full set of controls as in Column (3) of Table 2. That is, all regressions include subdistrict and year fixed effects, and interactions between year dummies and the 1991 composition and infrastructure variables. Table 3 reports the resulting coefficients on the post-period exposure interaction.

The estimates in Table 3 indicate that mutt exposure is associated with a substantial increase in non-farm employment per capita in the post-reform period, and that this increase is concentrated in manufacturing and private-sector jobs. Column (1) reproduces the total-employment from Table 2: a one standard-deviation increase in exposure is associated with about 10 additional non-farm jobs per 1,000 residents, relative to the 1990 baseline, after conditioning on fixed effects and flexible trends in 1991 characteristics and infrastructure.

Column (2) shows that roughly eight of these additional jobs per 1,000 residents arise in manufacturing. Column (4) shows a similarly large response in private-sector employment, with about 8.9 additional private non-farm jobs per 1,000 residents. The estimated effect on services employment in Column (3) is positive but smaller, around 1.5 jobs per 1,000 residents.

Overall, subdistricts with greater mutt exposure do not simply expand all forms of non-farm activity proportionally. Rather, they experience disproportionate gains in manufacturing and private-sector jobs. This is consistent with the idea that mutts help intermediate the transition into higher-productivity non-farm activities by providing local public goods, informal credit, and political connections that are particularly valuable for industrial and entrepreneurial ventures, rather than by expanding public-sector employment or low-productivity services.

Table 3: Mutts and non-farm employment by sector (Economic Census)

|                                  | (1)                | (2)               | (3)               | (4)               | (5)               |
|----------------------------------|--------------------|-------------------|-------------------|-------------------|-------------------|
|                                  | Total emp          | Manuf emp         | Services emp      | Private emp       | Government emp    |
| Post $\times$ exposure           | 10.03***<br>(1.80) | 7.99***<br>(1.56) | 1.54***<br>(0.28) | 8.90***<br>(1.64) | 0.60***<br>(0.20) |
| Subdistrict fixed effects        | Yes                | Yes               | Yes               | Yes               | Yes               |
| Year fixed effects               | Yes                | Yes               | Yes               | Yes               | Yes               |
| Additional controls <sup>a</sup> | Yes                | Yes               | Yes               | Yes               | Yes               |
| Observations                     | 668                | 668               | 668               | 668               | 668               |
| Subdistricts                     | 167                | 167               | 167               | 167               | 167               |
| Mean in 1990                     | 61.6               | 10.1              | 41.6              | 53.8              | 10.2              |

Notes: Each column reports a separate difference-in-differences regression at the subdistrict level using the Economic Census panel for 1990, 1998, 2005, and 2013. The dependent variables are sectoral non-farm employment per 1,000 residents (total, manufacturing, services, private, and government), constructed as in Table 2. “Exposure” is the standardised pre-1990 mutt exposure index, and “Post” is an indicator for Economic Census years 1998, 2005, and 2013. All regressions include subdistrict fixed effects and year fixed effects. Standard errors, in parentheses, are clustered at the subdistrict level.

<sup>a</sup>Additional controls are interactions between year dummies and the 1991 Scheduled Caste and Scheduled Tribe shares, literacy rate, a standardised ruggedness index, and 1991 per-capita measures of transport, postal and telephone facilities, markets, and electricity.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

### 6.3 Timing of the employment gains

To study when these gains arise and to probe the identifying assumptions, a dynamic version of (3) is estimated in which the single post-reform interaction is replaced with separate interactions between exposure and year dummies for 1998, 2005, and 2013. Specifically,

$$Y_{ct} = \alpha_c + \lambda_t + \delta_{1998}(\text{Exposure}_c \times \mathbf{1}\{t = 1998\}) + \delta_{2005}(\text{Exposure}_c \times \mathbf{1}\{t = 2005\}) + \delta_{2013}(\text{Exposure}_c \times \mathbf{1}\{t = 2013\}) + \theta'_t Z_c + u_{ct}, \quad (4)$$

where the omitted year is the 1990 Economic Census. Subdistrict and year fixed effects, together with the same full set of year-specific baseline controls used in Column (3) of Table 2, are included. Table 4 reports the estimated coefficients on the exposure-year interactions.

Table 4: Timing of mutt exposure effects on non-farm employment

|                                  | Total non-farm employment per 1,000 |
|----------------------------------|-------------------------------------|
| Exposure $\times$ 1{Year = 1998} | −0.94<br>(1.15)                     |
| Exposure $\times$ 1{Year = 2005} | 17.57***<br>(2.60)                  |
| Exposure $\times$ 1{Year = 2013} | 13.40***<br>(2.80)                  |
| Subdistrict fixed effects        | Yes                                 |
| Year fixed effects               | Yes                                 |
| Additional controls <sup>a</sup> | Yes                                 |
| Observations                     | 668                                 |
| Subdistricts                     | 167                                 |

Notes: The dependent variable is total non-farm employment per 1,000 residents, constructed as in Table 2. The table reports coefficients on interactions between standardised pre-1990 mutt exposure and year dummies for 1998, 2005, and 2013; 1990 is the omitted year. All regressions include subdistrict fixed effects, year fixed effects, and the full set of baseline covariates interacted with year dummies. Standard errors, in parentheses, are clustered at the subdistrict level.

<sup>a</sup>Baseline covariates are the 1991 Scheduled Caste and Scheduled Tribe shares, literacy rate, a standardised ruggedness index, and 1991 per-capita measures of transport, communications, markets, and electricity.

Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

The estimate for 1998 is essentially zero: a one standard-deviation increase in mutt exposure is associated with only −0.9 additional non-farm jobs per 1,000 residents in 1998, and an  $F$ -test for  $\delta_{1998} = 0$  yields a  $p$ -value of about 0.42. By contrast, the coefficients for

2005 and 2013 are large and precisely estimated: exposure is associated with about 18 additional jobs per 1,000 residents in 2005 and about 13 additional jobs per 1,000 in 2013, relative to 1990.

This result is not what one would expect if mutt exposure simply captured pre-existing differences in trends. In the first post-reform Economic Census (1998), high-exposure and low-exposure subdistricts look very similar in terms of non-farm employment growth relative to 1990. The divergence appears only between 1998 and 2005 and remains sizable in 2013. Combined with the fixed-effects structure of equation (3), the absence of an early-post difference supports a difference-in-differences interpretation in which pre-1990 mutt exposure shapes how subdistricts respond to post-1990 reforms, rather than proxying for unobserved upward trends.

To gauge how mutt exposure relates to pre-reform characteristics, we run cross-sectional regressions of each 1991 baseline variable: the Scheduled Caste and Scheduled Tribe shares, literacy, our ruggedness measure, and standardised per-capita measures of markets, transport facilities, and communications on the exposure index. Appendix Table A.1 shows that higher-exposure subdistricts have slightly lower Scheduled Caste shares, somewhat higher literacy, more rugged terrain, and better transport and communications, but weaker market infrastructure, while the Scheduled Tribe share is essentially unrelated to exposure. Quantitatively, a one standard-deviation increase in exposure is associated with about a 0.6 percentage-point lower Scheduled Caste share (around 3 percent of its mean), roughly a 2 percentage-point higher literacy rate (around 5 percent of its mean), and increases of about 0.21 standard deviations in ruggedness, transport and communications, together with a decline of about 0.15 standard deviations in market access. All  $R^2$  values are below five percent, however, indicating that these observables explain only a small fraction of the cross-sectional variation in exposure. These modest but non-trivial correlations motivate the rich control set in the baseline regressions, which allows 1991 composition, ruggedness, and infrastructure to have flexible year-specific slopes.

## 6.4 Robustness of the Economic Census results

### 6.4.1 Salience weights and unweighted exposure

A natural concern with the baseline exposure measure is that the salience weights are based on contemporary Google rating counts. These ratings are observed long after the 1990 pre-reform baseline and may themselves reflect post-1990 developments such as smartphone penetration, tourism, or modern economic activity. If so, one might worry that the results conflate the long-run influence of mutts with ex post visibility in the

digital era.

To address this, we construct an alternative, purely unweighted raw count of mutts. Specifically, this is the “Mutts within 20 km (count)” variable in Panel A of Table 1. We standardize the measure and follow the same specification as done previously to re-estimate the Economic Census difference-in-differences specification in equation (3), using the same sequence of control sets as in Table 2.

Table 5 reports the results. Column (1) includes only subdistrict and year fixed effects. Column (2) adds interactions between year dummies and the 1991 Scheduled Caste share, Scheduled Tribe share, and literacy rate. Column (3) further allows 1991 ruggedness and infrastructure (per-capita transport, communications, markets, and electricity) to have flexible year-specific slopes, mirroring the preferred specification in Table 2.

The unweighted specifications yield effects that are very similar in magnitude to the baseline salience-weighted results and remain statistically significant at conventional levels. In the preferred full-controls specification in Column (3), a one standard-deviation increase in the raw (unweighted) mutt count within 20 kilometres is associated with 8.7 additional non-farm jobs per 1,000 residents in the post-reform period, relative to 1990. Given a 1990 mean of 61.6 jobs per 1,000 residents, this corresponds to roughly a 14 percent higher level of non-farm employment. For an average subdistrict with about 170,000 residents in 1991, the point estimate implies on the order of  $8.65 \times 170 \approx 1,470$  additional non-farm jobs.

Columns (1) and (2) tell a similar story: depending on the control set, the implied gains range from about 8.8 to 10.4 jobs per 1,000 residents, which is very close to the 10.0 jobs per 1,000 residents estimated with the salience-weighted kernel in Table 2. The fact that the estimated effects are stable in size and precision when mutt exposure is based purely on physical counts, without any Google-ratings weights, suggests that the Economic Census results are not an artefact of ex post salience proxies. Instead, it is the underlying spatial distribution of mutts, rather than their contemporary online visibility, that appears to be doing the work.

#### 6.4.2 Alternative exposure measures

A natural concern is that the main results might be driven by the particular way in which mutt exposure is constructed. The baseline measure combines (i) a 20 kilometre kernel around the subdistrict centroid, which allows mutts just across an administrative boundary to influence a subdistrict, and (ii) salience weights that give more importance to mutts with many Google ratings. If the estimated effects disappear when mutt influence is restricted strictly within subdistrict boundaries, or when the influence of highly salient

Table 5: Non-farm employment and unweighted mutt exposure per 1,000 residents

|                                                | (1)<br>Emp        | (2)<br>Emp       | (3)<br>Emp       |
|------------------------------------------------|-------------------|------------------|------------------|
| Post $\times$ Mutts Raw Exposure               | 10.42**<br>(4.68) | 8.83**<br>(4.32) | 8.65**<br>(4.34) |
| Subdistrict fixed effects                      | Yes               | Yes              | Yes              |
| Year fixed effects                             | Yes               | Yes              | Yes              |
| 1991 SC/ST/literacy $\times$ year <sup>a</sup> | No                | Yes              | Yes              |
| Ruggedness $\times$ year                       | No                | No               | Yes              |
| 1991 infrastructure $\times$ year <sup>b</sup> | No                | No               | Yes              |
| Observations                                   | 668               | 668              | 668              |
| Subdistricts                                   | 167               | 167              | 167              |
| $R^2$                                          | 0.52              | 0.54             | 0.55             |
| Mean of $Y_{c,1990}$                           | 61.6              | 61.6             | 61.6             |

Notes: The dependent variable is non-farm employment per 1,000 residents in a subdistrict, measured in Economic Census years 1990, 1998, 2005, and 2013. For each year, Economic Census employment counts are divided by the closest available Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplied by 1,000. The mean level of the dependent variable in 1990 is 61.6 jobs per 1,000 residents. “Mutts Raw Exposure” is constructed as the number of mutts whose point locations lie within 20 kilometres of the subdistrict centroid, divided by the 1991 population and multiplied by 1,000, and then standardised to have mean zero and unit variance across subdistricts. The reported coefficients are on the interaction between this unweighted exposure index and a post-reform indicator for Economic Census years 1998, 2005, and 2013.

<sup>a</sup> “1991 SC/ST/literacy  $\times$  year” denotes interactions between year dummies and the 1991 shares of Scheduled Castes, Scheduled Tribes, and literate population.

<sup>b</sup> “1991 infrastructure  $\times$  year” denotes interactions between year dummies and 1991 per-capita measures of transport facilities, communications (postal and telephone) infrastructure, markets, and electricity. All specifications include subdistrict and year fixed effects. Standard errors, in parentheses, are clustered at the subdistrict level.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

mutts is compressed, this would cast doubt on the interpretation of the exposure gradient.

To probe this, we re-estimate the full Economic Census specification using three alternative pre-1990 exposure measures, holding the control set fixed. All three measures are defined per 1,000 residents using the 1991 population, and are then standardised to have mean zero and unit variance across subdistricts, so the coefficients are always interpretable as the effect of a one-standard-deviation difference in exposure. Table 6 reports the results.

Column (1) uses a within-boundary measure that sums the salience-weighted number of mutts whose point locations fall inside the subdistrict and divides by the 1991 population. This forces mutt influence to be purely local and ignores any spillovers from mutts in neighbouring subdistricts. Column (2) uses the measure used in the main specifications above. Column (3) applies a log transformation to the underlying kernel weight, constructing a per-capita index based on  $\log(1 + \text{kernel weight}_c)$  before standardisation. This compresses the right tail and reduces the influence of subdistricts that lie very close to a small number of highly salient mutts, placing more weight on proportional than on absolute differences in exposure.

The three measures yield very similar patterns. A one-standard-deviation increase in exposure is associated with about 8.9 additional non-farm jobs per 1,000 residents when mutt influence is restricted strictly within subdistrict boundaries (Column 1), about 10.0 additional jobs when using the 20 kilometre kernel (Column 2), and about 12.6 additional jobs when the log-transformed kernel is used (Column 3). Relative to the 1990 mean of 61.6 non-farm jobs per 1,000 residents, these estimates correspond to roughly 15–20 percent higher post-reform non-farm employment in more exposed subdistricts.

The comparison between Columns (1) and (2) shows that allowing mutts just across the border to affect a subdistrict is not crucial: the estimated effect is slightly larger with the kernel but of the same order of magnitude when exposure is defined solely within administrative boundaries. The comparison between Columns (2) and (3) shows that down-weighting very salient mutts through the log transformation also leaves the basic result intact. Taken together, these robustness checks suggest that the positive association between mutt exposure and non-farm employment is not an artefact of the specific geographic kernel or the linear use of salience weights.

### 6.4.3 Outliers in mutt exposure

A further concern is that the Economic Census results may be driven by a small number of subdistricts with extremely high mutt exposure. To examine this, the underlying 20

Table 6: Alternative exposure measures and non-farm employment per 1,000 residents

|                                  | (1)<br>Emp        | (2)<br>Emp         | (3)<br>Emp         |
|----------------------------------|-------------------|--------------------|--------------------|
| Post $\times$ exposure           | 8.90***<br>(1.03) | 10.03***<br>(1.80) | 12.61***<br>(3.95) |
| Subdistrict fixed effects        | Yes               | Yes                | Yes                |
| Year fixed effects               | Yes               | Yes                | Yes                |
| Additional controls <sup>a</sup> | Yes               | Yes                | Yes                |
| Observations                     | 668               | 668                | 668                |
| Subdistricts                     | 167               | 167                | 167                |
| Mean of $Y_{c,1990}$             | 61.6              | 61.6               | 61.6               |

Notes: The dependent variable is total non-farm employment per 1,000 residents in a subdistrict, constructed as in Table 2 by dividing Economic Census employment counts in 1990, 1998, 2005, and 2013 by the closest Population Census total (1991 for 1990, 2001 for 1998 and 2005, and 2011 for 2013) and multiplying by 1,000. Each column uses a different standardised pre-1990 mutt exposure measure interacted with a post-reform indicator for Economic Census years 1998, 2005, and 2013. In Column (1), exposure is the salience-weighted count of mutts whose point locations lie within the subdistrict boundary, divided by the 1991 population and multiplied by 1,000. In Column (2), exposure is the 20 kilometre log-kernel of mutt locations around the subdistrict centroid, divided by the 1991 population and multiplied by 1,000. In Column (3), exposure is a log-transformed kernel measure, constructed by applying  $\log(1 + \text{kernel weight}_c)$  to the raw kernel weight around the centroid, dividing by the 1991 population, and then standardising. All exposure measures are standardised to have mean zero and unit variance across subdistricts.

<sup>a</sup>Additional controls are interactions between year dummies and the 1991 Scheduled Caste and Scheduled Tribe shares, literacy rate, a standardised ruggedness index, and 1991 per-capita measures of transport, postal and telephone facilities, markets, and electricity. All specifications include subdistrict and year fixed effects. Standard errors, in parentheses, are clustered at the subdistrict level.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

kilometre exposure per 1,000 residents is used to drop all subdistrict–year observations above the 95th percentile of this raw distribution. On this trimmed sample we re-estimate our preferred specification. Table 7 compares the full–sample and trimmed–sample results. All regressions include subdistrict and year fixed effects and the same baseline controls as in Column (3) of Table 2.

Table 7: Outlier robustness: dropping the top 5% of mutt exposure

|                                  | (1)<br>Emp<br>Full sample | (2)<br>Emp<br>$\leq$ p95 exposure |
|----------------------------------|---------------------------|-----------------------------------|
| Post $\times$ exposure           | 10.03***<br>(1.80)        | 6.11*<br>(3.40)                   |
| Subdistrict fixed effects        | Yes                       | Yes                               |
| Year fixed effects               | Yes                       | Yes                               |
| Additional controls <sup>a</sup> | Yes                       | Yes                               |
| Observations                     | 668                       | 636                               |
| Subdistricts                     | 167                       | 159                               |

Notes: The dependent variable is total non-farm employment per 1,000 residents, constructed as in Table 2. Column (1) reproduces the baseline specification using the standardised 20 kilometre log-kernel exposure measure. Column (2) drops subdistricts in the top 5 percent of the raw exposure distribution, re-standardises exposure on the trimmed sample, and re-estimates the same regression. All specifications include subdistrict fixed effects and year fixed effects. Standard errors, in parentheses, are clustered at the subdistrict level.

<sup>a</sup>Additional controls are interactions between year dummies and the 1991 Scheduled Caste and Scheduled Tribe shares, literacy rate, a standardised ruggedness index, and 1991 per-capita measures of transport, postal and telephone facilities, markets, and electricity.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

In the full sample, a one standard–deviation increase in exposure is associated with about 10.0 additional non-farm jobs per 1,000 residents. After trimming the top 5 percent of exposure values, the point estimate remains positive and large at 6.1 jobs per 1,000 residents, but the standard error rises to about 3.4 and the coefficient is significant only at the 10 percent level. The trimming exercise therefore reduces statistical precision, as one would expect when discarding the most exposed subdistricts: the point estimate falls from about 10.0 to 6.1 jobs per 1,000 residents and is estimated less precisely and is significant at the 10 percent level. The effect remains positive and economically meaningful, but the loss of precision highlights that a substantial share of the identifying variation comes from subdistricts with very strong mutt presence.

#### 6.4.4 Log specifications for outcomes

For each outcome  $Y_{ct}$  (defined in Section 4.2) the following specification is estimated:

$$\log(1 + Y_{ct}) = \alpha_c + \lambda_t + \beta^{\log}(\text{Exposure}_c \times \text{Post}_t) + \theta'_t Z_c + u_{ct},$$

including the same subdistrict and year fixed effects and full baseline control set, and clustering standard errors at the subdistrict level. The log transformation reduces the influence of very high employment levels and allows effects to be interpreted in proportional rather than purely additive terms. Table 8 reports the coefficients on  $\text{Post} \times \text{exposure}$  for total non-farm employment and its components.

Table 8: Robustness: log specifications for Economic Census employment

|                                  | (1)<br>Log emp      | (2)<br>Log manuf emp | (3)<br>Log services emp | (4)<br>Log private emp | (5)<br>Log gov emp |
|----------------------------------|---------------------|----------------------|-------------------------|------------------------|--------------------|
| Post $\times$ exposure           | 0.072***<br>(0.016) | 0.114***<br>(0.024)  | 0.029***<br>(0.009)     | 0.071***<br>(0.017)    | 0.036**<br>(0.015) |
| Subdistrict fixed effects        | Yes                 | Yes                  | Yes                     | Yes                    | Yes                |
| Year fixed effects               | Yes                 | Yes                  | Yes                     | Yes                    | Yes                |
| Additional controls <sup>a</sup> | Yes                 | Yes                  | Yes                     | Yes                    | Yes                |
| Observations                     | 668                 | 668                  | 668                     | 668                    | 668                |
| Subdistricts                     | 167                 | 167                  | 167                     | 167                    | 167                |

Notes: Each column reports a separate difference-in-differences regression using Economic Census data for 1990, 1998, 2005, and 2013. The dependent variables are  $\log(1 + Y_{ct})$ , where  $Y_{ct}$  is the relevant non-farm employment measure per 1,000 residents (total, manufacturing, services, private, or government), defined as in Table 3. “Exposure” is the standardised pre-1990 mutt exposure index, interacted with a post-reform indicator for Economic Census years 1998, 2005, and 2013. All specifications include subdistrict fixed effects and year fixed effects. Standard errors, in parentheses, are clustered at the subdistrict level.

<sup>a</sup>Additional controls are interactions between year dummies and the 1991 Scheduled Caste and Scheduled Tribe shares, literacy rate, a standardised ruggedness index, and 1991 per-capita measures of transport, postal and telephone facilities, markets, and electricity.

Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Because  $\beta^{\log}$  is a semi-elasticity and mutt exposure is standardised, the coefficients in Table 8 can be interpreted as approximate percentage differences in the outcome in the post-reform period between subdistricts that differ by one standard deviation in pre-1990 mutt exposure, relative to 1990. The estimates imply that a one standard-deviation increase in exposure is associated with about a 7 percent higher level of total non-farm employment per capita, about a 12 percent higher level of manufacturing employment per capita, and roughly a 7 percent higher level of private employment per capita, with smaller and less precisely estimated effects on services and government employment. These log-specification results are fully consistent with the level estimates and confirm that the core Economic Census findings are not an artefact of the particular scaling of the outcome. These semi-elasticities are somewhat smaller than the effects implied by the level

specifications (for example, 10 additional non-farm jobs per 1,000 residents corresponds to roughly a 15–20 percent increase relative to the 1990 mean), so the log specifications can be viewed as providing conservative proportional estimates of the mutt exposure gradient.

## 6.5 Mutts and higher-education institutions

The Economic Census results show that subdistricts with greater mutt exposure experience larger gains in non-farm employment per capita. A natural question is whether these areas also accumulate more educational capacity that could support such a transition. The next step is therefore to examine how mutt exposure relates to the expansion of upper-secondary and tertiary education.

We measure local educational capacity using the higher-education institutions index ( $\text{HEI}_{ct}$ ) defined in Section 4: the number of secondary and senior secondary schools and colleges per 1,000 residents in subdistrict  $c$  at Census year  $t$ , excluding primary schools because enrollment is nearly universal by 2011 and the mechanism of interest concerns upper-secondary and tertiary education.

The same measure of mutt exposure  $\text{Exposure}_c$  as in the Economic Census analysis is used. Let  $\text{Post}_t$  be an indicator equal to one for Census years 2001 and 2011 and zero for 1991. The main specification is

$$\text{HEI}_{ct} = \alpha_c + \lambda_t + \beta (\text{Exposure}_c \times \text{Post}_t) + \theta'_t Z_c + u_{ct}, \quad (5)$$

where  $Z_c$  collects the 1991 baseline characteristics (caste composition, literacy, terrain, and infrastructure) as we have previously used, each interacted with year dummies. Equation (5) is estimated with subdistrict and year fixed effects, clustering standard errors at the subdistrict level.

Table 9: Mutt exposure and higher-education institutions per 1,000 residents

|                                   | (1)<br>HEI          | (2)<br>HEI          | (3)<br>HEI          |
|-----------------------------------|---------------------|---------------------|---------------------|
| Post $\times$ exposure            | 0.013***<br>(0.002) | 0.011***<br>(0.002) | 0.010***<br>(0.003) |
| Subdistrict fixed effects         | Yes                 | Yes                 | Yes                 |
| Year fixed effects                | Yes                 | Yes                 | Yes                 |
| 1991 SC/ST/literacy $\times$ year | No                  | Yes                 | Yes                 |
| Ruggedness $\times$ year          | No                  | Yes                 | Yes                 |
| Infrastructure $\times$ year      | No                  | No                  | Yes                 |
| Observations                      | 528                 | 528                 | 528                 |
| Subdistricts                      | 176                 | 176                 | 176                 |
| Mean of HEI in 1991               | 0.11                | 0.11                | 0.11                |

Notes: The dependent variable is the higher-education institutions index (HEI) per 1,000 residents, as defined in Section 3.2. Exposure is standardised mutt exposure per 1,000 residents. “Post” is an indicator for Census years 2001 and 2011. The reported coefficient is on the interaction between exposure and Post and can be interpreted as the change in HEI per 1,000 residents associated with a one standard-deviation higher exposure in the post period relative to 1991. Column (1) includes only subdistrict and year fixed effects. Column (2) additionally includes interactions between year dummies and the 1991 shares of Scheduled Castes and Scheduled Tribes, the literacy rate, and a standardised ruggedness index. Column (3) further adds interactions between year dummies and 1991 per-capita transport and communications infrastructure, market access, and electricity. Standard errors, in parentheses, are clustered at the subdistrict level.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Although the level effects are small in absolute terms, they are meaningful relative to the initial stock of higher-education institutions. In the preferred specification with full controls (Column 3), a one standard-deviation increase in mutt exposure is associated with an increase in the HEI of about 0.010 institutions per 1,000 residents. Given a 1991 mean of 0.11 institutions per 1,000, this represents a roughly 9 percent larger stock of higher-education institutions in high-exposure subdistricts by the 2001 and 2011 Censuses. For an average subdistrict with 170,000 residents in 1991, this translates into approximately 1.7 additional secondary schools or colleges.

### 6.5.1 Timing of higher-education expansion

To examine when the higher-education gap opens up, we also estimate a dynamic version of equation (5) that replaces the single post-reform interaction with separate interactions

between exposure and year dummies for 2001 and 2011:

$$\text{HEI}_{ct} = \alpha_c + \lambda_t + \delta_{2001}(\text{Exposure}_c \times \mathbf{1}\{t = 2001\}) + \delta_{2011}(\text{Exposure}_c \times \mathbf{1}\{t = 2011\}) + \theta'_t Z_c + \varepsilon_{ct}, \quad (6)$$

where the omitted year is 1991 and  $Z_c$  denotes the 1991 baseline covariates interacted with year dummies, as above. Table 10 reports the estimated interaction coefficients.

Table 10: Timing of higher-education expansion

|                                                      | HEI (per 1,000 residents) |
|------------------------------------------------------|---------------------------|
| Exposure $\times$ $\mathbf{1}\{\text{Year} = 2001\}$ | 0.0071***<br>(0.0014)     |
| Exposure $\times$ $\mathbf{1}\{\text{Year} = 2011\}$ | 0.0151***<br>(0.0040)     |
| Subdistrict fixed effects                            | Yes                       |
| Year fixed effects                                   | Yes                       |
| 1991 SC/ST/literacy $\times$ year                    | Yes                       |
| Ruggedness $\times$ year                             | Yes                       |
| Observations                                         | 528                       |
| Subdistricts                                         | 176                       |

Notes: The dependent variable is the higher-education institutions index (HEI) per 1,000 residents, defined as in Table 9. The table reports coefficients on interactions between standardised pre-1990 mutt exposure and year dummies for 2001 and 2011; 1991 is the omitted year. All specifications include subdistrict fixed effects, year fixed effects, and interactions between year dummies and the 1991 shares of Scheduled Castes and Scheduled Tribes, the literacy rate, and a standardised ruggedness index. Standard errors, in parentheses, are clustered at the subdistrict level.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

By 2001, subdistricts with one standard-deviation higher mutt exposure have about 0.007 more higher-education institutions per 1,000 residents than comparable low-exposure subdistricts, a gap of roughly 6–7 percent relative to the 1991 mean. By 2011 this gap more than doubles, to about 0.015 institutions per 1,000 residents (14 percent of the 1991 mean). The fact that the exposure gradient in HEI is zero in 1991 by construction, modest but precisely estimated by 2001, and substantially larger by 2011 is consistent with a gradual build-up of educational infrastructure in high-exposure areas. Combined with the Economic Census dynamics in Table 4, these results support a mechanism in which mutts facilitate the expansion of upper-secondary and tertiary education, which in turn may enable subdistricts to capitalise more fully on the opportunities created by post-1990

economic reforms.

## 6.6 Mutts and political competition

The third channel examined is political competition. Historical and contemporary accounts emphasise that mutts in Karnataka act as conveners in local politics, endorsing candidates, mediating disputes, and lobbying for public investment. If mutts make it harder for incumbents or dominant parties to win elections by large margins, they may strengthen incentives to supply public goods and support non-farm enterprise.

A panel of Karnataka legislative assembly elections from 1978 to 2018 is assembled. For each assembly constituency  $\ell$  and election year  $t$ , the data record the vote shares of the winner and runner-up and the turnout rate. Two measures of electoral competition are constructed. The first is the margin of victory, defined as the absolute difference between the winner's and runner-up's vote shares in percentage points. The second is competitiveness, defined as 100 minus the winner's vote share, so that higher values correspond to tighter races.

To measure the political reach of mutts, the same pre-1990 measure of mutt exposure introduced in Section 6.1 is used. The construction of this measure at the assembly constituency level has already been described.

The 1994 state election is treated as the beginning of the post-reform period, reflecting both the national liberalisation of the early 1990s and the implementation of the 73rd and 74th Constitutional Amendments. Let  $\text{Post94}_t$  be an indicator equal to one for elections from 1994 onwards and zero for earlier elections. For each electoral outcome  $E_{\ell t}$  the following specification is estimated:

$$E_{\ell t} = \alpha_{\ell} + \lambda_t + \gamma (\text{Exposure}_{\ell} \times \text{Post94}_t) + \psi'_t Z_{\ell} + \eta_{\ell t}, \quad (7)$$

where  $\alpha_{\ell}$  and  $\lambda_t$  are constituency and year fixed effects, and  $Z_{\ell}$  collects 1991 baseline characteristics which are the shares of the population who are Scheduled Castes and Scheduled Tribes, the literacy rate, and a standardised measure of ruggedness, each interacted with year dummies so that these characteristics can have their own time trends. Standard errors are clustered at the constituency level.

Table 11: Mutt exposure and assembly election competition

|                                  | (1)<br>Margin of victory | (2)<br>Competitiveness |
|----------------------------------|--------------------------|------------------------|
| Post-1994 $\times$ exposure      | −0.66***<br>(0.25)       | 0.70***<br>(0.21)      |
| Constituency fixed effects       | Yes                      | Yes                    |
| Year fixed effects               | Yes                      | Yes                    |
| Additional controls <sup>a</sup> | Yes                      | Yes                    |
| Observations                     | 1,765                    | 1,765                  |
| Constituencies                   | 353                      | 353                    |
| $R^2$                            | 0.26                     | 0.45                   |

Notes: The unit of observation is an assembly constituency in a Karnataka state election between 1978 and 2018. Column (1) uses the margin of victory (absolute difference between winner and runner-up vote shares, in percentage points) as the dependent variable. Column (2) uses competitiveness, defined as 100 minus the winner's vote share (in percentage points). "Exposure" is the constituency-level mutt exposure, constructed by aggregating subdistrict-level mutt exposure (the 20 kilometre log-kernel of mutt locations per 1,000 residents based on the 1991 population, standardised to mean zero and unit variance) using 1991 population weights and then re-standardising across constituencies. "Post-1994" is an indicator for elections from 1994 onwards; earlier elections form the baseline. The reported coefficients are on the interaction between Post-1994 and exposure, and can be interpreted as the change (in percentage points) in the outcome associated with a one-standard-deviation increase in exposure in the post-1994 period, relative to the pre-1994 period. All specifications include constituency fixed effects and year fixed effects. <sup>a</sup>Additional controls are interactions of election-year dummies with the 1991 shares of Scheduled Castes and Scheduled Tribes, the literacy rate, and a standardised ruggedness index. Standard errors, in parentheses, are clustered at the constituency level.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Column (1) of Table 11 shows that, in elections held after 1994, a one standard deviation increase in mutt exposure is associated with a reduction in the margin of victory of about 0.66 percentage points. Column (2) shows that the same increase in exposure is associated with a rise in competitiveness of about 0.70 percentage points. Although modest in magnitude, these effects are precisely estimated and consistent across the two indicators. They suggest that, in the post-reform period, constituencies with stronger mutt presence are characterised by tighter electoral races and somewhat weaker dominance by incumbent or leading parties. As a complementary check, Appendix A.2 reports results for two additional electoral outcomes: the effective number of parties (ENOP) and voter turnout. Using the same difference-in-differences design with constituency and year fixed effects, the coefficients on the interaction between mutt exposure and the post-1994 indicator are small and statistically indistinguishable from zero for both ENOP and turnout. In other

words, mutt-exposed constituencies do not systematically display a more fragmented party system or higher (or lower) voter participation after 1994. The political effects of mutts therefore appear to operate primarily through tightening margins in contests, rather than through broad changes in party entry or mass mobilisation.

In combination with the higher-education and Economic Census findings, these results support a common narrative: subdistricts with greater pre-1991 mutt exposure both build up more upper-secondary and tertiary institutions and experience more competitive legislative elections, and they also realise larger subsequent gains in non-farm employment. This is consistent with a mechanism in which mutts, as locally embedded religious institutions, foster human capital and intensify political competition, thereby making their regions better positioned to benefit from the opportunities created by economic liberalisation.

Read alongside the education and employment results, the election findings support a broader interpretation of mutts as local institutions that shape both the supply side and the political environment of development. Subdistricts with higher mutt exposure not only expand their stock of upper-secondary and tertiary institutions and generate more non-farm employment, they also exhibit more competitive legislative elections. These results are consistent with a mechanism in which mutts facilitate the emergence of educated elites and local organisational capacity, and in which more competitive politics and stronger local institutions in turn foster investment, business formation, and the diffusion of new economic opportunities in the post-liberalisation period.

## 7 Conclusion

This paper has examined how a long standing religious infrastructure shapes economic and political development in a large Indian state. We focus on Hindu monastic institutions in Karnataka, known as *mutts* or *mathas*, which combine religious authority with control over schools, hostels, land, and local political networks. Using a new geocoded inventory of 253 historically important mutts, linked to subdistrict and constituency level administrative data on employment, education, and elections, we study how pre-1990 mutt exposure is associated with subsequent economic change and political competition.

Three main findings emerge. First, mutt exposed subdistricts experience systematically larger gains in non-farm employment per capita in the post-reform Economic Census rounds. A one standard deviation increase in pre-1990 mutt exposure is associated with roughly 10 additional non-farm jobs per 1,000 residents, relative to a 1990 baseline of about 62 jobs per 1,000, conditional on baseline controls. These gains are concentrated in manufacturing

and private employment rather than services or public jobs. Log specifications imply somewhat smaller semi-elasticities, around 7 percent for total non-farm employment and about 11 to 12 percent for manufacturing, which can be viewed as conservative versions of the same result.

Second, subdistricts with greater mutt exposure build more upper secondary schools and colleges per capita. A one standard deviation higher exposure is associated with an increase of about 0.01 institutions per 1,000 residents in the post-reform period. Given a 1991 mean of about 0.11 institutions per 1,000, this corresponds to roughly a 9 percent higher stock of upper secondary and tertiary institutions in high exposure areas by 2001 and 2011.

Third, legislative assembly constituencies with higher mutt exposure experience somewhat tighter electoral competition after 1994. In a panel of Karnataka state elections from 1978 to 2018, a one standard deviation increase in mutt exposure is associated with a post-1994 reduction in the margin of victory of about 0.7 percentage points and a corresponding increase in a simple competitiveness index. There is no clear association with changes in turnout or in the effective number of parties, which suggests that the main effect operates through closer races rather than through large shifts in participation or party fragmentation.

Taken together, the evidence suggests that mutts in Karnataka act as local development institutions. Pre-1990 mutt exposure predicts where non-farm employment expands after liberalisation, where upper secondary and tertiary education capacity grows more quickly, and where legislative races become more competitive after electoral reforms. The timing and pattern of the results are consistent with a mechanism in which mutts use their organisational capacity, landholdings, and political connections to invest in human capital and social welfare, and to channel public resources and private opportunities toward their constituencies. By expanding schooling and hostels and by mediating between followers and the state, mutts appear to have helped their communities take advantage of new non-farm jobs while also raising the electoral stakes around those communities.

These findings speak to several strands of work. In the economics of religion, they add a new case in which religious organisations appear to support rather than block structural transformation, alongside evidence on Protestant and Catholic missions, Islamic charities, and shrines (e.g. Becker and Woessmann, 2009; Woodberry, 2012; Calvi, 2022; Valencia Caicedo, 2019; Chaudhary and Rubin, 2011; Malik and Mirza, 2022; Bazzi et al., 2020). In the literature on networks and development, they show how identity based organisations embedded in caste and regional networks can lower barriers to non-farm entry and expand access to education, at least for the groups they serve (Munshi and Rosenzweig,

2016; Munshi, 2019). In political economy, the fact that mutt exposed constituencies experience tighter elections links religious infrastructure to electoral competition and to work on how brokers, parties, and local intermediaries shape responsiveness and public goods provision (Besley and Burgess, 2002; Finan and Mazzocco, 2021; Asher and Novosad, 2017).

The analysis also has important limitations for any causal interpretation. First, the empirical strategy is reduced form and relies on the assumption that, conditional on rich pre-1990 controls and their interactions with time, variation in mutt exposure is as good as random with respect to post-reform changes in employment, education, and electoral outcomes. Historical evidence suggests that major mutts predate the period of analysis by decades or centuries, which supports treating pre-1990 exposure as predetermined. It is still not possible to rule out unobserved pre-1990 characteristics that both attracted powerful mutts and favoured later development, such as deeper traditions of learning, stronger local elites, or better access to markets. The extensive baseline controls mitigate this concern but do not eliminate it.

Second, the analysis is largely silent on distributional and normative questions within subdistricts. Mutts are closely aligned with specific caste communities, and the observed expansion of educational capacity and non-farm employment in high-exposure areas may disproportionately benefit followers and affiliates. Admissions to mutt run hostels and schools, informal job referrals, and access to welfare programmes may be skewed towards particular jatis. Without household or individual level microdata, it is not possible to assess whether mutt linked development narrows or widens gaps between groups. Evidence from work on deras and shrines suggests that religious organisations can both mitigate and reproduce inequality (Batra and Prabhakar, 2024; Malik and Mirza, 2022). The findings here should therefore be read as average effects at the subdistrict or constituency level, not as statements about equity.

Third, external validity is limited. The study focuses on one Indian state and on a set of Hindu monastic institutions that are relatively large, visible, and closely linked to state welfare schemes. The results need not carry over to smaller mutts, to mutts in other regions, or to Sikh deras, churches, mosques, and other religious organisations that operate in different regulatory and political environments. They may also differ in settings with weaker state capacity or different forms of religious authority. We therefore treat the paper as a detailed case study of one important context, rather than as evidence of a universal relationship between religious institutions and development outcomes.

Even with these caveats, the results have implications for both policy and research. For policy, the Karnataka case illustrates the potential of religious organisations with deep

local roots and extensive educational and welfare infrastructures to act as partners in delivering public goods and supporting structural change. Mutts appear to have expanded the supply of upper secondary and tertiary education and to have supported growth in manufacturing and private non-farm employment.

For research, several extensions are natural. One direction is to exploit heterogeneity across mutts, since large and politically active mutts may operate very differently from smaller and more inward looking establishments. Another is to study distributional effects within subdistricts using household surveys or administrative microdata on enrolment, exam performance, and labour market outcomes by caste, gender, and religion. A third is to unpack the channels linking mutts to non-farm employment, for example, by combining firm and worker level data on location, credit, migration, and occupational choice with qualitative work on how mutt networks connect entrepreneurs, workers, and politicians. A fourth is to extend the framework to other states and religious traditions, for instance by comparing mutt dense regions of Karnataka with Sikh areas of Punjab, regions with strong deras, or districts with dense Christian or Islamic institutions.

More broadly, the analysis shows how combining administrative microdata with geocoded information on religious institutions can illuminate the interaction between historical religious infrastructure and contemporary development. As similar data become available in other settings, it should be possible to build a comparative picture of when and how religious organisations act as providers of club goods, brokers of political voice, or gatekeepers of opportunity. The Karnataka mutt case suggests that, in a large Hindu majority democracy, religious institutions can serve as catalysts of human capital formation and structural transformation and can modestly intensify electoral competition, even as they pose familiar challenges of inclusion and accountability.

## References

- Montek S. Ahluwalia. Economic reforms in india since 1991: Has gradualism worked? *Journal of Economic Perspectives*, 16(3):67–88, 2002. doi: 10.1257/089533002760278721.
- Sam Asher and Paul Novosad. Politics and local economic growth: Evidence from india. *American Economic Journal: Applied Economics*, 9(1):229–273, 2017. doi: 10.1257/app.20150512.
- Sam Asher and Paul Novosad. Rural roads and local economic development. *American Economic Review*, 110(3):797–823, 2020. doi: 10.1257/aer.20180268.
- 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. *World Bank Economic Review*, 35(4):845–871, 2021. doi: 10.1093/wber/lhab003.
- Naheed Ataulla. Karnataka: The importance of being lingayat. National Herald, 2022. URL <https://www.nationalheraldindia.com/india/karnataka-the-importance-of-being-lingayat>. Accessed 25 November 2025.
- Emmanuelle Auriol, Julie Lassébie, Amma Panin, Eva Raiber, and Paul Seabright. God Insures those Who Pay? Formal Insurance and Religious Offerings in Ghana. *The Quarterly Journal of Economics*, 135(4):1799–1848, 07 2020. ISSN 0033-5533. doi: 10.1093/qje/qjaa015. URL <https://doi.org/10.1093/qje/qjaa015>.
- Abdul Aziz. Democratic decentralisation: Experience of Karnataka. *Economic and Political Weekly*, 35(39):3521–3526, 2000.
- Corry Azzi and Ronald Ehrenberg. Household allocation of time and church attendance. *Journal of Political Economy*, 83(1):27–56, 1975. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/1833272>.
- Harpreet Bajwa. Deras: Where religion and power coincide. The New Indian Express, 2017. URL <https://www.newindianexpress.com/nation/2017/Apr/28/deras-where-religion-and-power-coincide-1598834.html>. Accessed 25 November 2025.
- Robert J. Barro and Rachel M. McCleary. Religion and economic growth across countries. *American Sociological Review*, 68(5):760–781, 2003.

- Kartikeya Batra and Avantika Prabhakar. Divine intervention? religious organizations and public service delivery in india. Working paper, 2024.
- Samuel Bazzi, Gabriel Koehler-Derrick, and Benjamin Marx. The institutional foundations of religious politics: Evidence from indonesia. *Quarterly Journal of Economics*, 135(2): 845–911, 2020. doi: 10.1093/qje/qjz040.
- Sascha O. Becker and Ludger Woessmann. Was weber wrong? a human capital theory of protestant economic history. *Quarterly Journal of Economics*, 124(2):531–596, 2009. doi: 10.1162/qjec.2009.124.2.531.
- Sascha O. Becker, Amma Serwaah-Panin, Steven Pfaff, and Jared Rubin. Religion and economic development: Past, present, and future. Technical report, 2025. IZA Discussion Paper No. 17747.
- Jeanet Sinding Bentzen. Acts of God? Religiosity and Natural Disasters Across Subnational World Districts. *The Economic Journal*, 129(622):2295–2321, 05 2019. ISSN 0013-0133. doi: 10.1093/ej/uez008. URL <https://doi.org/10.1093/ej/uez008>.
- Eli Berman. Sect, subsidy, and sacrifice: An economist’s view of ultra-orthodox jews. *The Quarterly Journal of Economics*, 115(3):905–953, 2000. ISSN 00335533, 15314650. URL <http://www.jstor.org/stable/2586899>.
- Eli Berman. *Radical, Religious, and Violent: The New Economics of Terrorism*. MIT Press, Cambridge, MA, 2009.
- Eli Berman and David D. Laitin. Religion, terrorism and public goods: Testing the club model. *Journal of Public Economics*, 92(10–11):1942–1967, 2008.
- Timothy Besley and Robin Burgess. The political economy of government responsiveness: Theory and evidence from india. *Quarterly Journal of Economics*, 117(4):1415–1451, 2002. doi: 10.1162/003355302320935061.
- Timothy Besley, Rohini Pande, Lupin Rahman, and Vijayendra Rao. The politics of public good provision: Evidence from indian local governments. *Journal of the European Economic Association*, 2(2-3):416–426, 05 2004. doi: 10.1162/154247604323068104. URL <https://doi.org/10.1162/154247604323068104>.
- Vani K. Borooah and Sriya Iyer. Vidya, veda, and varna: The influence of religion and caste on education in rural india. *Journal of Development Studies*, 41(8):1369–1404, 2005. doi: 10.1080/00220380500186960.

- Rossella Calvi. The protestant legacy: Missions, gender, and human capital in india. *Journal of Human Resources*, 57(6):1946–1980, 2022. doi: 10.3368/jhr.58.2.0919-10437 R2.
- Jean-Paul Carvalho. Veiling. *Quarterly Journal of Economics*, 128(1):337–370, 2013. doi: 10.1093/qje/qjs045.
- Jean-Paul Carvalho. Identity-based organizations. *American Economic Review*, 106(5): 410–414, 2016. doi: 10.1257/aer.p20161039.
- Jean-Paul Carvalho and Mark Koyama. Jewish emancipation and schism: Economic development and religious change. *Journal of Comparative Economics*, 44(3):562–584, 2016. doi: 10.1016/j.jce.2016.06.002.
- Jean-Paul Carvalho, Sriya Iyer, and Jared Rubin. Introduction. In *Advances in the Economics of Religion*, pages 1–21. Palgrave Macmillan, 2019.
- Mark C. Casson, Marina Della Giusta, and Uma S. Kambhampati. Formal and informal institutions and development. *World Development*, 38(2):137–141, 2010. doi: 10.1016/j.worlddev.2009.10.008.
- Latika Chaudhary and Jared Rubin. Reading, writing, and religion: Institutions and human capital formation. *Journal of Comparative Economics*, 39(1):17–33, 2011. doi: 10.1016/j.jce.2010.06.001.
- Daniel L. Chen. Club goods and group identity: Evidence from islamic resurgence during the indonesian financial crisis. *Journal of Political Economy*, 118(2):300–354, 2010. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/10.1086/652462>.
- ET Online. Why mutts matter so much in karnataka’s electoral politics. The Economic Times, 2023. URL <https://economictimes.indiatimes.com/news/elections/assembly-elections/karnataka/why-mutts-matter-so-much-in-karnatakas-electoral-politics/articleshow/99940128.cms>. Accessed 25 November 2025.
- Frederico Finan and Maurizio Mazzocco. Electoral incentives and the allocation of public funds. *Journal of the European Economic Association*, 19(5):2467–2512, 2021. doi: 10.1093/jeea/jvaa055.
- Andrew D. Foster and Mark R. Rosenzweig. Agricultural productivity growth, rural economic diversity, and economic reforms: India, 1970–2000. *Economic Development*

- and Cultural Change*, 52(3):509–542, 2004. doi: 10.1086/420968. URL <https://doi.org/10.1086/420968>.
- Ejaz Ghani, Arti Grover Goswami, and William R. Kerr. Highway to success: The impact of the golden quadrilateral project for the location and performance of indian manufacturing. *The Economic Journal*, 126(591):317–357, 2016. doi: <https://doi.org/10.1111/ecoj.12207>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12207>.
- K. Giriprakash. Why are karnataka’s religious mutts so powerful? *The Hindu Business Line*, April 2018. URL <https://www.thehindubusinessline.com/news/national/why-are-karnatakas-religious-mutts-so-powerful/article23576630.ece>. Online article; accessed 28 November 2024.
- Government of India. The constitution (seventy-third amendment) act, 1992. Ministry of Law and Justice, Government of India, 1992a.
- Government of India. The constitution (seventy-fourth amendment) act, 1992. Ministry of Law and Justice, Government of India, 1992b.
- Government of Karnataka. The Karnataka Panchayat Raj act, 1993. Act No. 14 of 1993, Government of Karnataka, 1993. As amended.
- Avner Greif. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, Cambridge, 2006.
- Jonathan Gruber and Daniel M. Hungerman. Faith-based charity and crowd-out during the great depression. *Journal of Public Economics*, 91(5-6):1043–1069, 2007.
- Luigi Guiso, Paola Sapienza, and Luigi Zingales. People’s opium? religion and economic attitudes. *Journal of Monetary Economics*, 50(1):225–282, 2003.
- Luigi Guiso, Paola Sapienza, and Luigi Zingales. Does culture affect economic outcomes? *Journal of Economic Perspectives*, 20(2):23–48, 2006. doi: 10.1257/jep.20.2.23.
- Thomas Blom Hansen and Finn Stepputat. *States of Imagination: Ethnographic Explorations of the Postcolonial State*. Duke University Press, 2001.
- Joseph Henrich. The evolution of costly displays, cooperation and religion: Credibility enhancing displays and their implications for cultural evolution. *Evolution and Human Behavior*, 30(4):244–260, 2009.

- Chang-Tai Hsieh and Peter J. Klenow. Misallocation and manufacturing tfp in china and india\*. *The Quarterly Journal of Economics*, 124(4):1403–1448, 11 2009. ISSN 0033-5533. doi: 10.1162/qjec.2009.124.4.1403. URL <https://doi.org/10.1162/qjec.2009.124.4.1403>.
- Daniel Hungerman. Religious Institutions and Economic Wellbeing. *The Future of Children*, 30(2020):9–28, 2020. ISSN 1550-1558. doi: 10.1353/foc.2020.0003. URL <https://muse.jhu.edu/article/764023>.
- Daniel M. Hungerman. Are church and state substitutes? evidence from the 1996 welfare reform. *Journal of Public Economics*, 89(11-12):2245–2267, 2005.
- Laurence R. Iannaccone. Religious practice: A human capital approach. *Journal for the Scientific Study of Religion*, 29(3):297–314, 1990.
- Laurence R. Iannaccone. Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy*, 100(2):271–297, 1992.
- Laurence R. Iannaccone. Introduction to the economics of religion. *Journal of Economic Literature*, 36(3):1465–1495, 1998. ISSN 00220515. URL <http://www.jstor.org/stable/2564806>.
- Aya Ikegame. Why do backward castes need their own gurus? the social and political significance of new caste-based monasteries in karnataka. *Contemporary South Asia*, 18(1):57–70, 2010. doi: 10.1080/09584930903561630.
- Aya Ikegame. Guru governance: Political anthropology on religious leaders in south india. *Annual Review of Social Anthropology*, 37:87–111, 2011.
- Aya Ikegame. The governing guru: Hindu mathas in liberalising india. In Jacob Copeman and Aya Ikegame, editors, *The Guru in South Asia: New Interdisciplinary Perspectives*, pages 46–63. Routledge, London and New York, 2012a.
- Aya Ikegame. Mathas, gurus and citizenship: The state and communities in colonial india. *Citizenship Studies*, 16(5-6):689–703, 2012b. doi: 10.1080/13621025.2012.698494.
- Aya Ikegame. *Princely India Re-imagined: A Historical Anthropology of Mysore from 1799 to the Present*. Routledge, London, 2013.
- Sriya Iyer. The new economics of religion. *Journal of Economic Literature*, 54(2):395–441, June 2016. doi: 10.1257/jel.54.2.395. URL <https://www.aeaweb.org/articles?id=10.1257/jel.54.2.395>.

- Sriya Iyer. *The Economics of Religion in India*. Harvard University Press, Cambridge, MA, 2018. doi: 10.4159/9780674989313.
- Sriya Iyer, Chander Velu, and Melvyn Weeks. Divine competition: Religious organisations and service provision in india. Cambridge Working Papers in Economics 1409, Faculty of Economics, University of Cambridge, 2014. URL <https://www.econ.cam.ac.uk/publications/cwpe/1409>.
- Francesca R. Jensenius and Gilles Verniers. Studying indian politics with large-scale data: Indian election data 1961–today. *Studies in Indian Politics*, 5(2):269–275, 2017. doi: 10.1177/2321023017727984.
- Craig Johnson. Decentralisation in India: Poverty, politics and Panchayati Raj. Working Paper 199, Overseas Development Institute, 2003.
- JSS Mahavidyapeetha. Jss mahavidyapeetha: In pursuit of a progressive society. Jagadguru Sri Veerasimhasana Mahasamsthana Math, Suttur, 2022. URL <https://sutturmath.org/jss-mahavidyapeetha/>. Accessed 25 November 2025.
- Adeel Malik and Rinchan Ali Mirza. Pre-colonial religious institutions and development: Evidence through a military coup. *Journal of the European Economic Association*, 20(2):907–956, 2022. doi: 10.1093/jeea/jvab050.
- A. Mathew. Inclusion and privatization strands in karnataka higher education policy. *The Indian Journal of Technical Education*, 39(3):71–92, 2016.
- Michael McBride. Why churches need free-riders: Religious capital formation and religious group survival. *Journal of Behavioral and Experimental Economics*, 58:77–87, 2015. ISSN 2214-8043. doi: <https://doi.org/10.1016/j.socec.2015.07.001>. URL <https://www.sciencedirect.com/science/article/pii/S2214804315000828>.
- Michael McBride. A rational choice theory of religious authority. *Rationality and Society*, 28(4):410–438, 2016.
- Michael McBride. Spatial models of religious market competition: A critical assessment. In Jean-Paul Carvalho, Sriya Iyer, and Jared Rubin, editors, *Advances in the Economics of Religion*, pages 43–59. Palgrave Macmillan, 2019.
- Rachel M. McCleary and Robert J. Barro. Religion and economy. *Journal of Economic Perspectives*, 20(2):49–72, 2006.

- Vandana Menon. Karnataka mutts offer a model for temple freedom. bjp raises the slogan but priests do not. ThePrint, 2024. URL <https://theprint.in/ground-reports/karnataka-mutts-offer-a-model-for-temple-freedom-bjp-raises-the-slogan-but-priests-dont/1993034/>. Accessed 25 November 2025.
- Kaivan Munshi. Caste and the indian economy. *Journal of Economic Literature*, 57(4): 781–834, 2019. doi: 10.1257/jel.20171307.
- Kaivan Munshi and Mark Rosenzweig. Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy. *American Economic Review*, 96(4):1225–1252, 2006. doi: 10.1257/aer.96.4.1225.
- Kaivan Munshi and Mark Rosenzweig. Networks and misallocation: Insurance, migration, and the rural–urban wage gap. *American Economic Review*, 106(1):46–98, 2016. doi: 10.1257/aer.20131365.
- Douglass C. North. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge, 1990.
- Mancur Olson. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, Cambridge, MA, 1965.
- Arvind Panagariya. India in the 1980s and 1990s: A triumph of reforms. IMF Working Paper WP/04/43, International Monetary Fund, 2004.
- Jean-Philippe Platteau. Religion, politics, and development: Lessons from the lands of islam. *Journal of Economic Behavior & Organization*, 68(2):329–351, 2008. ISSN 0167-2681. doi: <https://doi.org/10.1016/j.jebo.2008.06.003>. URL <https://www.sciencedirect.com/science/article/pii/S0167268108001248>.
- Sharan Poovanna. Education, justice, poll tickets — what are karnataka’s mutts and how they became power centres. ThePrint, 2023. URL <https://theprint.in/india/education-justice-poll-tickets-what-are-karnatakas-mutts-how-they-became-power-centres/1313148/>. Accessed 25 November 2025.
- Mohamed Saleh. The reluctant transformation: State industrialization, religion, and human capital in nineteenth-century egypt. *Journal of Economic History*, 75(1):65–94, 2015.
- Mohamed Saleh. Public mass modern education, religion, and human capital in twentieth-century egypt. *The Journal of Economic History*, 76(3):697–735, 2016. URL <https://doi.org/10.1017/S0022216X16000031>.

[//www.cambridge.org/core/journals/journal-of-economic-history/article/public-mass-modern-education-religion-and-human-capital-in-twentiethcentury-egypt/9B444DB3385B59B2C6169A7BD1BBCC3B](http://www.cambridge.org/core/journals/journal-of-economic-history/article/public-mass-modern-education-religion-and-human-capital-in-twentiethcentury-egypt/9B444DB3385B59B2C6169A7BD1BBCC3B).

Shri Dharmasthala Manjunatheshwara Educational Society. About sdme society. SDM Educational Society, Ujire, 2018. URL <https://www.sdmesociety.in/about-us-2/>. Accessed 25 November 2025.

A. D. Sood. *The Matha State: Kinship, Asceticism and Institutionalality in the Public Life of Karnataka*. University of Chicago, Department of South Asian Languages and Civilizations and Department of Anthropology, 2006.

Sree Siddaganga Education Society. Educational services. Sree Siddaganga Math, 2025. URL <https://siddagangamath.org/siddaganga/institute.html>. Accessed 25 November 2025.

Rodney Stark. Micro foundations of religion: A revised theory. *Sociological Theory*, 17(3): 264–289, 1999.

Guillermo Trejo. Religious competition and ethnic mobilization in latin america: Why the catholic church promotes indigenous movements in mexico. *American Political Science Review*, 103(3):323–342, 2009.

Felipe Valencia Caicedo. The mission: Human capital transmission, economic persistence, and culture in south america. *Quarterly Journal of Economics*, 134(1):507–556, 2019. doi: 10.1093/qje/qjy024.

Robert D. Woodberry. The missionary roots of liberal democracy. *American Political Science Review*, 106(2):244–274, 2012. doi: 10.1017/S0003055412000093.

## A Additional empirical diagnostics

### A.1 Baseline balance between mutt exposure and covariates

This subsection reports cross-sectional balance regressions that relate mutt exposure to pre-reform characteristics. If exposure were simply a proxy for richer or more advantaged subdistricts, the Economic Census results could reflect pre-existing growth paths rather than the role of mutts as institutions.

For each baseline variable  $Z_c$  measured in 1991 we estimate

$$Z_c = \gamma_0 + \gamma_1 \text{Exposure}_c + \varepsilon_c,$$

where  $\text{Exposure}_c$  is the standardised mutt exposure index per 1,000 residents, constructed as a 20 kilometre log kernel of mutts around the subdistrict centroid, weighted by mutt salience and divided by the 1991 population.

We consider seven baseline characteristics: the 1991 shares of the population who are Scheduled Castes and Scheduled Tribes, the literacy rate, a standardised measure of terrain ruggedness, and standardised per-capita measures of markets, transport facilities, and communications (postal and telephone). Table A.1 reports the estimated coefficients.

Table A.1: Balance of mutt exposure with 1991 baseline characteristics

|                            | (1)                 | (2)              | (3)                 | (4)                 | (5)                  | (6)                 | (7)                 |
|----------------------------|---------------------|------------------|---------------------|---------------------|----------------------|---------------------|---------------------|
|                            | SC share            | ST share         | Literacy            | Ruggedness (z)      | Markets (z)          | Transport (z)       | Communications (z)  |
| Exposure index             | −0.006**<br>(0.003) | 0.001<br>(0.002) | 0.020***<br>(0.004) | 0.211***<br>(0.038) | −0.153***<br>(0.038) | 0.207***<br>(0.038) | 0.214***<br>(0.038) |
| Observations               | 668                 | 668              | 668                 | 668                 | 668                  | 668                 | 668                 |
| $R^2$                      | 0.008               | 0.000            | 0.038               | 0.045               | 0.024                | 0.044               | 0.047               |
| Mean of dependent variable | 0.179               | 0.052            | 0.399               | 0.000               | 0.000                | 0.000               | 0.000               |

Notes: Each column reports a separate cross-sectional regression at the subdistrict level in 1991. The dependent variables are the share of the 1991 population who are Scheduled Castes (Column 1), the share who are Scheduled Tribes (Column 2), the literacy rate (Column 3), a standardised index of terrain ruggedness (Column 4), and standardised per-capita measures of markets (Column 5), transport facilities (Column 6), and communications (postal and telephone) (Column 7). The regressor is the standardised mutt exposure index per 1,000 residents, constructed as a 20 kilometre log kernel of mutts around the subdistrict centroid, weighted by mutt salience and divided by the 1991 population. “Mean of dependent variable” reports the sample mean of each baseline characteristic (zero for the standardised indices). Standard errors, in parentheses, are conventional ordinary least squares standard errors. Significance levels: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

The patterns in Table A.1 show that mutts are not randomly located. Subdistricts with higher mutt exposure have slightly lower Scheduled Caste shares, somewhat higher literacy, and more rugged terrain. Quantitatively, a one standard-deviation increase in exposure is associated with a reduction of about 0.6 percentage points in the Scheduled Caste share (roughly 3 per cent of its mean), an increase of about 2 percentage points in literacy (about 5 per cent of its mean), and an increase of roughly 0.21 standard deviations in ruggedness. Exposure is also positively correlated with transport and communications,

by about 0.21 standard deviations, and negatively correlated with markets, by about 0.15 standard deviations. The coefficient on the Scheduled Tribe share is very close to zero and statistically insignificant.

At the same time, all  $R^2$  values are below five per cent, which indicates that these observables explain only a small fraction of the cross-sectional variation in exposure. These correlations motivate the rich control set in the Economic Census regressions. In our preferred specification we allow the 1991 Scheduled Caste share, Scheduled Tribe share, literacy rate, ruggedness index, and per-capita infrastructure measures to have their own year-specific slopes. This means that any time-varying effect of these baseline characteristics on nonfarm employment is absorbed by the control set. The coefficient on exposure in the difference-in-differences regressions is therefore identified from comparisons of subdistricts that share the same 1991 caste composition, literacy, terrain, and infrastructure, but differ in mutt exposure.

## A.2 Party fragmentation and turnout

The main elections analysis focuses on the margin of victory and a simple competitiveness index (100 minus the winner’s vote share). These outcomes capture how closely contested assembly races are, but they do not speak directly to two other dimensions of electoral politics: the fragmentation of the party system and the level of voter participation. This section therefore reports results for two additional constituency-level outcomes: the effective number of parties (ENOP) and voter turnout.

The effective number of parties is a standard summary of vote fragmentation, constructed from candidate-level vote shares in each assembly election. Higher values indicate that votes are more evenly spread across several candidates or parties, whereas lower values correspond to more concentrated, two-cornered contests. Turnout is measured as the percentage of registered voters who cast a ballot in the constituency.

For both outcomes, the specification mirrors the main elections difference-in-differences design in equation (7). Let  $E_{\ell t}$  denote either ENOP or turnout in constituency  $\ell$  in election year  $t$ . The estimating equation is

$$E_{\ell t} = \alpha_{\ell} + \lambda_t + \gamma (\text{Exposure}_{\ell} \times \text{Post94}_t) + \psi'_t Z_{\ell} + \eta_{\ell t},$$

where  $\alpha_{\ell}$  and  $\lambda_t$  are constituency and year fixed effects,  $\text{Exposure}_{\ell}$  is the standardised constituency-level mutt exposure index described in Section 4.3, and  $\text{Post94}_t$  is an indicator equal to one for elections from 1994 onwards. The vector  $Z_{\ell}$  collects 1991 baseline constituency characteristics (the shares of the population who are Scheduled Castes and

Scheduled Tribes, the literacy rate, and a standardised ruggedness index), each interacted with year dummies so that these characteristics can have flexible time paths. Standard errors are clustered at the constituency level. As in the main specifications, the main effects of exposure and the post-1994 dummy are absorbed by the fixed effects, leaving the interaction term as the parameter of interest.

Table A.2 reports the estimated coefficients on  $\text{Post94}_t \times \text{Exposure}_\ell$  for ENOP and turnout. For ENOP, the point estimate is small (about 0.017) and statistically indistinguishable from zero, with a tight confidence interval around the null. For turnout, the estimate is likewise close to zero (about 0.04 percentage points) and far from statistical significance. In both cases, the within-constituency  $R^2$  attributable to the exposure interaction is effectively zero.

These results suggest that, while constituencies with higher mutt exposure experience somewhat tighter post-1994 races in terms of margins and competitiveness (as shown in the main text), there is no evidence that mutt exposure systematically alters the fragmentation of the party system or overall voter participation. Put differently, mutt-exposed constituencies do not exhibit systematically more (or fewer) effective parties or higher (or lower) turnout after 1994. This pattern is consistent with an interpretation in which mutts affect electoral bargaining, candidate selection, and the closeness of contests, rather than broadly reshaping party entry or mass turnout.

Table A.2: Mutt exposure, party fragmentation, and turnout (assembly elections)

|                                              | (1)<br>ENOP      | (2)<br>Turnout (%) |
|----------------------------------------------|------------------|--------------------|
| Post-1994 $\times$ exposure <sup>a</sup>     | 0.017<br>(0.019) | 0.037<br>(0.141)   |
| Constituency fixed effects                   | Yes              | Yes                |
| Year fixed effects                           | Yes              | Yes                |
| Baseline controls $\times$ year <sup>b</sup> | Yes              | Yes                |
| Observations                                 | 1,765            | 1,765              |
| Constituencies                               | 353              | 353                |
| $R^2$                                        | 0.52             | 0.88               |
| Clusters (constituencies)                    | 353              | 353                |

*Notes:* The unit of observation is an assembly constituency in a Karnataka state election between 1978 and 2018. Column (1) uses the effective number of parties (ENOP), computed from candidate-level vote shares, as the dependent variable. Column (2) uses turnout, measured as the percentage of registered voters casting a ballot. All regressions are estimated using `reghdfe` with constituency and year fixed effects and standard errors clustered at the constituency level.

<sup>a</sup> “Exposure” is the standardised constituency-level mutt exposure index, constructed by aggregating the subdistrict-level 20 kilometre log-kernel exposure per 1,000 residents (based on the 1991 population) using 1991 population weights and re-standardising across constituencies. “Post-1994” is an indicator equal to one for elections from 1994 onwards. The table reports coefficients on the interaction between Post-1994 and exposure; the main effects of exposure and the post dummy are absorbed by the fixed effects.

<sup>b</sup> Baseline controls are the 1991 shares of Scheduled Castes and Scheduled Tribes, the literacy rate, and a standardised ruggedness index, each interacted with election-year dummies.

*Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .