MISSIONARY INFRASTRUCTURE AND THE DEEPENING OF OTTOMAN STATE SCHOOLING

ORHAN TORUL*

July 8, 2026

Abstract

How did the Ottoman state expand schooling when confronted with rival foreign institutions? Using sanjak-level data for Anatolia (1876–1902/08), I estimate whether districts hosting American missionary main stations saw faster growth of male *rüşdiye* schools—the tier that trained provincial officials. Conditional on demographics, baseline schooling, and province indicators, station sanjaks added roughly one more school (control mean 1.9). The estimate is stable across reweighting, alternative rosters, and spatial falsification tests, but statistically insignificant under design-based inference ($p = 0.15$–0.31), reflecting 16 treated sanjaks in a 15-province sample. The pattern suggests defensive deepening in established hubs rather than frontier entry.

Keywords: Educational competition, missionary schooling, state capacity, Ottoman Empire, administrative constraint

JEL Codes: I25, N33, O15, H75, Z12

*Address: Boğaziçi University, Department of Economics, 34342 Bebek, Istanbul, Turkey. E-Mail: orhan.torul@bogazici.edu.tr. I thank K. Kıvanç Karaman for his valuable feedback. I gratefully acknowledge the financial support of the Young Scientist Award (BAGEP) of Türkiye's Science Academy. All remaining errors are mine.

1 Introduction

How do historical states respond to the spatial presence of rival non-state institutions? In the late nineteenth century, the Ottoman Empire expanded secular schooling as part of a broader program of administrative modernization often described as "defensive modernization" (Berkes, 1964; Zürcher, 2004). This process, aimed at safeguarding sovereignty through the adoption of Western organizational forms, has parallels in other historical settings where states faced external pressures (Tilly, 1990). During the same period, Western missionary organizations established a durable institutional presence across Anatolia, concentrated in "main stations" that served as regional administrative and educational hubs (Kocabaşoğlu, 1989; Stone, 2006). This paper studies whether Ottoman educational investments were spatially responsive to these concentrated sites of foreign institutional presence.

A large literature on historical state-building emphasizes the expansion of state institutions under military, fiscal, and social pressure (Besley and Persson, 2011; Evered, 2012). Yet this literature often treats territorial expansion as the default form of state response. Under conditions of administrative scarcity, however, states may not broaden provision uniformly across space. They may instead deepen provision where they already have officials, infrastructure, and monitoring capacity (Ortaylı, 2000). This distinction motivates the central question of the paper: whether Ottoman educational competition with missionaries operated through the extensive margin of first entry into previously unserved districts or through the intensive margin of additional provision in districts where the state was already present.

The paper addresses this question with a sanjak-level dataset covering Ottoman Anatolia between 1876 and 1902–08 (Amasyalı, 2020). The sanjak is the administrative unit at which school counts are recorded in Ottoman provincial yearbooks, and therefore the finest spatial resolution at which the dependent variable is measured. I measure the intensive margin as the change in the number of male *rüşdiye* schools over the period and the extensive margin as the transition from zero to positive provision among sanjaks that lacked a male *rüşdiye* in 1876. The main explanatory variable is an indicator for whether a sanjak hosted a missionary main station, defined from missionary records as a parent hub with permanent foreign personnel and substantial institutional infrastructure. Conditioning on population, local religious composition, baseline schooling, and vilayet fixed effects, the analysis estimates the localized association between missionary infrastructure and subsequent Ottoman school expansion.

The core finding is that sanjaks hosting major missionary stations added approximately 1.1 more male *rüşdiye* schools over the period than observably similar sanjaks, against a control-group mean change of 1.9. The point estimate is stable—across sequential controls, reweighting, count models, alternative station rosters, and the exclusion of any treated sanjak or vilayet—and it survives conditioning on the density of Armenian and Greek community schools, on Armenian religious infrastructure, and on proxies for administrative and commercial importance. It is, however, statistically imprecise: under vilayet clustering $p = 0.31$, under randomization inference $p = 0.15$, and under the wild cluster bootstrap $p = 0.26$, a constraint imposed by 16 treated sanjaks spread across 13 of the sample's 15 provinces. Spatial falsification tests in the spirit of Conley and Kelly (2025) indicate that these design-based p-values

are well calibrated rather than understated. Formal sensitivity bounds locate the design's genuine exposure: a confounder as strong as the Christian-share control would leave the estimate essentially intact, but one as predictive as log population—an unobserved "development potential" factor—would absorb it (Section 6.2). By contrast, all five in-sample frontier sanjaks that lacked state schooling in 1876 gained provision by 1902–08 whether or not missionaries were present—a descriptive pattern consistent with an empire-wide floor on provision, though five observations cannot establish more than that. The data support a positive association between missionary main stations and subsequent male *rüşdiye* expansion, but the estimate is too imprecise to distinguish confidently between no effect and a moderately large positive one.

This argument differs in emphasis from two related bodies of work. The first documents the long-run effects of missionary activity on human capital, institutional development, and economic growth (Woodberry, 2012; Bai and Kung, 2015; Cagé and Rueda, 2016; Valencia Caicedo, 2019). These studies typically examine outcomes decades or centuries after missionary contact and identify missions as a source of lasting institutional change. The present paper asks a different question: not whether missionaries shaped long-run development, but whether the contemporary state responded to their presence. The second body of work interprets late Ottoman schooling through broader modernization narratives (Somel, 2001; Fortna, 2002; Pamuk, 2006). What remains less clear is whether the state's educational response took the form of broad spatial diffusion or targeted reinforcement of administratively important locations.

The closest empirical work is Amasyalı (2022a) and Amasyalı (2022b), which introduced the historical GIS data this paper builds on. Amasyalı (2022a) takes the Ottoman and Armenian educational reaction to missionary incursion as its historical premise—documented through the qualitative record—and estimates the long-run legacy of that educational competition for present-day development, measured by night-time luminosity. Amasyalı (2022b) studies the diffusion of women's education and, in its mechanism analysis, relates 1914 levels of *ibtidaî* and female *rüşdiye* counts to missionary presence in kaza-level cross-sections. Neither paper estimates the state's schooling response on the margin the Hamidian government itself prioritized: the growth of male *rüşdiye* provision, the tier that trained provincial officials. This paper supplies that missing step. It measures the change in male *rüşdiye* counts between 1876 and 1902–08 at the sanjak level—the level at which these counts are actually recorded—conditions on the 1876 baseline, identifies within provinces, and subjects the association to a design-based inference and sensitivity battery that this literature's Ottoman branch has not previously applied. Where the two literatures meet, the results align: consistent with Amasyalı (2022b), station presence predicts no differential in female state schools, while primary-school counts of both sexes show large, imprecise positive associations.

The paper contributes to three related literatures. First, it contributes to the political economy of education by showing that institutional competition shaped not simply whether schools expanded, but where and how expansion occurred (e.g., Gallego, 2010). Second, it contributes to the literature on state capacity by showing that administrative constraints can produce hub-centric responses to rival insti-

tutional presence. Third, it contributes to the study of late Ottoman reform by offering a district-level empirical account of educational expansion that takes provincial administrative structure seriously (Kodaman, 1988). The paper does not claim that missionary station placement was exogenous or that the design identifies a causal effect in a strong quasi-experimental sense. Its contribution is instead to document a structured pattern of association that is consistent with selective fortification under administrative constraint.

The remainder of the paper proceeds as follows. Section 2 describes the historical and institutional background of Ottoman educational reform and missionary organization. Section 3 introduces the dataset and defines the schooling margins and missionary treatment variables. Section 4 outlines the empirical strategy and inference framework. Section 5 presents the main results on the intensive and extensive margins of school expansion. Section 6 discusses credibility and sensitivity checks. Section 7 concludes.

2 Historical Background

The Hamidian era (1876–1909) was a period of intensified institutional change in the Ottoman Empire. Facing external military pressure, fiscal weakness, and internal separatist pressures, the central government sought to consolidate sovereignty through a broader expansion of administrative infrastructure (Ortaylı, 2000; Zürcher, 2004). One important component of this effort was the *Rüşdiye* system, a tier of secular middle schooling intended to produce literate and disciplined personnel for the provincial administration while also fostering a shared imperial framework (Somel, 2001). The spread of these schools therefore had educational, political, and administrative significance.

Male *rüşdiye* schools are a particularly useful outcome for studying this process. They occupied a central place in late Ottoman educational planning, were closely linked to bureaucratic training and state presence, and are observed with sufficient consistency across the relevant years to support district-level comparison (Kodaman, 1988). Other schooling categories, especially female schooling, are historically important but less suitable as the paper's main outcome because they appear later, are sparser, and are more vulnerable to discretionary or locally idiosyncratic placement.

Importantly, all parent missionary stations in the Anatolian sample were established before the baseline year of 1876, most during the 1830s–1860s (Stone, 2006). The temporal ordering—station founding, then school expansion—is therefore consistent with the paper's framing, though it does not by itself establish causation.

Missionary organization across Anatolia followed a distinct institutional hierarchy. The American Board of Commissioners for Foreign Missions and related organizations established "main stations" as permanent and heavily resourced centers (Stone, 2006). These hubs often included schools, hospitals, printing presses, boarding facilities, and foreign missionary residences (Kocabaşoğlu, 1989). From these nodes, missionary organizations supervised wider networks of "outstations" staffed more lightly and often through local intermediaries (Kieser, 2010). This distinction matters because main stations represented concentrated, visible, and durable foreign institutional presence, whereas outstations reflected

more diffuse and lower-capacity missionary activity.

Missionary location choice was not random, and the direction of its selection is historically specific. The ABCFM effort in Anatolia was directed overwhelmingly at Armenian and other Eastern Christian communities; conversion of Muslims was legally foreclosed, so stations followed Armenian population (Stone, 2006; Kieser, 2010). This matters because Armenian and Greek presence in this geography is not a neutral covariate: minority communities carried disproportionate human capital and commercial networks whose developmental legacy is well documented (Arbatlı and Gökmen, 2023). A sanjak attractive to missionaries on these grounds might have drawn state schools for reasons unrelated to competition—the more so because *rüşdiye* construction depended partly on local fiscal means (Somel, 2001). Two features of the setting cut against a purely mechanical selection-on-prosperity account. First, station placement tracked Armenian settlement into the landlocked eastern interior, which was poorer and less commercially integrated than the coastal west; the most developed maritime provinces held almost none of the network (Amasyalı, 2022a). Second, ten of the sixteen station sanjaks were provincial capital seats, so administrative centrality and treatment overlap substantially and must be conditioned on rather than assumed away. Sections 6.4, 6.5, and 6.2 address both fronts directly: with capital-seat, commercial-center, and administrative-size controls; with the density of Armenian and Greek community schools and Armenian religious infrastructure as competing regressors; with a disaggregated Armenian-share control; and with formal sensitivity bounds for what remains unobserved. None of this converts the design into a quasi-experiment, and the paper does not claim otherwise.

The Ottoman state had stronger reason to respond to main stations than to outstations. Main stations signaled a long-term organizational commitment and often clustered multiple forms of influence—educational, medical, and religious—within administratively significant places (Kieser, 2010). They were therefore more likely to be perceived as challenges to the state's claims over schooling, loyalty, and public authority (Fortna, 2002). Outstations, by contrast, were less capital intensive and more weakly institutionalized. If the state reacted strategically to missionary competition, one would therefore expect the response to be more closely tied to main stations than to diffuse missionary presence in general.

That officials monitored and reacted to missionary schooling is not merely a modeling assumption; it is documented in the administrative record. Following the 1869 Regulation of Public Education, foreign schools were required to hold permits and submit curricula and teachers to public inspection, and the government's recognition that it could not close missionary schools produced a concerted effort to improve the quantity and quality of state schools (Deringil, 1998). Deringil quotes the governor of Beirut in 1905—an Arab province outside this paper's Anatolian sample, but indicative of the empire-wide administrative mindset—to the effect that the only way to ensure Muslim children were "saved from the harmful clutches" of the Jesuits and Protestants was to establish a network of modern schools capable of competing with them (Deringil, 1998, 131), and Fortna (2002) documents the same competitive framing in Maarif correspondence. The question this paper asks is whether that documented administrative attention left a measurable spatial imprint on where male *rüşdiye* provision deepened.

Two contemporaneous shocks complicate the comparison and deserve explicit statement. The 1877–

78 Russo-Ottoman war and the ensuing resettlement of Muslim refugees (*muhacir*) altered population and fiscal burdens differentially across Anatolia, and the massacres of 1894–96 devastated Armenian communities precisely in several eastern station provinces. Both events fall inside the 1876–1902/08 outcome window, and both plausibly affected schooling through channels unrelated to educational competition. The vilayet fixed effects absorb province-level components of these shocks, and the 1881/93 census predates the massacres, but sanjak-level incidence within provinces remains a caveat that the design cannot eliminate.

This logic interacted with the administrative geography of the empire. Vilayets and their subordinate sanjaks were central to budgetary planning, personnel allocation, and implementation. Sanjaks that already possessed officials, routines of supervision, and some educational infrastructure were better positioned to absorb additional schools at relatively low marginal cost. By contrast, extending schooling to sanjaks without an existing foothold required more than constructing a building. It involved teacher recruitment, supervision, financing, and the creation of new administrative routines in places where the state had limited educational presence (Evered, 2012).

For that reason, the distinction between hub deepening and frontier expansion is not merely descriptive. It captures two different margins of state-building under constraint. Expanding into frontier sanjaks required fixed-cost investments in administrative reach, whereas adding schools in established hubs leveraged existing structures of authority. The paper tests whether missionary competition was associated primarily with the latter pattern: selective fortification of sanjaks where the state already had traction, rather than broad movement into previously unserved educational frontiers.

3 Data

The analysis uses a sanjak-level dataset covering Ottoman Anatolia (Amasyalı, 2020). This dataset is constructed by aggregating kaza-level demographic and geographic data to the sanjak, which is the administrative level at which school counts are recorded in Ottoman provincial yearbooks (*salnames*). The underlying kaza-level information covers 333 districts grouped into 57 sanjaks across 16 vilayets. School data in the salnames report school totals at the sanjak (or occasionally vilayet) level rather than at the kaza level (Somel, 2001; Kodaman, 1988). For this reason, the sanjak is both the natural unit of analysis and the finest spatial resolution at which the dependent variable exhibits independent variation. Both cross-sections are mapped onto a single, time-invariant boundary layer: the source digitizes kaza polygons in their late-nineteenth-century (circa 1893–1899) configuration and holds this geography fixed, so the analysis compares school counts within a constant set of administrative units, at the cost of abstracting from boundary changes over the window. Istanbul Vilayet is excluded from the regression sample as a structural outlier: its population exceeds the sample median by a factor of 15, and its school expansion (0 to 112) is an order of magnitude larger than any other sanjak.

Two properties of the source data deserve statement up front. First, salname school counts are administrative self-reports, and reporting quality plausibly improved with administrative centrality; if sta-

tion sanjaks were better documented, part of any measured differential could reflect recording rather than construction. The change specification differences out time-invariant reporting propensities, and the baseline-placebo check in Section 6.9 shows that stations do not predict the 1876 *level* of recorded provision, but time-varying reporting differences cannot be ruled out. Second, the "1902–08" endpoint reflects the salname volumes from which the later cross-section was compiled, which vary by province within that range; the outcome is therefore a change over roughly a quarter century rather than between two exact dates. Because the endpoint year varies at the province level, the vilayet fixed effects included in every specification absorb endpoint-timing differences across provinces.

Table 1 summarizes the treatment distribution and baseline characteristics of the resulting sample of 56 sanjaks. Panel A shows that 16 sanjaks hosted a main missionary station, 31 had outstations only, and 9 had no missionary presence. Panel B indicates that sanjaks hosting main stations had higher Christian population shares at baseline (normalized difference 0.83), highlighting the importance of conditioning on demographics and addressing imbalance through inverse-probability weighting and through re-estimation on the region of common support (Section 6.6).

Table 1: Descriptive Summary of Sample and Treatment Distribution (Sanjak Level)

Panel A: Missionary Exposure	Sanjaks	Share (%)
No missionary presence	9	16.1
Outstations only	31	55.4
Main station present	16	28.6
Panel B: Baseline Characteristics	Control	Main Station
Observations	40	16
Baseline male *rüşdiye* (1876)	3.5	3.9
Log population (1881)	11.94	12.23
Christian share (1881)	0.14	0.24
Number of kazas	5.4	6.8
Zero baseline school share	0.07	0.12

Figure 1 displays the spatial distribution of school expansion and missionary infrastructure across Ottoman Anatolia. Panel (a) shows the change in male *rüşdiye* schools by sanjak, with the sixteen station-hosting sanjaks marked by triangles. Panel (b) displays the treatment hierarchy, distinguishing sanjaks with no missionary presence, outstations only, and main stations.

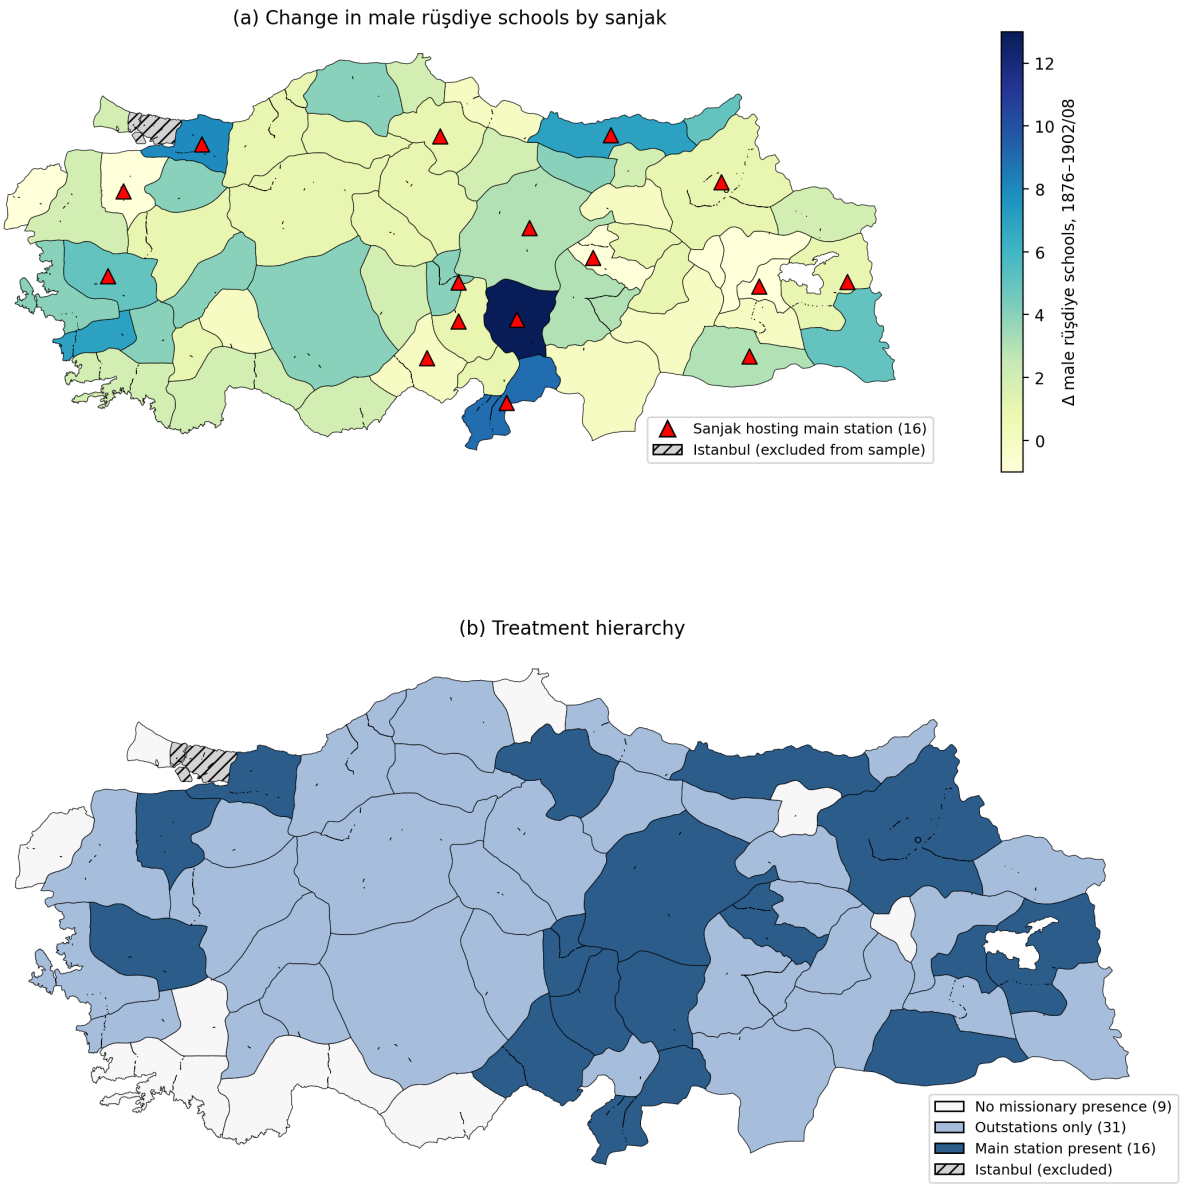


Figure 1: (a) Change in male *rüşdiye* schools (1876–1902/08) by sanjak. Triangles mark the sixteen sanjaks hosting main missionary stations, including the nine hubs matched manually to their kaza locations (Appendix Table A.2). Istanbul (hatched) is excluded from the regression sample, and the color scale spans in-sample values only. (b) Treatment hierarchy: no missionary presence, outstations only, and main station present.

3.1 Educational Outcomes

The main outcomes capture changes in male *Rüşdiye* provision between 1876 and 1902–08. The primary intensive-margin outcome is the change in the number of male *rüşdiye* schools per sanjak, recovered as integer counts from the salname data. This measure captures the degree to which the state deepened

educational provision within a sanjak over time. The change is not uniformly positive: six sanjaks record a net decline of one school over the period, a reminder that recorded provision could contract as well as expand.

Six sanjaks had zero male *rüşdiye* schools in 1876. One of them is Istanbul, which is excluded from the regression sample as an outlier; the five in-sample frontier sanjaks are Cebelibereket, Mardin, Hakkâri, Çatalca, and İzmit, of which two (Mardin and İzmit) hosted main stations. All five—and Istanbul—transitioned to positive provision by 1902–08. Because the outcome is constant, no regression is possible on this margin, and the paper accordingly treats the frontier pattern as descriptive: it is consistent with an empire-wide mandate to establish baseline provision everywhere, but five observations cannot discriminate among explanations for it. These five sanjaks remain in the intensive-margin regression sample; excluding them (the strict intensive margin, Appendix Table A.9) leaves the estimate essentially unchanged.

3.2 Missionary Infrastructure

The main explanatory variable is an indicator for whether a sanjak hosted a major missionary station. These locations are identified through spatial overlay of missionary records onto Ottoman administrative boundaries, using the network hierarchy recorded in the source data: the baseline coding treats as a parent hub any location on which at least one outstation depended (the "Dependent" field of the station records). Appendix Table A.2 documents every parent station and its sanjak assignment, generated programmatically from the replication code. The source records also contain a sparsely filled explicit main-station flag whose coverage conflicts with the station rosters of the missionary literature; Section 6.8 shows the result under three alternative codings, including a reconciled roster that adds İzmir (Smyrna), whose ABCFM station the baseline field misses. The paper distinguishes these hubs from smaller outstations, which reflected more diffuse missionary activity.

To test whether any observed association is specific to concentrated institutional presence rather than missionary activity in general, the analysis also constructs a treatment hierarchy with three mutually exclusive categories: no missionary presence, outstations only, and main station present. Sanjaks are classified by their highest level of missionary presence: main-station sanjaks typically also contained outstations, and "outstations only" denotes outstation presence with no main station. This hierarchy provides a more theory-consistent measure of institutional intensity and helps distinguish durable missionary infrastructure from weaker local exposure.

3.3 Controls and Fixed Effects

The baseline specifications include the natural logarithm of total sanjak population in 1881 and the Christian population share. These variables account for demographic scale and for local religious composition that may have shaped both missionary presence and the placement of state schools. Demographic data are based on the 1881/93 Ottoman census, as compiled and adjusted by Karpat (1985).

Intensive-margin specifications also control for the baseline number of male *rüşdiye* schools in 1876.

All specifications include vilayet fixed effects. This choice is substantively important because the vilayet was the primary level at which provincial administration, budgeting, and implementation were organized. Vilayet fixed effects therefore absorb broad regional heterogeneity in governance, fiscal capacity, and administrative priorities that could otherwise confound sanjak-level comparisons. The paper's identifying variation is thus explicitly within-vilayet rather than across broad provincial regions.

4 Empirical Strategy

The empirical design is associational and comparative rather than quasi-experimental. The paper does not assume that missionary station placement was exogenous. Instead, it asks whether sanjaks exposed to concentrated missionary infrastructure display different patterns of state-school expansion once one conditions on observed demographics, baseline schooling, and provincial administrative structure.

The main intensive-margin specification is

$$\Delta MaleRüşdiye_s = \alpha + \beta_1 MainStation_s + \gamma X_s + \mu_v + \epsilon_s, \tag{1}$$

where the dependent variable is the change in the number of male *rüşdiye* schools in sanjak s between 1876 and 1902–08, $MainStation_s$ indicates the presence of a missionary main station, X_s contains the demographic and baseline controls, and μ_v denotes a full set of vilayet indicators (fixed effects in the cross-sectional sense). Identification is therefore within-vilayet: 13 of the 15 sample vilayets contain both treated and control sanjaks, every treated sanjak lies in such a mixed vilayet, and the two remaining vilayets (Kastamonu and Konya) host no station and contribute only to the estimation of control coefficients. Because the 1876 baseline count enters X_s, this change specification is algebraically equivalent to regressing the 1902–08 level on treatment and the baseline: the treatment coefficient is identical, so the estimate is not an artifact of differencing or of mean reversion.

Inference is handled conservatively through a battery of procedures. The baseline tables report both HC1 heteroskedasticity-robust standard errors and standard errors clustered at the vilayet level to account for within-province correlation, together with exact p-values. Conley spatial standard errors with a 100-kilometer cutoff address spatial dependence across neighboring sanjaks (Conley, 1999). Given the small number of clusters (15 vilayets) and treated units (16 sanjaks), the wild cluster bootstrap with Rademacher weights provides a finite-sample correction, and its 95% confidence interval is obtained by test inversion (Cameron et al., 2008). Randomization inference with 10,000 permutations is reported both unrestricted and stratified within vilayets, the latter respecting the fixed-effects design. Because a cross-sectional spatial design can manufacture significance out of spatially trending unobservables, Section 6.3 additionally reports the two falsification tests proposed by Conley and Kelly (2025): placebo treatments drawn as spatial noise with the autocorrelation structure of the actual treatment, and synthetic outcomes generated as a spatial trend plus noise with no treatment role. To make the power con-

straint explicit: at 80% power and a 5% two-sided test, the design can detect effects no smaller than about 3.0 schools under clustering (2.2 under HC1); these figures use a standard-normal approximation on the clustered and HC1 standard errors, so the detectable effect under the bootstrap and randomization procedures is larger still.

As a complement to the binary treatment indicator, I also estimate a continuous distance specification:

$$\Delta MaleRüşdiye_s = \alpha + \beta_1 \log(1 + \text{Distance}_s) + \gamma X_s + \mu_v + \epsilon_s, \tag{2}$$

where Distance_s is the great-circle distance in kilometers from the sanjak centroid to the nearest main-station sanjak centroid; treated sanjaks have distance zero, so the $\log(1+\cdot)$ transformation requires no arbitrary offset. This specification exploits continuous spatial variation across all 56 sanjaks rather than relying solely on 16 treated units. A negative coefficient would indicate that proximity to main stations is associated with greater school expansion, consistent with a localized competitive response.

To characterize robustness to omitted variables, the paper uses two complementary frameworks. Oster's (2019) proportional-selection parameter δ measures how much more important unobserved confounders would have to be, relative to the included observables, to drive the coefficient to zero; it is reported for a range of $R_{\max}$ values up to 1 (Oster, 2019). Because the Oster framework presumes the included controls are uncorrelated with the omitted factor (Diegert et al., 2022)—a strong assumption here, where the leading unobservable (minority-linked development potential) is plainly correlated with the Christian-share control—Section 6.2 also reports sensitivity bounds in the framework of Cinelli and Hazlett (2020), which parameterize the hypothetical confounder directly by its partial correlations with treatment and outcome and require no such exogeneity of the controls.

5 Results

The central empirical pattern is a contrast between localized deepening and universal frontier expansion. Conditional on observed covariates and vilayet fixed effects, missionary infrastructure is associated with additional school provision in sanjaks where the state already had traction, while the state expanded into every in-sample frontier sanjak regardless of missionary presence. Demographic scale and regional administrative structure account for much of the broader variation in school provision.

5.1 The Intensive Margin: Selective Deepening

The intensive-margin estimates indicate that sanjaks hosting a missionary main station added approximately 1.1 more male *rüşdiye* schools over the period than sanjaks without such a station, conditional on observed covariates and vilayet fixed effects (Table 2). Given a control-group mean change of approximately 1.9 schools, the point estimate corresponds to a differential of nearly 60 percent of the typical expansion—economically large if taken at face value.

The point estimate is unchanged across inference procedures, but its precision depends on the assumed error structure, and the paper reports this dependence in full rather than selecting among procedures. Under HC1 robust standard errors the p-value is 0.16; Conley spatial standard errors at 100 km are nearly identical (SE = 0.79). Under vilayet clustering the standard error widens to 1.08 and the p-value is 0.31; randomization inference yields $p = 0.145$ (0.152 when permutations are stratified within vilayets), and the wild cluster bootstrap yields $p = 0.256$ with a test-inversion 95% confidence interval of $[-0.8, 3.1]$. The estimate is thus not statistically significant at conventional levels under the design-based procedures the paper considers most appropriate for 16 treated sanjaks in 15 clusters. The interval does exclude associations more negative than about -0.8 schools and effects larger than about three; it does not exclude zero. The Moran's I statistic for residual spatial autocorrelation is small and insignificant ($I = 0.05$, $p = 0.80$ at 100 km).

Table 2: Missionary Infrastructure and State-School Expansion (Sanjak Level)

	(1) Intensive Margin (Δ Male Rüşdiye)	(2) Any Increase (Full Sample LPM)
Main Station Dummy	1.101	0.006
HC1 SE	(0.779)	(0.157)
Clustered SE (Vilayet)	(1.082)	(0.202)
Clustered p-value	0.309	0.976
Log Population	Yes	Yes
Christian Share	Yes	Yes
Baseline Schools (1876)	Yes	Yes
Vilayet Fixed Effects	Yes	Yes
Observations	56	56
Dependent Variable Mean	2.29	0.77
Control-Group Mean	1.88	0.78
R^2	0.608	0.423

Notes: Unit of observation is the sanjak. School counts are recorded at the sanjak level in Ottoman administrative records and are recovered as integers from the source data. Istanbul Vilayet excluded as an outlier (population 15× sample median). Column 1 reports the change in the number of male *rüşdiye* schools between 1876 and 1902–08. Column 2 reports a linear probability model for any increase in male *rüşdiye* provision. The extensive margin admits no regression: all 5 in-sample frontier sanjaks (zero schools in 1876; a sixth, Istanbul, is excluded) transitioned to positive provision by 1902–08, with and without missionary stations, so the outcome is constant.

5.2 The Extensive Margin: Universal Frontier Expansion

Among the five in-sample sanjaks that lacked a male *rüşdiye* school in 1876 (a sixth, Istanbul, is excluded from the sample), all five transitioned to positive provision by 1902–08, including both treated and untreated units. Because the outcome is constant, no regression coefficient can be estimated, and with five observations the pattern is descriptive: it is consistent with an empire-wide administrative floor on

provision, but it cannot rule out alternatives. Column 2 of Table 2 offers the complementary full-sample evidence: missionary presence does not predict *whether* a sanjak expanded at all ($\beta = 0.006$, clustered SE = 0.202): treated and control sanjaks recorded an increase at nearly identical rates (75 versus 78 percent), so any differential lies in *how much* provision grew, which is what one would expect if competition operated on the depth rather than the incidence of provision. Uniform entry regardless of missionary presence implies that first establishment followed an empire-wide administrative logic; the margin on which stations could matter—and the only one the design can estimate—is how much provision deepened where it already existed.

5.3 Treatment Hierarchy and Interpretation

The treatment-hierarchy specifications show an ordering of point estimates consistent with the institutional distinction between concentrated and diffuse missionary presence: relative to sanjaks with no missionary presence, those with outstations only show $\beta = 0.95$ and those hosting main stations $\beta = 1.94$ (Table 3). The ordering, however, is not statistically established: a clustered Wald test cannot reject equality of the two coefficients ($p = 0.385$), so the hierarchy should be read as suggestive rather than as a demonstrated dose-response gradient. A direct comparison restricted to missionary-selected sanjaks—main station versus outstations only, with the same controls and fixed effects—yields a similar differential ($\beta = 0.92$, clustered SE = 1.29, $N = 47$; Appendix Table A.9); because both groups were chosen by the missionary organizations themselves, this contrast nets out site-selection criteria common to all missionary locations, though not those specific to main stations. Note also that the reference category here is sanjaks with no missionary presence at all, so these coefficients are not directly comparable to the pooled-control estimate of Table 2.

Table 3: Treatment Hierarchy: Outstations vs. Main Stations (Sanjak Level)

	(1) HC1 Δ Male Rüşdiye	(2) Clustered Δ Male Rüşdiye
Outstation Only	0.949	0.949
	(0.691)	(0.894)
Main Station	1.941	1.941
	(0.834)	(1.100)
Controls	Yes	Yes
Vilayet FE	Yes	Yes
N	56	56
R^2	0.618	0.618
Wald test, main = outstation (p, clustered)		0.385

Notes: Unit of observation is the sanjak. Column 1 reports HC1 robust standard errors. Column 2 clusters at the vilayet level. Controls include log population, Christian share, baseline male *rüşdiye* count, and vilayet fixed effects. Istanbul excluded. The reference category is sanjaks with no missionary presence, so coefficients are not directly comparable to the pooled-control specification of Table 2. The Wald test evaluates equality of the outstation and main-station coefficients under vilayet clustering; the ordering of point estimates is therefore suggestive rather than statistically established.

The pattern across school categories is informative about the channel, with two qualifications. Female *rüşdiye* schools—which were not part of the bureaucratic-training pipeline that missionaries most directly threatened—show no detectable association with main-station presence ($\beta = 0.21$, SE = 0.23, 95% CI $[-0.23, 0.65]$; Appendix Table A.7). Because female provision was sparse (control mean below one school), this is a noisy rather than a precise null: it is consistent with a response confined to the male tier, but also with binding constraints on female expansion—teacher supply, social norms—irrespective of competitive intent. It does, however, align exactly with Amasyalı (2022b), whose kaza-level analysis likewise finds no significant association between missionary *station* presence and female *rüşdiye* counts; her positive findings for women's education operate through missionary schools and competitive dyads, a margin this paper's treatment does not measure. Second, male and female *ibtidaî* (primary) schools both show large but statistically insignificant positive associations ($\beta = 43.8$, SE = 28.3; $\beta = 46.7$, SE = 30.6; Appendix Table A.7). These coefficients are too imprecise to establish a primary-tier response; if one existed, the point estimates suggest it operated broadly rather than through the station contrast alone—at the tier where enrollment fears about Christian conationals were most acute. The cross-category pattern is therefore supportive of a response concentrated in the male bureaucratic-training tier, but it is consistency evidence under the competitive-response reading, not an independent falsification test.

5.4 Distance Gradient

Table 4 presents the continuous distance specification alongside the binary indicator. Column 2 shows that log distance to the nearest main station is negatively associated with school expansion ($\beta = -0.29$,

vilayet-clustered SE = 0.23), indicating that sanjaks closer to main stations experienced greater increases in male *rüşdiye* provision. Column 3 restricts the sample to sanjaks without a station: the gradient among non-hosts alone is negative but imprecise ($\beta = -0.58$, SE = 0.95), so the continuous specification should be understood as largely re-expressing the treated–control contrast rather than as independent evidence of a smooth spatial decay. Figure 2 shows the partial correlation after residualizing on controls.

Table 4: Distance to Nearest Main Station and School Expansion

	(1) Binary	(2) Log Distance (Full Sample)	(3) Log Distance (Control Sanjaks Only)
Main Station Dummy	1.101		
	(1.082)		
Log(1 + Distance in km)		-0.286	-0.575
		(0.231)	(0.952)
Controls	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes
N	56	56	40
R^2	0.608	0.618	0.614

Notes: Dependent variable is the change in male *rüşdiye* schools (1876–1902/08). Distance is the great-circle distance in km between sanjak centroids and the nearest main-station sanjak centroid; the regressor is log(1 + km), so treated sanjaks (distance 0) require no arbitrary offset. Column 3 restricts the sample to sanjaks without a main station, isolating the proximity gradient from the treated–control contrast. Standard errors clustered at the vilayet level in parentheses. Istanbul excluded. A specification combining the binary indicator with log distance is not reported: every treated sanjak sits at distance zero, so the two regressors are nearly collinear in this sample and the indicator's coefficient would be identified only by extrapolating the control-group distance gradient to zero distance.

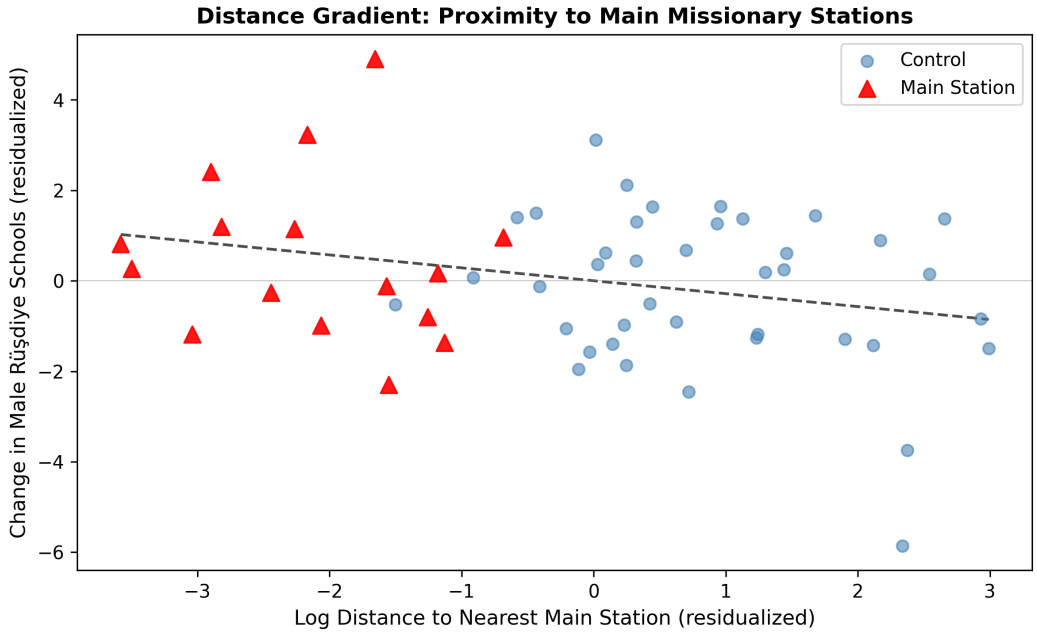


Figure 2: Added-variable plot: log distance to nearest main station and change in male *rüşdiye* schools, residualized on controls and vilayet fixed effects.

6 Robustness and Credibility

The paper evaluates the intensive-margin result using a conservative battery of credibility checks. These exercises do not convert the design into a quasi-experiment, but they characterize—quantitatively where possible—the extent to which the main association is sensitive to inference choices, spatial dependence, unobserved confounding, competing local densities, treatment coding, influential observations, and functional form.

6.1 Inference Sensitivity

Table 5 summarizes the sensitivity of the main intensive-margin estimate to alternative inference procedures. The point estimate of 1.1 additional schools is unchanged across procedures; precision varies as documented in Section 5.1. Conley spatial standard errors are stable across bandwidths (SE = 0.79 at 100 km, 0.79 at 150 km, 0.80 at 200 km), unrestricted randomization inference yields $p = 0.145$ and its within-vilayet stratified counterpart $p = 0.152$ (Figure 3), and the wild cluster bootstrap yields $p = 0.256$ with a test-inversion 95% confidence interval of $[-0.8, 3.1]$. A single reading unifies these numbers: the design is informative about magnitude and sign stability but cannot deliver conventional significance under cluster-robust inference with 15 groups.

Table 5: Inference and Credibility Summary (Intensive Margin, Sanjak Level)

Diagnostic	Estimate
Coefficient (Main Station)	**1.101**
OLS SE	0.769
HC1 Robust SE	0.779
Clustered SE (Vilayet)	1.082
Conley Spatial SE (100 km)	0.787
Randomization Inference, unrestricted (10,000 perms)	$p = 0.145$
Randomization Inference, within-vilayet strata	$p = 0.152$
Wild Cluster Bootstrap (9,999 reps)	$p = 0.256$
95% CI (test inversion)	[-0.82, 3.05]
Oster δ, bivariate → full ($R_{\max} = 0.79$)	9.87
$R_{\max} = 0.9$	6.15
$R_{\max} = 1.0$	4.59
Oster δ, within-vilayet ($R_{\max} = 0.79$)	22.60
Bias-adjusted β^* at $\delta = 1$ ($R_{\max} = 0.79$)	0.989
Identified set $[\beta^*, \hat{\beta}]$	[0.99, 1.10]
LOO Vilayet range	[0.50, 1.41]

Notes: This table summarizes the sensitivity of the main intensive-margin estimate to alternative inference procedures. Conley SEs use a rectangular kernel with cutoff at 100 km (≈ 1 degree). Wild cluster bootstrap follows Cameron et al. (2008) with Rademacher weights, restricted (WCR), clustered at the vilayet level; the 95% confidence interval is obtained by test inversion (collecting all null values not rejected at the 5% level), since the percentile interval of the null-imposed bootstrap distribution is an acceptance region rather than a confidence interval. Stratified randomization inference permutes treatment within vilayets, holding each vilayet's number of treated sanjaks fixed. Oster δ (bivariate → full) compares the regression without controls to the full specification and reports the degree of proportional selection on unobservables required to drive the coefficient to zero, at $R_{\max} \in \{1.3\tilde{R}^2, 0.9, 1.0\}$; the within-vilayet δ conditions on vilayet FEs in the short regression (Oster, 2019).

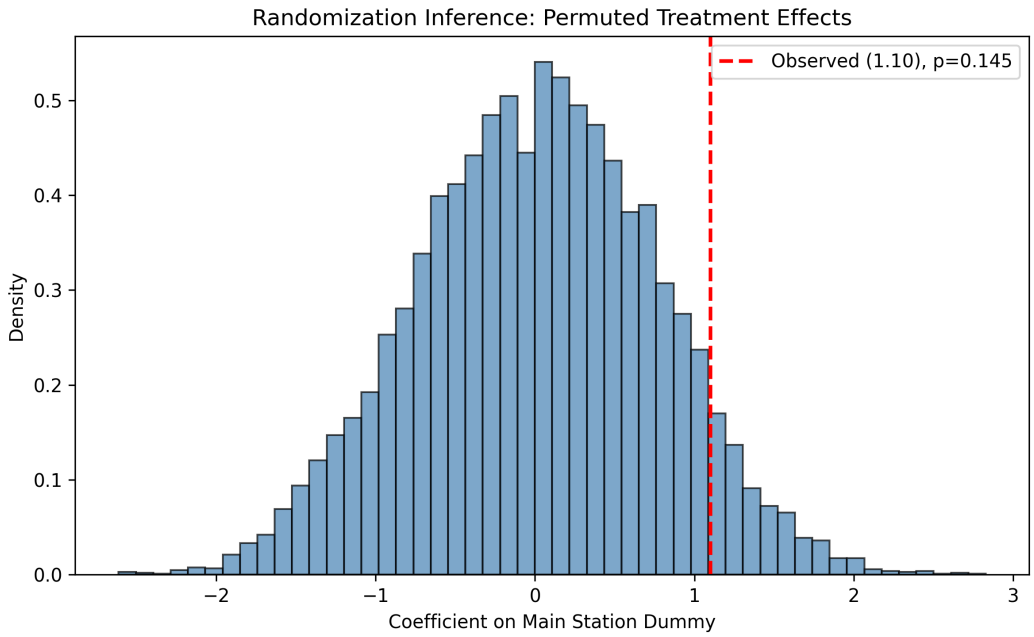


Figure 3: Randomization inference: distribution of permuted treatment effects (10,000 permutations). The dashed red line marks the observed coefficient. The exact two-sided p-value is 0.145.

6.2 Selection on Unobservables

Two complementary frameworks bound the scope for unobserved confounding. In the Oster (2019) framework, moving from the bivariate regression ($\beta = 1.44$, $R^2 = 0.06$) to the full specification ($\beta = 1.10$, $R^2 = 0.61$) leaves the coefficient nearly unchanged while explaining ten times more variance, so the implied degree of proportional selection required to drive the estimate to zero is large: $\delta = 9.9$ at the conventional $R_{\max} = 1.3\tilde{R}^2$, and still 6.2 and 4.6 at $R_{\max} = 0.9$ and 1.0. The bias-adjusted coefficient at equal selection ($\delta = 1$) is $\beta^* = 0.99$, and the identified set $[0.99, 1.10]$ excludes zero. Conditioning on vilayet indicators in the short regression yields $\delta \approx 23$, though much of that figure reflects variance mechanically absorbed by the indicators and should not be read as independent evidence. These calculations assume the included controls are uncorrelated with the omitted factor—exactly the assumption a minority-linked development confounder would violate, since it correlates with Christian share.

The Cinelli and Hazlett (2020) bounds in Table 6 therefore parameterize the confounder directly and dispense with that assumption. The robustness value is 0.21: a confounder would need to explain 21 percent of the residual variance of both treatment and outcome to drive the point estimate to zero. Benchmarking against observed covariates makes this concrete. A confounder as strong as Christian share—which predicts treatment powerfully but adds almost nothing to the outcome residual—barely moves the estimate (adjusted $\beta = 1.00$; even at three times that strength, 0.73). The estimate's genuine vulnerability is to a confounder with the joint strength of log population, the strongest outcome predic-

tor in the model: at that benchmark the adjusted estimate falls to 0.44, and at twice it the sign flips. An unobserved factor resembling "local development potential" at that intensity—net of population itself, Christian and Armenian shares, baseline schooling, and province indicators—cannot be excluded, and the paper flags this as the binding limitation of the design.

Table 6: Sensitivity to Unobserved Confounding (Cinelli and Hazlett, 2020)

Hypothetical confounder	$R^2_{D\sim Z\|X}$	$R^2_{Y\sim Z\|X,D}$	Adjusted $\hat{\beta}$	Adjusted SE
None (baseline)	—	—	1.101	0.769
1× Christian share	0.129	0.003	0.998	0.834
2× Christian share	0.258	0.006	0.879	0.902
3× Christian share	0.388	0.010	0.734	0.991
1× Log population	0.112	0.159	0.437	0.759
2× Log population	0.225	0.318	-0.320	0.731
3× Log population	0.337	0.477	-1.205	0.693
Robustness value (drive $\hat{\beta}$ to 0)			0.209	
Partial R^2 of treatment with outcome			0.052	

Notes: Each row posits an unobserved confounder Z whose partial R^2 with the treatment (given controls) and with the outcome (given controls and treatment) equals the stated multiple of the corresponding partial R^2 of an observed covariate, and reports the worst-case bias-adjusted estimate (Cinelli and Hazlett, 2020). The robustness value is the share of residual variance of both treatment and outcome that a confounder would need to explain to drive the point estimate to zero. Unlike coefficient-stability arguments, these bounds are parameterized directly by the confounder's partial correlations and do not assume the observed controls are uncorrelated with the unobservable. Computed from the OLS covariance of the main specification; under vilayet clustering the estimate is already statistically indistinguishable from zero, so the corresponding robustness value for statistical significance is zero by construction. Istanbul excluded.

6.3 Spatial Falsification Tests

Conley and Kelly (2025) show that spatial cross-sections can produce apparently strong results when both treatment and outcome carry spatial trends, and that low residual autocorrelation does not by itself rule this out. Two features of this setting already work against that failure mode: the treatment is not in fact spatially clustered (Moran's $I = -0.03$ at 100 km, reflecting the deliberate dispersion of mission hubs, roughly one per field), and the paper's inference never rested on asymptotic spatial standard errors alone. Table 7 nonetheless implements both falsification tests. Placebo treatments drawn as spatial noise with the calibrated autocorrelation structure of the actual treatment produce $|t|$-statistics exceeding the observed one in 27 percent of draws; synthetic outcomes generated as a quadratic spatial trend plus correlogram-matched noise, regressed on the *actual* treatment, do so in 25 percent of draws. Both rejection rates sit close to the paper's design-based p-values (0.26–0.31), indicating that those p-values do not understate the probability of obtaining such an association from spatial structure alone. Adding a quadratic spatial basis to the main specification, the constructive remedy Conley and Kelly (2025) recommend, leaves the coefficient at 1.04 (clustered SE 1.20).

Table 7: Spatial Falsification Tests (Conley and Kelly, 2025)

Diagnostic	Value
Moran's I of treatment indicator (100 km)	-0.029
Calibrated spatial range of placebo fields	25 km
(a) Spatial-noise placebo treatments: $\Pr(\lvert t_{\text{placebo}}\rvert \geq \lvert t_{\text{obs}}\rvert)$	0.271
(b) Synthetic outcomes (trend + matched noise): $\Pr(\lvert t\rvert \geq \lvert t_{\text{obs}}\rvert)$	0.245
(c) Estimate with quadratic spatial basis	1.043 (1.195)
Baseline estimate for comparison	1.101 (1.082)

Notes: Test (a) replaces the main-station indicator with placebo treatments generated by thresholding Gaussian random fields with exponential covariance (range 25 km, calibrated so placebo assignments match the observed spatial autocorrelation of the true treatment), holding the number of treated sanjaks fixed at 16; the entry is the share of 2,000 placebo draws whose cluster-robust $|t|$ exceeds the observed $|t| = 1.02$. Test (b) regresses synthetic outcomes—a quadratic spatial trend in longitude/latitude plus noise with the correlogram of the detrended outcome—on the *actual* treatment and controls; it asks whether a spatial trend-plus-noise process with no treatment effect reproduces the observed test statistic. Row (c) adds a quadratic polynomial in longitude and latitude to the main specification. Clustered standard errors in parentheses. Istanbul excluded.

6.4 Administrative and Commercial Importance

A natural concern is that missionaries selected administratively important sanjaks—provincial capitals, trade hubs, or larger administrative units—that would have received more schools regardless of missionary presence. The overlap is real: ten of the sixteen treated sanjaks were vilayet capital seats.

Table 8 conditions on the observable dimensions of this concern: an indicator for the vilayet capital seat itself, an indicator for the most populous sanjak in each vilayet, the number of constituent kazas, and all three jointly; the horse-race table below adds an indicator for the commercial centers recorded in the source GIS (İzmir, Bursa, and Ankara within the sample). The main-station coefficient ranges from 1.05 to 1.66 across these specifications and never attenuates materially. Because these are coarse proxies, this stability bounds only the observable component of administrative importance; it cannot rule out confounding through dimensions the proxies miss, which is precisely what the sensitivity bounds of Section 6.2 are for. Measured facets of administrative and commercial centrality do not absorb the association; unmeasured facets remain the design's principal vulnerability.

Table 8: Controlling for Administrative Importance

	(1) Baseline	(2) + Capital Seat	(3) + Largest	(4) + N Kazas	(5) + All
Main Station	1.101	1.211	1.400	1.053	1.663
	(1.082)	(1.103)	(1.058)	(1.033)	(1.240)
Controls	Yes	Yes	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes	Yes	Yes
Vilayet capital seat	No	Yes	No	No	Yes
Largest sanjak in vilayet	No	No	Yes	No	Yes
Number of kazas	No	No	No	Yes	Yes
R^2	0.608	0.609	0.647	0.654	0.685
N	56	56	56	56	56

Notes: This table addresses the concern that missionaries selected administratively important sanjaks that would have received more schools regardless. Column 2 adds an indicator for the vilayet capital seat (the *merkez* sanjak hosting the provincial administration; independent sanjaks are coded as their own seats). Column 3 adds an indicator for the most populous sanjak within each vilayet, Column 4 the number of constituent kazas, and Column 5 all three. Because these are coarse proxies, coefficient stability across columns bounds only the observable component of administrative importance. Standard errors clustered at the vilayet level in parentheses. Istanbul excluded.

6.5 Missionary Hubs versus Indigenous Christian Schooling

A distinct alternative reading of the main estimate is that stations merely mark places with dense Christian—especially Armenian—educational activity, and that the state was responding to that local educational environment (or to the communities themselves) rather than to missionary infrastructure specifically (Amasyalı, 2022a). The source GIS package permits a direct test: it includes 1,313 geocoded Ottoman-era Armenian community schools and 1,067 Greek schools from the Hrant Dink Foundation cultural-heritage layer, together with Armenian churches (religious rather than educational density) and the 1881/93 census counts of Protestant converts.

Table 9 runs the horse race. Conditioning on the log density of Armenian schools leaves the main-station coefficient essentially unchanged (1.03, SE 1.16); adding Greek schools (1.12), Armenian churches (1.01), Protestant population share (0.90), or commercial-center status (1.11) does not alter the picture, and none of the density measures itself predicts *rüşdiye* growth conditional on the controls. Table 10 performs the complementary demographic exercise, splitting the pooled Christian-share control into Armenian and non-Armenian components: the main-station coefficient rises slightly (to 1.24–1.26) rather than falling. Within the limits of observables, then, the association is specific to missionary hubs and is not a re-expression of local Christian educational, religious, or demographic intensity. This also sharpens the interpretation of the treatment-hierarchy ordering in Section 5.3: if stations proxied only for the salience of the local Christian community, the Armenian-school and church densities should have absorbed the station coefficient, and they do not.

Table 9: Missionary Hubs vs. Indigenous Christian Educational Density

	(1)	(2)	(3)	(4)	(5)	(6)
Main station	1.101	1.026	1.123	1.012	0.901	1.109
	(1.082)	(1.158)	(1.117)	(1.183)	(1.089)	(1.117)
Log(1 + Armenian schools)		0.250	0.211			
		(0.292)	(0.322)			
Log(1 + Greek schools)			0.448			
			(0.441)			
Log(1 + Armenian churches)				0.239		
				(0.322)		
Protestant population share					43.069	
					(202.885)	
Commercial center						0.341
						(2.255)
Controls, Vilayet FE	Yes	Yes	Yes	Yes	Yes	Yes
N	56	56	56	56	56	56
R^2	0.608	0.611	0.622	0.611	0.610	0.608

Notes: This table separates the missionary-hub association from the density of indigenous Christian education and religious infrastructure. Armenian schools (1,313 geocoded Ottoman-era community schools) and Greek schools (1,067) are from the Hrant Dink Foundation cultural-heritage layer included in the source GIS package (Amasyalı, 2020); Armenian churches proxy religious rather than educational density. Protestant population share is from the 1881/93 census; it is a proportion with almost no sample variation (mean 0.003), so its coefficient is estimated on negligible support and only its statistical insignificance is informative. Commercial centers (İzmir, Bursa, Ankara within the sample) are from the package's commercial-centers layer. If the main-station association merely re-expressed the intensity of local Christian educational activity, conditioning on these densities should absorb it. Dependent variable: change in male *rüşdiye* schools, 1876–1902/08. All specifications include log population, Christian share, baseline schools, and vilayet fixed effects. Standard errors clustered at the vilayet level. Istanbul excluded.

Table 10: Disaggregating the Christian-Share Control: Armenian vs. Other Christian Population

	(1) Pooled	(2) Disaggregated	(3) Armenian Only
Main Station	1.101	1.243	1.258
	(1.082)	(1.127)	(1.095)
Armenian Share	—	-3.10	-3.17
		(3.41)	(3.26)
Other Christian Share	—	0.53	—
		(2.86)	
Other controls, Vilayet FE	Yes	Yes	Yes
N	56	56	56
R^2	0.608	0.613	0.613

Notes: Missionary main stations targeted Armenian communities specifically, and Armenian presence proxies for minority-linked human capital and commercial development (Arbatlı and Gökmen, 2023). Column 2 splits the pooled Christian share into Armenian and non-Armenian Christian shares; Column 3 retains only the Armenian share. Standard errors clustered at the vilayet level. Istanbul excluded.

6.6 Selection on Observables: Reweighting

Treated and control sanjaks differ on baseline observables: the normalized difference on Christian population share is 0.83, far beyond the 0.25 threshold of Imbens and Rubin (2015) (Appendix Table A.4), which means treated and control units overlap poorly on this dimension. Reweighting exercises should therefore be read with the overlap problem in view. Inverse-probability weighting on log population, Christian share, and baseline schools yields 1.03 (HC1 SE = 0.71), with a Kish effective sample size of 38 of 56 and a maximum normalized weight of 3.1 (Appendix Table A.6); entropy balancing, which matches treated covariate means exactly, yields 1.25 (clustered SE 1.08). The most conservative comparison restricts the sample to the common support of Christian share, where treated and control sanjaks are genuinely comparable: $\beta = 0.82$ (clustered SE 1.21, $N = 36$). The reweighted point estimates thus range from 0.8 to 1.25—consistently positive, uniformly imprecise, and somewhat smaller where the comparison is cleanest, a gradient the reader should weigh.

6.7 Heterogeneity and Influential Units

Leave-one-vilayet-out checks show that the intensive-margin coefficient remains positive across all 15 subsamples, ranging from 0.50—when Halep, the most influential province, is dropped—to 1.41 (Figure 4). Leave-one-treated-sanjak-out diagnostics across all 16 treated sanjaks show that no single unit drives the result: the coefficient ranges from 0.76 (dropping Maraş) to 1.36 (dropping Bursa). Re-including Istanbul leaves the coefficient at 1.10 by construction—Istanbul is the sole sanjak of its vilayet, so its own indicator absorbs it completely and the accompanying jump in R^2 (to 0.99) is mechanical; the check is reported for completeness but is uninformative, which is why Istanbul is excluded throughout.

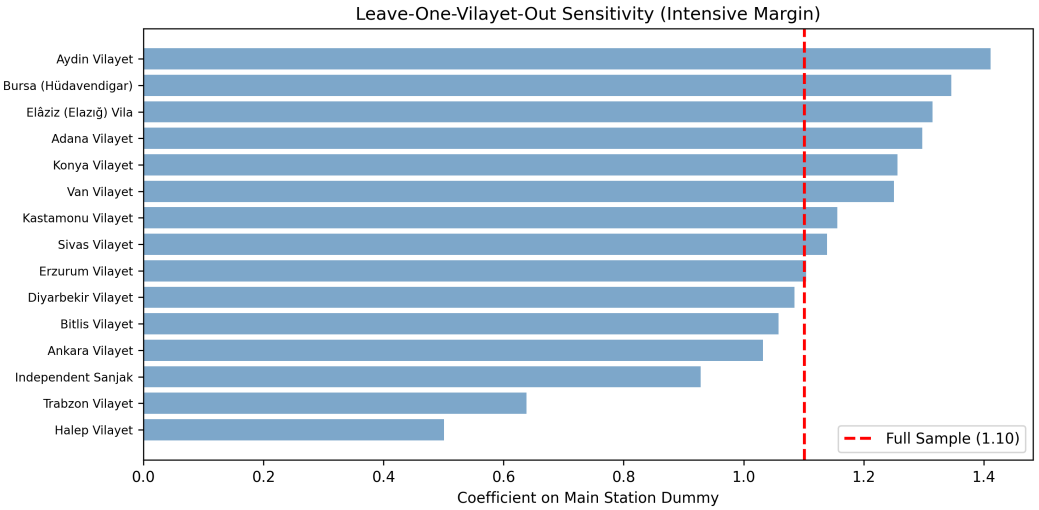


Figure 4: Leave-one-vilayet-out coefficient estimates for the intensive-margin specification. The dashed red line marks the full-sample estimate.

6.8 Treatment Coding and Alternative Rosters

The treatment coding derives from the network hierarchy in the source records, and Table 11 shows the estimate under the three codings those records support. The baseline Dependent-field roster (16 treated sanjaks) yields 1.10. The reconciled roster that adds İzmir—whose Smyrna station is documented in the missionary literature and flagged in the source attribute table, but missed by the Dependent field—yields 1.39 (SE 1.14); reclassifying Saruhan (Manisa), whose main-station status is the least certain, changes nothing (1.39). A roster restricted to the source file's sparsely filled explicit flag (nine stations) yields a null (−0.35), but that coding moves eight sanjaks hosting documented stations—Maraş, Kayseri, Trabzon, Bitlis, Adana, Kozan, İzmit, Saruhan—into the control group, so its contrast is not station-versus-no-station; the paper reports it for transparency and reads the reconciled roster as the historically defensible upper bound. A further variant using only the seven geometry-joined stations (dropping all nine manually matched hubs) yields 1.60 (SE 1.42; Appendix Table A.9). The full station-by-station assignment, generated programmatically from the replication code, is in Appendix Table A.2; validating the roster against the ABCFM annual station lists remains the natural next step for this data infrastructure.

Table 11: Alternative Station Rosters from the Source Records

Roster	$\hat{\beta}$	Clust. SE	Treated	N
Baseline: Dependent-field roster (16 treated)	1.101	(1.082)	16	56
Explicit 'Main Station' flag (9 stations)	-0.352	(1.020)	9	56
Union / reconciled roster (+İzmir, 17 treated)	1.387	(1.139)	17	56

Notes: The source GIS records encode the missionary network's hierarchy in two imperfectly overlapping ways. The baseline coding treats as a hub any location on which at least one outstation depended (the "Dependent" field). An explicit "Main Station" attribute flag marks only nine stations; it adds Smyrna (İzmir), which the Dependent field misses, but it also classifies stations that the missionary literature documents as full stations—Marash, Cesarea (Kayseri), Trebizond, Bitlis, Adana, Hadjin—as outstations or leaves them uncoded, moving eight sanjaks that hosted documented stations into the control group; its estimate is therefore not interpretable as a station-versus-no-station contrast. The union roster adds İzmir to the baseline coding and corresponds most closely to the ABCFM station rosters in the missionary literature (Stone, 2006). Dependent variable: change in male *rüşdiye* schools, 1876–1902/08; all specifications include log population, Christian share, baseline schools, and vilayet fixed effects. Standard errors clustered at the vilayet level. Istanbul excluded.

6.9 Additional Robustness

Appendix Table A.9 collects the remaining checks in one place; three deserve comment. First, a placebo regression of the *1876 baseline level* on treatment yields a null (−0.30, SE 0.89): station sanjaks did not simply start with more recorded schools, conditional on controls. Second, continuous-intensity measures behave as the presence-based reading predicts: missionary locations per capita show no association (−0.01, SE 0.07)—consistent with Amasyalı (2022b), who finds the same for density measures—while the presence contrast carries the association. Third, an SLX term counting treated sanjaks within 150 km enters positively (0.57, SE 0.35) with the own-station coefficient intact (1.41), suggesting if anything mild positive spillovers to neighbors, which would attenuate rather than inflate the treated–control contrast. The remaining rows—per-capita outcome, strict intensive margin on the 51 positive-baseline sanjaks, excluding capital seats or the most populous quartile, and the treatment-by-baseline interaction—tell the same story of a stable, imprecise positive association.

Count models match the linear results: Poisson and negative binomial specifications on the 1902–08 level (the change cannot be modeled as a count, as six sanjaks record declines) yield an average marginal effect of 1.14 with a delta-method standard error of 1.02 (Appendix Table A.8). An exposure variant that forces a unit elasticity on Muslim population reverses the sign, but the data reject that restriction decisively (free elasticity 0.37, SE 0.13): one *rüşdiye* served an entire district, so school counts rise far less than proportionally with population, and the free-coefficient specification is the appropriate one. Kaza-level regressions with sanjak-clustered standard errors produce a proportionally consistent coefficient once one accounts for the distribution of sanjak-level school counts across constituent kazas (Appendix Table A.10).

7 Conclusion

This paper asks where the Hamidian state deepened its flagship secular schooling tier, and finds that the geography of that deepening was aligned with the geography of its most visible institutional rivals. Sanjaks hosting American missionary main stations added roughly one more male *rüşdiye* school over 1876–1902/08 than observably similar sanjaks in the same province—a differential of nearly 60 percent of the typical expansion—and that association is specific to missionary hubs: it survives conditioning on the density of Armenian and Greek community schools, on Armenian religious infrastructure, on Protestant converts, on capital-seat and commercial status, and on alternative rosters of the stations themselves. Meanwhile the state established at least minimal provision everywhere, including every frontier sanjak without missionaries. These patterns are consistent with the documented administrative anxiety about missionary schools (Deringil, 1998; Fortna, 2002) having left a spatial imprint on the depth, rather than the reach, of the state's response; they do not establish that competition caused it.

The central limitation is statistical power: with 16 treated sanjaks and 15 clusters, the association is not statistically significant under cluster-robust or design-based inference (p between 0.15 and 0.31), and the smallest effect the design could reliably detect is about three schools. The spatial falsification tests show these p-values are not understated, and the sensitivity analysis quantifies the exposure that remains: a confounder as strong as the Christian-share control would leave the estimate intact, but one with the joint strength of log population—an unobserved "development potential" factor—would absorb it. Two further caveats stand. The design observes no pre-1876 trajectory, so a steeper prior trend in station sanjaks cannot be excluded, although stations do not predict the 1876 level of provision; and the war of 1877–78, refugee resettlement, and the massacres of 1894–96 fall inside the outcome window with sanjak-level incidence the design cannot fully absorb.

Two extensions would move this literature decisively. Digitizing intermediate salname years would convert the long difference into a panel and allow trajectories, not just changes, to be compared around the missionary geography. And matching the station roster to the ABCFM annual reports would put the treatment coding—here validated only against the source records' internal hierarchy—on independent documentary footing. Both are feasible with sources already identified, and both would test whether the pattern documented here bears the weight of a causal reading.

Data Availability

The analysis uses the publicly available Ottoman Anatolia Schooling Historical GIS (Amasyalı, 2020). A replication package containing all code that constructs the estimation dataset from the source GIS files and regenerates every table and figure in the paper accompanies the manuscript and will be deposited in the Harvard Dataverse repository upon acceptance, with the DOI cited here.

References

Amasyalı, E. (2020). Ottoman anatolia schooling hgis. https://github.com/emreamasyali/ottoman-anatolia-schooling-hgis.

Amasyalı, E. (2022a). Indigenous responses to protestant missionaries: Educational competition and economic development in ottoman turkey. *European Journal of Sociology*, 63(1):39–86.

Amasyalı, E. (2022b). Protestant missionary education and the diffusion of women's education in ottoman turkey: A historical gis analysis. *Social Science History*, 46(1):173–222.

Arbatlı, C. E. and Gökmen, G. (2023). Human capital transfers and sub-national development: Armenian and greek legacy in post-expulsion turkey. *Journal of Economic Growth*, 28(1):1–43.

Bai, Y. and Kung, J. K.-s. (2015). Diffusing knowledge while spreading god's message: Protestantism and economic prosperity in china, 1840–1920. *Journal of the European Economic Association*, 13(4):669–698.

Berkes, N. (1964). *The Development of Secularism in Turkey*. McGill University Press.

Besley, T. and Persson, T. (2011). *Pillars of Prosperity: The Political Economics of Development Clusters*. Princeton University Press.

Cagé, J. and Rueda, V. (2016). The long-term effects of the printing press in sub-Saharan Africa. *American Economic Journal: Applied Economics*, 8(3):69–99.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.

Cinelli, C. and Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B*, 82(1):39–67.

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

Conley, T. G. and Kelly, M. (2025). The standard errors of persistence. *Journal of International Economics*, 153:104027.

Deringil, S. (1998). *The Well-Protected Domains: Ideology and the Legitimation of Power in the Ottoman Empire, 1876–1909*. I.B. Tauris.

Diegert, P., Masten, M. A., and Poirier, A. (2022). Assessing omitted variable bias when the controls are endogenous. arXiv:2206.02303.

Evered, E. Ö. (2012). *Empire and Education under the Ottomans: Politics, Reform, and Resistance from the Tanzimat to the Young Turks*. IB Tauris.

Fortna, B. C. (2002). *Imperial Classroom: Islam, the State, and Education in the Late Ottoman Empire*. Oxford University Press.

Gallego, F. A. (2010). Historical origins of schooling: The role of democracy and political decentralization. *The Review of Economics and Statistics*, 92(2):228–243.

Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

Karpat, K. H. (1985). *Ottoman Population, 1830–1914: Demographic and Social Characteristics*. University of Wisconsin Press.

Kieser, H.-L. (2010). *Nearest East: American Millennialism and Mission to the Middle East*. Temple University Press.

Kocabaşoğlu, U. (1989). *Kendi Belgeleriyle Anadolu'daki Amerika: 19. Yüzyılda Osmanlı İmparatorluğu'ndaki Amerikan Misyoner Okulları*. Arba.

Kodaman, B. (1988). *II. Abdülhamid Devri Eğitim Sistemi*. Türk Tarih Kurumu.

Ortaylı, İ. (2000). *İmparatorluğun En Uzun Yüzyılı*. İletişim.

Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.

Pamuk, Ş. (2006). Estimating economic growth in the middle east since 1820. *The Journal of Economic History*, 66(3):809–828.

Somel, S. A. (2001). *The Modernization of Public Education in the Ottoman Empire, 1839–1908: Islamization, Autocracy, and Discipline*. Brill.

Stone, F. A. (2006). *Academies for Anatolia: A Study of the Rationale, Program, and Impact of the Educational Institutions Sponsored by the American Board in Turkey, 1830–2005*. Caddo Gap Press.

Tilly, C. (1990). *Coercion, Capital, and European States, AD 990–1990*. Blackwell.

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

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

Zürcher, E.-J. (2004). *Turkey: A Modern History*. IB Tauris.

Appendices

A Variable Definitions and Data Construction

Table A.1: Variable Definitions and Sources

Variable	Definition	Source
Panel A: Outcomes		
Δ Male Rüşdiye	Change in male *rüşdiye* school count, 1876–1902/08, at the sanjak level	Maarif Salnamesi via Amasyali (2020)
Zero-to-Positive	Indicator: sanjak had 0 male *rüşdiye* in 1876 and >0 by 1902/08	Maarif Salnamesi via Amasyali (2020)
Any Increase	Indicator: sanjak experienced any positive change in male *rüşdiye*	Maarif Salnamesi via Amasyali (2020)
Panel B: Treatment Variables		
Main Station	Indicator: ≥1 parent missionary station in sanjak. Parent stations identified as hubs listed in the "Dependent" field of missionary records	Missionary Locations shapefile (Amasyali 2020)
Outstation Only	Indicator: missionary presence but no main station	Missionary Locations shapefile (Amasyali 2020)
Log Distance	Log geodesic distance (km) from sanjak centroid to nearest main station	Computed from GIS data
Panel C: Controls		
Log Population	Log of total sanjak population (sum of kaza-level totals)	Karpat (1985), 1881/93 Census
Christian Share	Christians as share of total sanjak population	Karpat (1985), 1881/93 Census
Baseline Schools	Number of male *rüşdiye* schools in 1876	Maarif Salnamesi via Amasyali (2020)
Vilayet FE	Fixed effects for 15 vilayets (incl. independent sanjaks)	Ottoman administrative geography

Table A.2 documents the identification and sanjak assignment of all parent missionary stations, generated programmatically from the replication code so that it reflects the treatment assignment actually used in estimation. Of the 17 parent locations identified from the "Dependent" field in the missionary records, 7 have their own point geometry in the GIS data and are matched to sanjaks via spatial join; the remaining 9 Anatolian stations are matched manually to their documented kaza locations. Constantinople is excluded with Istanbul Vilayet. Alternative rosters are examined in Section 6.8.

Table A.2: Parent Missionary Stations: Identification and Sanjak Assignment

Station (source spelling)	Method	Kaza	Sanjak
Adana	Spatial join	Adanamerkez Kaza	Adana Sanjak
Aintab	Manual match	Antep Kaza	Halep Sanjak
Bitllis	Manual match	Bitlismerkez Kaza	Bitlis Sanjak
Broosa	Spatial join	Bursamerkez Kaza	Bursa Sanjak
Cesarea	Spatial join	Kayserimerkez Kaza	Kayseri Sanjak
Constantinople	Excluded (Istanbul)	—	—
Erzroom	Manual match	Erzurummerkez Kaza	Erzurum Sanjak
Hadjin	Spatial join	Hacin Kaza	Kozan Sanjak
Harpoot	Manual match	Harput Kaza	Elaziz Sanjak
Manissa	Spatial join	Manisa Kaza	Saruhan Sanjak
Marash	Spatial join	Marasmerkez Kaza	Maraş Sanjak
Mardin	Manual match	Mardinmerkez Kaza	Mardin Sanjak
Marsovan	Manual match	Merzifon Kaza	Amasya Sanjak
Nicomedia	Spatial join	Izmitmerkez Kaza	Izmit (Special) Sancak
Sivas	Manual match	Sivasmerkez Kaza	Sivas Sanjak
Trabzon	Manual match	Trabzonmerkez Kaza	Trabzon Sanjak
Van	Manual match	Vanmerkez Kaza	Van Sanjak

Notes: Parent (main) stations are the locations named in the "Dependent" field of the missionary GIS records (Amasyalı, 2020), i.e., hubs on which at least one outstation depended. Stations with their own point geometry are assigned to kazas by spatial join; the remainder are matched manually to their historically documented kaza locations. Station spellings follow the source records. This table is generated programmatically from the replication code, so it reflects the treatment assignment actually used in estimation. Notably, İzmir (Smyrna) does not appear as a parent hub in the source records' Dependent field and is therefore untreated in the baseline coding; robustness checks reclassifying İzmir as treated are reported in Table 11.

B Full Regression Output

Table A.3: Full Regression Output with Control Coefficients

	(1) HC1	(2) Vilayet Clustered
Main Station	1.101	1.101
SE	(0.779)	(1.082)
Log Population	2.069	2.069
SE	(0.761)	(0.805)
Christian Share	-0.980	-0.980
SE	(1.979)	(2.589)
Baseline Schools (1876)	-0.606	-0.606
SE	(0.153)	(0.187)
Vilayet FE	Yes	Yes
R^2	0.608	0.608
N	56	56
Dep. Var. Mean	2.29	2.29

Notes: Full coefficient estimates from the main intensive-margin specification. The dependent variable is the change in male *rüşdiye* schools (1876–1902/08). Vilayet fixed effects included but not reported. Istanbul Vilayet excluded.

C Covariate Balance

Table A.4: Balance of Covariates: Main Station vs. Control Sanjaks

	Control Mean (SD)	Treated Mean (SD)	Norm. Diff.
Log Population	11.94	12.23	0.46
	(0.62)	(0.62)	
Christian Share	0.14	0.24	0.83
	(0.14)	(0.12)	
Armenian Share	0.07	0.17	0.85
	(0.09)	(0.14)	
Baseline Schools (1876)	3.48	3.94	0.16
	(2.63)	(3.07)	
Number of Kazas	5.42	6.81	0.51
	(2.24)	(3.15)	
Population	183751.64	243712.69	0.46
	(113879.97)	(145472.61)	
Observations	40	16	

Notes: Normalized differences computed as $(\bar{x}_T - \bar{x}_C)/\sqrt{(s_T^2 + s_C^2)/2}$ following Imbens and Rubin (2015). Values above 0.25 in absolute value suggest meaningful imbalance.

D Sequential Controls

Table A.5: Coefficient Stability Under Sequential Addition of Controls

Specification	$\hat{\beta}$	HC1 SE	R^2	N
Bivariate	1.437	(1.055)	0.058	56
+ Log Pop	1.083	(1.050)	0.134	56
+ Christian Share	1.418	(1.164)	0.148	56
+ Baseline Schools	1.040	(1.019)	0.363	56
+ Vilayet FE	1.101	(0.779)	0.608	56

Notes: Each row adds one control to the previous specification. The stability of the coefficient across rows indicates robustness to observable confounders, complementing the Oster (2019) bound. Istanbul excluded.

E Inverse-Probability Weighting

Table A.6: Addressing Covariate Imbalance: IPW and Trimmed Specifications

	(1) Unweighted	(2) IPW	(3) Trimmed
Main Station	1.101	1.026	0.821
	(1.082)	(0.712)	(1.210)
Controls	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes
N	56	56	36
Effective sample size	—	38.0	—
Max. normalized weight	—	3.13	—
Post-wt. norm. diff. (Christian share)	—	0.06	—
R^2	0.608	0.719	0.807

Notes: Column 1 reproduces the baseline estimate. Column 2 uses inverse-probability (ATE) weights from a logit model of treatment on log population, Christian share, and baseline schools; weights are capped at their 95th percentile and normalized to mean one. Effective sample size is the Kish measure $(\sum w)^2 / \sum w^2$. Column 3 restricts the sample to the common support of the Christian-share distribution (dropping control sanjaks below the treated minimum and above the treated maximum, and vice versa); because treated and control sanjaks overlap poorly on Christian share, Column 3 is the comparison that weights only genuinely comparable units and should be read as the more conservative estimate. Standard errors: vilayet-clustered (Cols. 1, 3), HC1 robust (Col. 2). Istanbul excluded.

F Placebo and Secondary Outcomes

Table A.7: Placebo and Secondary Outcomes: Female and Primary Schooling

	(1)	(2)	(3)	(4)	(5)
	Female Rüşdiye (Placebo)		Female *İbtidaî*		Male *İbtidaî*
	Binary	Distance	Binary	Distance	Binary
Main Station	0.207		46.740		43.774
	(0.225)		(30.588)		(28.338)
Log(1 + Distance km)		-0.046		-10.297	
		(0.047)		(8.239)	
Controls	Yes	Yes	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes	Yes	Yes
N	56	56	56	56	56
Dep. Var. Mean	0.54	0.54	65.0	65.0	95.1
R^2	0.495	0.496	0.610	0.611	0.554

Notes: Outcomes are 1902–08 levels (female and primary school counts are not recorded for 1876, so no change specification is possible for these categories). Columns 1–2 use female *rüşdiye* schools as the placebo outcome: under the interpretation that the state response operated through the male bureaucratic-training channel, no association with female *rüşdiye* provision is expected. Columns 3–4 use female *ibtidaî* (primary) schools and Column 5 male *ibtidaî* schools as secondary outcomes. Standard errors: vilayet-clustered (Cols. 1, 3, 5), OLS (Cols. 2, 4). Istanbul excluded.

G Count Data Models

Table A.8: Count Data Models: Poisson and Negative Binomial

	(1) Poisson	(2) Poisson (Exposure)	(3) Neg. Binomial
Main Station	0.187	-0.158	0.187
	(0.161)	(0.085)	(0.151)
Controls	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes
Log(Muslim pop.) offset	No	Yes	No
N	56	56	56
Avg. Marginal Effect	1.14	—	1.10
Delta-method SE	(1.02)	—	—

Notes: Dependent variable is the number of male *rüşdiye* schools in 1902–08 (count level; the change cannot be modeled as a count because six sanjaks record declines). Baseline schools (1876) included as control, so the specification captures growth conditional on the starting stock. Coefficients are log-rate ratios. Column 2 replaces the free log-population coefficient with a log Muslim population exposure offset, expressing provision per capita of the primary served population; the data strongly reject the implied unit population elasticity (free elasticity 0.37, SE 0.13: school counts rise far less than proportionally with population, as one *rüşdiye* served an entire district), so Column 1 is the preferred count specification. The average marginal effect for the binary treatment is $\frac{1}{N}\sum_i[\hat{\mu}_i(D=1)-\hat{\mu}_i(D=0)]$ with a delta-method standard error based on the cluster-robust covariance. Standard errors: vilayet-clustered (Cols. 1–2), ML (Col. 3). Istanbul excluded.

H Curated Robustness Battery

Table A.9: Curated Robustness Battery (Intensive Margin)

Specification	Coefficient	$\hat{\beta}$	Clust. SE	N
Baseline	Main station	1.101	(1.082)	56
Per-capita outcome (Δ per 100k)	Main station	0.353	(0.622)	56
Strict intensive margin (baseline > 0)	Main station	1.056	(1.556)	51
Main station vs. outstations only	Main station	0.921	(1.289)	47
Alt. roster: İzmir (Smyrna) treated	Main station	1.387	(1.139)	56
Alt. roster: +İzmir, −Saruhan	Main station	1.392	(1.186)	56
Geometry-joined stations only	Main station	1.602	(1.418)	56
Exclude vilayet capital seats	Main station	2.168	(1.880)	39
Exclude top-quartile population	Main station	1.504	(1.329)	42
Entropy balancing	Main station	1.247	(1.081)	56
SLX: own treatment (neighbor count included)	Main station	1.412	(1.152)	56
SLX: treated neighbors within 150 km	Neighbor count	0.566	(0.345)	56
Intensity dose: locations per 100k	Missionary density	-0.009	(0.068)	56
Intensity dose: log(1 + locations)	Log locations	0.310	(0.401)	56
Placebo: 1876 baseline level as outcome	Main station	-0.296	(0.889)	56
Interaction: main station × baseline (centered)	Interaction	-0.307	(0.241)	56
main-station effect at mean baseline	Main station	1.207	(1.055)	56

Notes: All rows include log population, Christian share, baseline schools (except the interaction and baseline-placebo rows, as indicated), and vilayet fixed effects; standard errors are clustered at the vilayet level. Alternative rosters address the historically uncertain status of İzmir (Smyrna hosted an ABCFM station city but does not appear as a parent hub in the source records' Dependent field) and of Saruhan (Manisa). The main-station-versus-outstations-only row restricts the sample to sanjaks with some missionary presence, so the comparison group shares the missionaries' own site-selection criteria. The geometry-joined-only row drops the nine manually matched hubs from the treatment. Entropy balancing reweights control sanjaks to match treated covariate means exactly. The SLX rows add the count of treated sanjaks within 150 km. The dose rows replace the binary indicator with continuous exposure measures. The 1876-placebo row asks whether main stations predict the pre-period *level* of provision. Istanbul excluded throughout.

I Kaza-Level Robustness

Table A.10: Kaza-Level Robustness (Sanjak-Clustered Standard Errors)

	(1) Baseline	(2) Hierarchy	(3) Armenian	(4) Any Increase
Main Station	0.053	0.049	0.029	-0.040
	(0.078)	(0.078)	(0.076)	(0.060)
Outstation Only		0.015		
		(0.044)		
Controls	Yes	Yes	Yes	Yes
Vilayet FE	Yes	Yes	Yes	Yes
SE Clustering	Sanjak	Sanjak	Sanjak	Sanjak
N	326	326	326	326
R^2	0.494	0.494	0.496	0.539

Notes: Unit of observation is the kaza. School data is recorded at the sanjak level in Ottoman administrative records; standard errors are therefore clustered at the sanjak level to reflect the true unit of variation. Istanbul excluded. Controls include log population, Christian share (or Armenian share in Col. 3), baseline male *rüşdiye* count, and vilayet fixed effects.