Identifying the Diffusion and Suppression of Political Naming in Turkey with a Matched Interrupted Time Series Design

Ramses Llobet*

Abstract

What can naming practices reveal about political expression under authoritarian rule? We argue that a name's political association can encourage adoption, while repression may increase its perceived costs. Using Turkish birth registration records from 1960 to 1990, we examine Deniz naming around the persecution and execution of the revolutionary activist Deniz Gezmiş and the period of repression surrounding the 1980 military coup. To assess this account against broader cultural influences, we combine province and district models of departures from a projected trajectory with comparisons against forty names associated with the 1930s language reform, matched on pre-1971 naming patterns. Deniz's naming rate rises during the 1970s and contracts around the coup while remaining above the matched-name benchmark through 1990. Its joint growth and contraction statistic ranks first within the selected pool in both analyses. Larger modeled contractions occur where the earlier increase was greater and prior support for the Turkish Workers' Party was stronger. These findings support a political interpretation of naming, conditional on the projection rules, comparability of the selected names, and registry coverage. Individual intentions and the mechanisms producing the contraction remain unresolved. The paper shows how administrative records of everyday choices can inform the study of political expression without requiring explicit declarations of political allegiance.

*Ph.D. candidate, Department of Political Science, University of Washington. This paper presents the author's independent methodological contribution as part of a doctoral dissertation. The underlying data collection and broader substantive project are joint work with Asli Cansunar, Can Bodrumlu, and Nela Mrchkovska; any errors reported here are the author's own.

1. Introduction

What can naming practices reveal about political expression under authoritarian rule? We examine this question through changes in Deniz naming across 636 Turkish districts from 1960 to 1990, across two political episodes: the growing political significance of the revolutionary student activist Deniz Gezmiş's name around his persecution and execution, and the period of repression surrounding the military coup of September 12, 1980. Sentenced to death in October 1971 and executed on May 6, 1972, Gezmiş became a martyr for the Turkish left, and the name *Deniz* ("sea" in Turkish) a symbol of political identification. Figure 1 previews the episode in a single series: the Deniz naming rate rises sharply in the early 1970s, remains elevated through the decade, and falls back around the 1980 coup without disappearing.

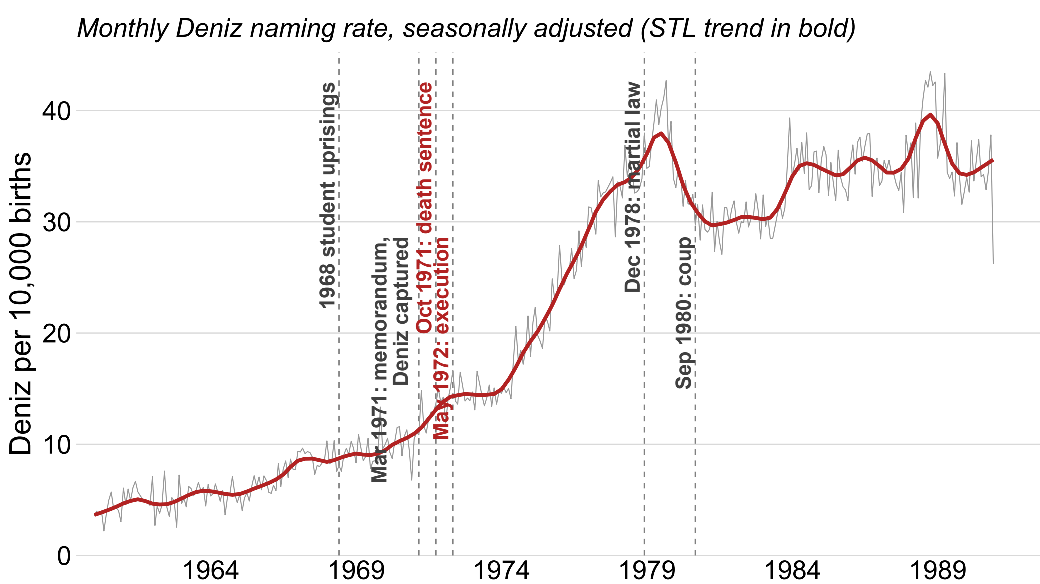

Figure 1: Monthly Deniz naming rate per 10,000 births, 1960–1990, STL-deseasonalized with trend in bold. Vertical lines mark political events.

How can we distinguish political expression from broader cultural influences on naming practices? Names can gain or lose popularity as cultural preferences change, so the temporal sequence alone does not establish a political interpretation.

Naming can reflect social identification while also changing with cultural preferences. Names have been used to study racial identification and immigrant assimilation (Abramitzky et al., 2020; Fryer & Levitt, 2004), but their popularity can also shift with cultural tastes (Lieberson, 2000). The associations conveyed by names can affect employers' responses (Bertrand & Mullainathan, 2004), and parents may adjust naming choices to perceived economic penalties or pressures to assimilate (Algan et al., 2022; Saavedra, 2021). Together, these findings suggest that the appeal of a name depends partly on what it conveys and how others may respond to it. We use *political signal* to describe the political association conveyed by a name.

Political events can give an established name a recognizable political meaning, allowing its adoption to express identification with a political cause. Repression may subsequently discourage that adoption by increasing the perceived costs of the association. We call this the *political signal hypothesis*.

The evidence comes from the Turkish civil birth register: more than 28 million domestic births assigned to geographic units from 1960 to 1990. The register also supplies its own comparison set. The state's language reform of 1934 to 1935 coined a cohort of secular Turkish names (Nişanyan, 2025) that share Deniz's era and its modernizing tailwind, and from that cohort we build a matched pool of non-revolutionary donors, frozen before any estimate is formed. The comparison names provide a realized benchmark for Deniz's trajectory. The panel models estimate changes relative to a projected path and examine how those changes vary across places. The annual analysis uses 1971 and 1980 to distinguish the political episodes; the monthly evidence limits attribution to individual events (Section 2.1).

The evidence supports a political interpretation of Deniz naming through the combination of its temporal trajectory, its unusually pronounced rise and subsequent contraction relative to matched names, and the political geography of its contraction. Together, these findings are consistent with naming as political expression and with repression constraining its adoption. The records leave individual motivations and the processes producing the contraction unresolved.

The paper makes two contributions. Substantively, it develops an account of political expression through a family decision whose political meaning can be recognizable without being exclusive. Work on hidden transcripts and preference falsification explains why conduct under pressure can diverge from privately held views (Kuran, 1997; Scott, 1990). Naming connects this problem to a choice that enters social life while retaining several possible meanings. By examining its adoption across political periods and places, the paper brings evidence from everyday decisions to the study of the possibilities and constraints of expression under authoritarian rule.

Methodologically, the paper combines projected trajectories with comparisons across names to study national political episodes without an unexposed geographic comparison. The panel models quantify departures from a specified continuation of the earlier trajectory and describe how responses vary across places. Applying the same specification to matched names assesses how unusual Deniz's response is within that cohort, while their observed trajectories provide a benchmark for naming counts. The combination assesses the size of the modeled change alongside the distinctiveness of Deniz's trajectory among comparable names. Their interpretation remains conditional on the projection rule and the comparability of the selected names.

Section 2 develops the historical context, the political signal argument, and its observable implications before describing the political geography of naming. Section 3 explains the comparison benchmarks, temporal assumptions, and estimands. Section 4 describes the registry and geographic coverage. Section 5 specifies the estimating strategies, and Section 6 reports the findings. Section 7 considers their scope and interpretation, followed by a short conclusion.

2. Political Context and Descriptive Patterns in Deniz Naming

Deniz is the Turkish word for *sea*, one of a set of secular given names that the language reform of 1934 to 1935 brought into use (Nişanyan, 2025). Its association with Deniz Gezmiş gave this established name a political significance that could develop through his public prominence, persecution, and execution. This history motivates the argument linking naming to political expression. The section develops that argument and its observable implications before examining the practice's geographic distribution.

2.1. Political events and the transformation of Deniz naming

Deniz Gezmiş rose to national prominence as a student leader in the radical wave of 1968 to 1971 and as a founder of the People's Liberation Army of Turkey (Türkiye Halk Kurtuluş Ordusu). The military memorandum of March 1971 brought down the government and opened a period of martial law. Gezmiş was captured that month and tried before a military tribunal, which announced death sentences for him, Yusuf Aslan, and Hüseyin İnan in October 1971. Their execution on May 6, 1972 made Gezmiş a martyr for the Turkish left. His prominence, capture, trial, sentence, and execution formed a sequence through which the name could acquire a recognizable political association, allowing its use for a newborn to express political identification.

The later contraction also unfolded across a sequence of political events. Martial law was introduced in thirteen provinces in December 1978 amid escalating political violence (Central Intelligence Agency, 1978). The coup of September 12, 1980 marked a further change in the political conditions under which the name was adopted. The junta banned political parties, closed trade unions, and arrested leftist activists on a mass scale (Sayari, 2010). It governed through a coercive legal regime that outlasted the return to civilian rule (Öztan & Bezci, 2015) and used penal provisions that restricted leftist political expression (Human Rights Watch, 1999). Naming remained legal within this restrictive political environment.

Martial law was lifted gradually and ended in 1987, but emergency powers and restrictions on political expression continued (Commission on Security and Cooperation in Europe, 1988, p. 5).

Figure 1 places these events against the national naming practice. The deseasonalized rate already rose from about 4 to 9 per ten thousand births during the 1960s, a secular climb that any counterfactual must accommodate (Section 3.3). Against that climb, the series accelerates through 1971, reaches a 1979 peak near 38, then declines around the coup, and begins to recover in the mid-1980s.

The monthly evidence does not locate the onset of this acceleration at the death sentence. A four-break Bai–Perron partition of the seasonally adjusted gap between Deniz and the matched cohort places the first break in March 1971, with a 95 percent confidence interval from November 1970 to July 1971 (Bai & Perron, 1998, 2003). The interval covers the military memorandum and Gezmiş's capture, but excludes both the October sentence and the May 1972 execution. The estimate cannot distinguish responses to the events within this sequence or establish when the name acquired political meaning.

In the same monthly partition, the later break falls in November 1979, with a 95 percent confidence interval from October to December 1979. The interval precedes the coup, placing the change within the earlier martial-law period. Earlier repression is one possible explanation, but the naming series does not establish it.

2.2. Observable implications of the political signal

A name can carry political meaning without serving only a political purpose. Families may choose the same name for different reasons, while its association with a public figure makes a political interpretation available to others. Naming can therefore express identification with a cause without constituting an explicit declaration of allegiance.

This ambiguity helps explain why repression may affect naming even when the name itself remains legal. Although chosen within a family, a name enters official records and everyday social interactions. If families expect its political association to attract unwanted attention, they may become less willing to adopt it. The political signal hypothesis thus concerns how political conditions shape the appeal and perceived costs of a name whose other meanings remain available.

We evaluate this argument through four observable implications, while distinguishing aggregate patterns from the mechanisms behind individual choices. First, national diffusion accelerates during the early 1970s relative to its projected prior path and remains elevated through 1979. Second, the period surrounding the 1980 coup brings a downward departure from that path, followed by a slow recovery. The annual model summarizes these changes across the 1971 and 1980 boundaries; the monthly evidence limits their attribution to individual events. Third, diffusion spreads beyond its initial locations rather than merely increasing within them. This remains a theoretical expectation: the geographic patterns and fitted propagation examined here do not directly establish expansion into new locations or transmission between them. The hypothesis does not require the post-1971 acceleration to increase monotonically with prior left support because a political name may diffuse beyond districts with strong left support. A deeper post-coup contraction in districts with stronger prior left support would provide related contextual evidence, without testing geographic spread itself. Fourth, Deniz should distinguish itself from the matched pool of comparison names in the conjunction of acceleration and reversal examined in Section 6.1.1.

2.3. The political geography of naming prevalence

Political context is measured by the Türkiye İşçi Partisi vote share in the 1969 general election, expressed in percentage points. This election was the last that TİP contested before its 1971 ban. The measure predates the first intervention by two years and cannot reflect subsequent naming. Electoral returns are observed for 634 historical districts; the two later boundary districts inherit the vote share of their parent district.

Figure 2 relates Deniz prevalence in 1971 to 1979 to this prior political geography. Clusters of high Deniz prevalence are disproportionately located above the median TİP vote share, and this association persists after adjustment for multiple testing. Section E reports the detailed classifications and geographic concentrations.

The party comparisons also limit the interpretation. Birlik Partisi produces nearly the same level concordance as TİP, so the map supports a broader political geography but does not isolate TİP as its mechanism. Section E reports the full comparison with the parties that contested the 1969 election.

Prevalence and response have different geographies. The local clusters of post-1971 acceleration do not survive the same multiple-testing correction, even though clusters of naming prevalence remain (Section E). The maps therefore describe where the practice was prevalent, while the panel analyses in Section 5 estimate within-unit change.

2.4. Naming trajectories by pre-intervention left support

Figure 3 groups districts by their 1969 TİP vote share, using quartiles and a separate band for the top five percent. The series are ordered from the mid-1960s onward, with Deniz rates in the groups with the highest support roughly two to three times those in the group with the lowest support. This persistent separation primarily describes differences in naming levels. The gaps appear to widen during the early 1970s, while the post-1980 contraction appears deepest in the groups with the highest support, which recover their pre-coup level only late in the decade. These visual contrasts do not establish that district responses vary systematically with TİP support. The province and district panel analyses in Section 5 estimate that moderation directly; the matched comparison across names does not. The bands are descriptive; estimation uses continuous vote share. The series are birth-weighted rates over the observed district totals in each group, rather than averages of district rates, so the district attrition described in Section 4.2 remains embedded in the display.

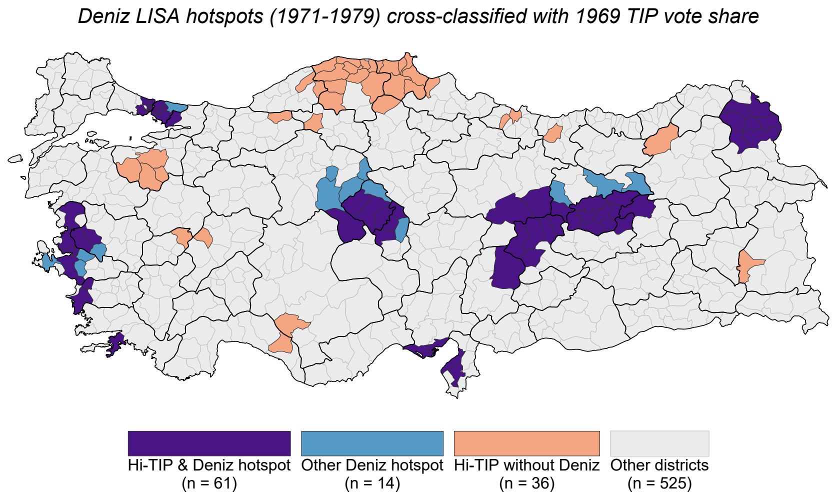

Figure 2: Deniz prevalence clusters, 1971–1979, by 1969 TİP support at the nominal five-percent threshold. Hi-TİP denotes vote shares at or above the district median.

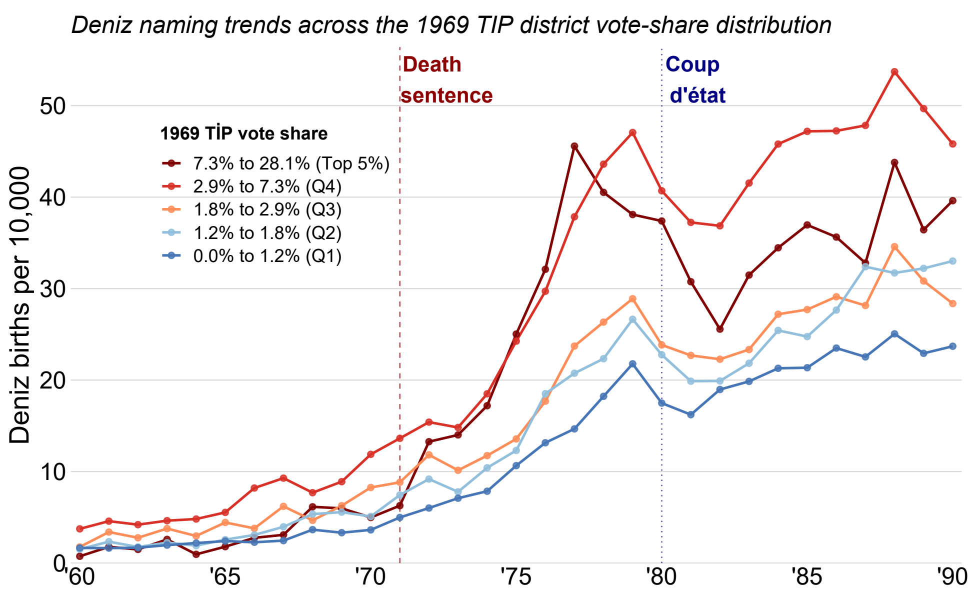

Figure 3: Birth-weighted Deniz rates per 10,000 by 1969 TİP vote-share quartile, with the top five percent separate. Vertical lines mark the sentence and coup.

These descriptions make the political interpretation plausible, but do not separate it from broader cultural influences on naming. The design therefore compares Deniz with other names and states the projected trajectory against which its changes are evaluated.

3. Research Design: Panel Interrupted Time Series and Matched-Name Comparisons

Interpreting the temporal and geographic patterns requires a comparison with the changes that might have occurred without Deniz's political association. The research design combines two complementary benchmarks in a setting with no geographically unexposed place: a panel interrupted time series (ITS) projects Deniz's pre-intervention path, while a matched-name comparison supplies a realized benchmark from the same registry and calendar. The panel analyzes annual naming trajectories at province and district scales.[1]

3.1. Counterfactual logic and comparison benchmarks

Both interventions reached the entire country, so no geographic control group remained untreated. The panel ITS instead compares the observed post-intervention trajectory with the path implied by the pre-intervention regime (McDowall et al., 2019, pp. 1–10). Its control condition is therefore an extrapolation: the canonical strategy carries the pre-existing trend "continued unchanged" into the post-intervention period (Lopez Bernal et al., 2017, p. 349). This is a modeling object rather than a group of observed units.

The register also supplies a realized comparison: names with similar origins and pre-intervention trajectories. Their subsequent naming patterns provide a benchmark for assessing whether Deniz's rise and contraction were unusual within that cohort. The two benchmarks answer related questions, but the observed comparison path and the model

[1] ITS draws on epidemiology (Lopez Bernal et al., 2017, p. 349), econometric work on dynamics (Hausman & Rapson, 2018, p. 542), and recent potential-outcome formalization (Menchetti et al., 2023, pp. 5–7).

projection rest on different assumptions.

Three threats organize the design: events concurrent with the interventions, changes in registry coverage or recorded population composition, and misspecification of the counterfactual trend. These threats are cataloged by Baicker and Svoronos (2019, pp. 16–17) and Hausman and Rapson (2018, pp. 547–548). Section 4.2 evaluates the measurement and coverage problem. The panel absorbs time-invariant geographic differences and estimates heterogeneous responses, but it does not remove nationwide concurrent shocks or determine the correct projection. The analysis therefore states the projection rule explicitly and complements it with the comparison across names developed in Section 3.2.

Figure 4 holds the observed trajectory fixed to show how the two designs construct different counterfactual paths. An interrupted series projects the treated series from its own past. Difference-in-differences uses the change in a contemporaneous untreated group, assuming that the treated group would otherwise have changed in parallel. This assumption concerns the two untreated paths, including shared changes that would occur without treatment. Such changes are reflected in the DiD counterfactual. Without that comparison, an interrupted series cannot separate an intervention response from a concurrent departure from the projected trend.

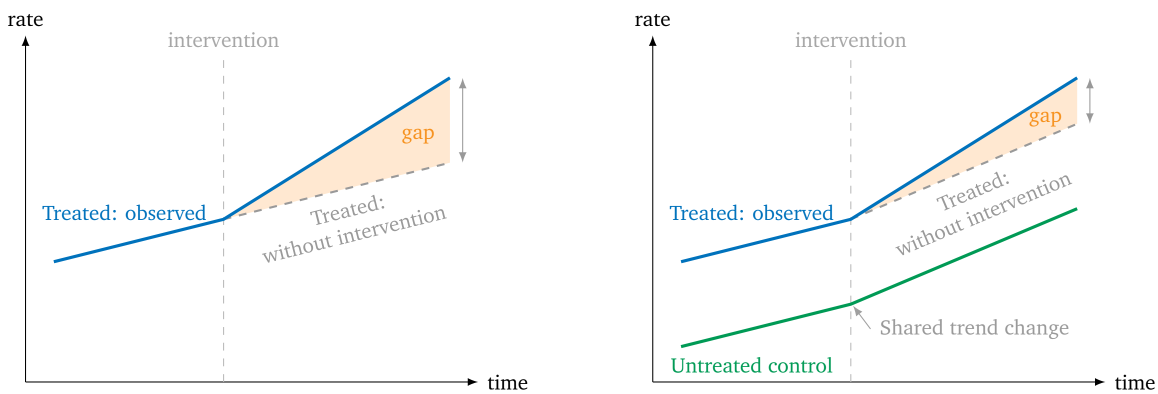

(a) Interrupted time series (no control group)

(b) Difference-in-differences (control group)

Figure 4: Counterfactual construction: (a) project the earlier trend; (b) use the change in the control group under parallel trends. Shading marks estimated gaps; the shared DiD trend change occurs without treatment.

3.2. Matched-name benchmarks and conditional exchangeability

A naming trend shared with other names remains a competing explanation for the temporal pattern. A modern Turkish name could enter use in the 1950s, spread through the 1960s, crest in the late 1970s, and decline as naming fashions changed around 1980. Fitted to the same design, that trajectory would produce a positive post-1971 diffusion loading and a negative post-1980 shift even without political activation. Unit effects, temporal persistence, spatial propagation, and other within-series adjustments do not exclude this possibility. The extrapolation problem in Section 3.3 therefore requires a comparison population that could share the secular naming trajectory without sharing Deniz's specific political association.

A controlled interrupted-series design ideally uses a comparison series closely aligned before the intervention (McDowall et al., 2019, p. 4). Deniz belongs to a cohort of names associated with the republican language reforms. The Turkish Language Association, founded in 1932, promoted Turkic-root vocabulary, while the 1934 Surname Law required citizens to register family names drawn from the new lexicon (Türköz, 2017). The Nişanyan etymological dictionary identifies the corresponding Yeni Türk given-name cohort, including *deniz*, the word for sea (Nişanyan, 2025). Candidate names from this cohort share reform-era provenance and exposure to registry-wide shocks. Matching then selects those with pre-1971 trajectories most similar to Deniz's. Their realized post-1971 path supplies an alternative benchmark without extrapolating Deniz's own pretrend.

The matched-name comparison serves two restricted purposes. First, within each permutation analysis, every name is evaluated under the same specification, placing Deniz within a pool-based null distribution. An extreme rank rejects the generic-trend null only conditional on within-pool exchangeability; it does not estimate causal magnitude. Second, the comparison names' realized trajectory supports the cohort-benchmarked excess-birth accounting reported in Section 6.1.2. This quantity describes Deniz relative to the specified cohort benchmark. It is not an identified effect of the political events.

Matching does not establish *exchangeability*. Names are outcomes observed in the same registry, not units assigned to conditions, and matching on pre-1971 trajectories cannot guarantee equal exposure to every political or cultural influence that shapes subsequent naming. Exact enumeration removes Monte Carlo error, but its inferential interpretation remains conditional on the political label being exchangeable within the realized pool. A rejection therefore does not generalize beyond that pool. The comparison rests on transparent construction and falsification rather than assignment. Section 5.3 defines eligibility, fixes the pool, evaluates balance, and implements the permutation procedure.

3.3. Intervention coding and counterfactual projection

The extrapolation begins where the baseline ends. The annual design uses 1971, the year of Gezmiş's capture, trial, and sentence, and 1980, the year of the coup, to distinguish the two political episodes. These historical boundaries do not identify the onset of naming responses. The monthly breaks precede both the sentence and the coup, and the interval around the early break extends into the pre-1971 baseline (Section 2.1). Coding the entire event year as treated cannot rule out earlier responses or separate events within that year. The design therefore evaluates the broader sequence predicted in Section 2.2: acceleration during the early 1970s followed by contraction around the coup. A rival account must reproduce both movements, but need not place either one at the precise date of the sentence or coup.

An ITS must specify both when an intervention begins and how its effect develops (Hausman & Rapson, 2018, p. 542). The impact taxonomy distinguishes abrupt from gradual onsets and permanent from temporary changes (McDowall et al., 2019, pp. 99–112). For each shock $k \in \{1971, 1980\}$, the design separates a step-level change from a bounded change that accumulates over its institutional window:

$$\underbrace{D_t^k = \mathbb{1}\{t \geq t_k^*\}}_{\text{level shift: abrupt and permanent}} \qquad \underbrace{S_t^k = \min\{\max(0,\ t - t_k^*),\ w_k\}\,/\,w_k}_{\text{saturating dose: gradual and permanent}}, \tag{1}$$

where t_k^* denotes the coded starting year and w_k the accumulation window. The step switches on and remains, whereas the dose rises in equal increments before remaining at one. The observable implications in Section 2.2 associate the first period primarily with diffusion and the period around the coup with contraction followed by recovery. A third bounded dose represents the pre-1971 secular increase. Section 4.3 gives the instantiated windows and variable definitions.[2]

The distinction between steps and doses matters because counterfactual projection is assumption-sensitive. A step cannot be reproduced by a smooth trend at the intervention date, so its contrast is less exposed to the pretrend rule. The dose coefficients and cumulative gaps are more exposed because both depend on how the untreated trajectory is projected through the same post-intervention years.

This sensitivity is clearest on the outcome's natural scale. The province analysis models log rates, so extending a linear log trend implies constant proportional growth and an exponentially increasing naming rate. The district Bayesian analysis instead models the log odds of naming (Section 5), so an extended linear trend compounds the odds rather than the probability itself. In both analyses, an open-ended secular component increasingly changes the post-1971 benchmark as the horizon lengthens.

Capping the pre-intervention dose holds its own contribution constant after 1970. It does not by itself hold the naming share at exactly its 1970 value because the fitted dynamics and other model components continue to operate. The cap therefore replaces an open-ended continuation of the secular component with a transparent held-contribution assumption. The time series does not establish that the secular increase had ended by 1970.

If secular momentum would have persisted without politicization, part of it can load onto S_t^{71} under the capped specification; under the uncapped specification, the continuing

[2] Capping also separates the terms. Under an uncapped segmented trend, the secular and diffusion terms correlate at 0.95 and the maximum variance-inflation factor is 33.6, against 0.64 and 6.87 once both terms are capped.

trend instead competes with that dose. The direction and magnitude of the natural-scale gap are therefore properties of the chosen benchmark, not of the observed counts alone.

3.4. Panel structure and spatio-temporal heterogeneity

The panel contributes cross-unit variation in responses rather than intervention timing: every province and district crosses the same two temporal thresholds. The province-panel analysis uses fixed effects to estimate average response dynamics, whereas the district Bayesian analysis partially pools heterogeneous local responses. This division follows the measurement considerations of Section 4.2, not differences in intervention exposure. Naming can spread through kinship, media, and political networks, so dependence across places is part of the phenomenon under study. The province-panel analysis represents temporal persistence and lagged spatial propagation, while the district Bayesian analysis estimates spatially structured heterogeneity in local responses. Neither restores a no-interference condition or creates unexposed places.

3.5. Dynamic, geographic, and comparative estimands

Four estimands describe the dynamic magnitude, cumulative count, geographic structure, and comparative distinctiveness of the naming response. They rely on different comparisons and therefore require distinct scales and population bases. For the province model, write $\ell_t^{(k)}(z_0)$ for the projected log naming rate with only episode k included and $\ell_t^{(0)}(z_0)$ for its no-shock counterpart, both at the common TİP reference $z_0 = 2.666$ percentage points. Dynamic responses and cumulative counts are

$$\begin{aligned}
\tau_k^{\text{dyn}}(h) &= \ell_{t_k^*+h}^{(k)}(z_0) - \ell_{t_k^*+h}^{(0)}(z_0), \\
E_T^M &= \sum_{t=1971}^{T} \left(Y_t - \hat{Y}_t^M\right), \\
E_T^a &= \sum_{t=1971}^{T} \sum_{i\in\mathcal{I}_{a,t}} n_{it}^a \left(\hat{p}_{it}^{a,\text{on}} - \hat{p}_{it}^{a,\text{off}}\right), \qquad a \in \{P, B\}.
\end{aligned} \tag{2}$$

Here t_k^* is the coded starting year for episode k. The log response at $h = 0$ is the *impact effect*, also called the point effect of Menchetti et al. (2023, p. 7). Subsequent horizons propagate the episode's step and dose through the fitted temporal and spatial dynamics (Li, 2017, eqs. 6.1–6.4). The reported percentage response is $100\{\exp[\tau_k^{\text{dyn}}(h)] - 1\}$. The combined path adds the two episodes' log responses at each year before exponentiation. These annual gaps are distinct from direct dose coefficients, steady-state contributions, and cumulative birth counts.

For the matched-name account, Y_t is the observed national domestic Deniz count and $\hat{Y}_t^M$ the count implied by the cohort's realized trajectory, including geographically unresolved domestic records. For each model account a, $\mathcal{I}_{a,t}$ indexes units with an observed denominator in year t, n_{it}^a is that birth denominator, and $\hat{p}_{it}^{a,\text{on}}$ and $\hat{p}_{it}^{a,\text{off}}$ are fitted naming rates with both episodes included and removed. These counts use each unit's covariates and response on the corresponding resolved population, rather than the common reference used for the dynamic summary.

The province no-shock path removes the four intervention terms, their TİP interactions, and their recursively propagated contributions while retaining the secular component, its TİP interaction, coastal trend, and observed exposure. The Bayesian path likewise removes the intervention contribution while retaining other fitted components and observed denominators. Both model accounts therefore condition on observed demographic composition and inherit their projection assumptions. The matched-name account requires no extrapolation of Deniz's pretrend, but its counterfactual interpretation depends on within-pool exchangeability (Section 3.2). None is a causal magnitude by construction, so the results report the three cumulative quantities separately.

The remaining estimands shift the question from magnitude to geography and comparative distinctiveness. Write Δ_i^{71} and Δ_i^{80} for district i's diffusion and coup responses, each measured as a deviation from the national trajectory, and consider a statistic computed identically for every name in the comparison pool of n names, Deniz included. They are

$$\underbrace{\rho_\Delta = \mathrm{Corr}_i\left(\Delta_i^{71},\, \Delta_i^{80}\right)}_{\text{geographic coupling of the responses}}, \qquad \underbrace{p_{\text{exact}} = (m+1)/n}_{\text{exact position in the pool}}, \tag{3}$$

where ρ_Δ measures the association between district deviations in diffusion and suppression. The political signal hypothesis predicts a negative correlation: districts with the steepest diffusion climb should experience the deepest post-coup drop. Heterogeneous responses that are unassociated across periods would instead imply a correlation near zero. Because ρ_Δ is a correlation among latent responses rather than an average response, it is estimated jointly with those responses at the district scale. Its primary evidence is the posterior magnitude, while its position in the matched-name pool is corroborating evidence.

The comparative estimand p_{exact} records Deniz's position in the frozen pool. Here m counts donor names whose statistic is at least as extreme as Deniz's under the specified tail rule. Component ranks use directional tails; joint-distance ranks use the upper tail. The fraction is exact with respect to enumeration, not random assignment: the analysis evaluates every eligible name position, so no Monte Carlo error enters the rank. The interpretation remains conditional on the within-pool exchangeability stated in Section 3.2.

The directional conjunction and comparative extremity answer different questions. The province intersection–union test evaluates whether the diffusion dose is positive and the coup shift negative. The primary matched-name comparison combines these coordinates in an unsigned multivariate distance, which measures joint extremity and may be driven by one coordinate. Component ranks retain the directional information, while the geographic coupling rank assesses the response association. A corresponding multivariate rank based on covariance-block correlations evaluates a separate claim about geographic structure and does not establish distinctiveness of that broader pattern (Section 6.1.1). Section 5.3 defines both statistics, their leave-one-out construction, their prespecified sides, and their numerical diagnostics. Section C maps the four estimands to the comparison carrying each, its empirical strategy, and its reporting location; Section C.2 records the corresponding assumptions.

4. Registry Data, Geographic Coverage, and Measurement

The scale division of Section 3.4 rests on what the registry can resolve. This section documents how individual civil registry records are cleaned, assigned where possible to constant-boundary birthplace units, and aggregated into province-year and district-year panels. It then establishes the coverage limits that govern scale assignment and defines the variables used in estimation.

4.1. Civil registry data and panel construction

The source is an extract of individual birth-registration records from the Turkish civil registry for 1960 to 1990. Each row describes one registered person and records a given name, sex, year of birth, and free-text birthplace. Given-name frequencies in administrative birth records are an established measure of naming behavior (Abramitzky et al., 2020; Fryer & Levitt, 2004). The free-text `birth_city` field is the only reliable source of birthplace geography. Separate registration fields identify the office holding the family's administrative civil record, an inherited association that can change later in life. They do not record place of birth or contemporaneous residence and are never used to assign a record to an analysis unit.

The cleaning pipeline normalizes name and geographic strings, recodes Deniz spelling variants and approved concatenated forms to a canonical indicator, classifies overseas birthplaces, and matches domestic birthplaces to the constant 1961 hierarchy of 67 provinces and 636 districts. It also removes 1,013,640 exact duplicate records, 3.22 percent of the raw extract. The resulting file contains 30,455,765 records: 29,995,824 domestic registrations and 459,941 registrations of births abroad. The latter remain available for auditing but do not enter the domestic panels. Section A documents string normalization, variant recovery, foreign-birth classification, historical boundary harmonization, and the correction files that make each assignment reproducible.

Figure 5 separates total annual Deniz registrations from registrations of births abroad. The latter remain part of the descriptive record but are excluded from the domestic province and district panels.

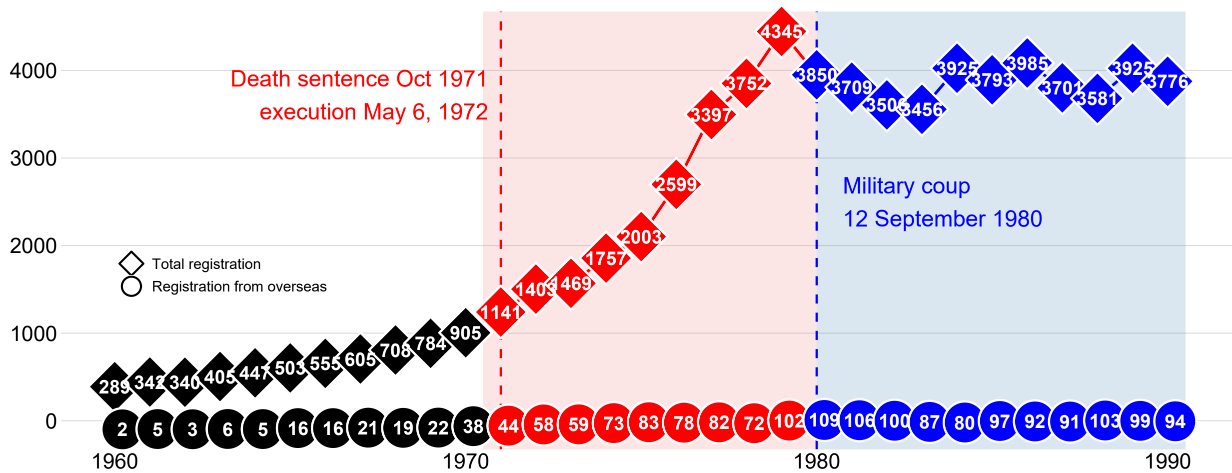

Figure 5: Annual Deniz registrations, 1960–1990: total births (diamonds) and births abroad (circles). Red and blue shading mark 1971–1979 and 1980–1990.

Panel construction applies two geographic admission rules to the domestic records. A record enters the province-year panel when its province of birth resolves and the district-year panel only when its district also resolves. The province panel contains all $67 \times 31 = 2{,}077$ cells and 28,394,417 assigned registrations. The complete district grid contains $636 \times 31 = 19{,}716$ cells. Of these, 19,434 contain at least one district-resolved domestic registration, jointly holding 20,061,493 registrations. The remaining 282 cells are held missing rather than coded as zero. Thus an observed district-year cell means that at least one birthplace resolved to that unit; it does not mean that every birth in the province received a district assignment. That distinction produces the coverage problem examined in Section 4.2.

4.2. Incomplete district identifiers and measurement consequences

The principal measurement loss occurs between province and district assignment. Of 67,094 domestic Deniz births, 65,732 resolve to a province and 38,340 to a district, giving analysis coverage of 98.0 and 57.1 percent. Among all domestic births, the corresponding shares are 94.7 and 66.9 percent (Table 8). Most loss therefore consists of province-resolved births without a district identifier, not absent district-year units. Only 282 of the 19,716 possible district-years contain no district-resolved domestic birth.

The loss reflects historical recording practice rather than residual matching error. Metropolitan offices often recorded only a province in the birthplace field, and Deniz births were concentrated in these urban records. When a district string is present, Deniz and other births resolve at essentially the same rate. Urban composition explains much, but not all, of the retention gap, and the latent districts of province-only records cannot be recovered by assigning them to *merkez* or another observed stratum (Section B).

Let R indicate that a district identifier resolves, N the given name, G^* the true district of birth, and t the year. Extending estimates from district-resolved births to all births requires name-neutral proportional thinning within each true district-year:

$$\Pr(R = 1 \mid N, G^*, t) = \Pr(R = 1 \mid G^*, t). \tag{4}$$

If this condition holds, the conditional distribution of names among resolved births matches that among all births within a true district-year. Resolution then reduces the effective denominator without changing the conditional naming probability. The condition cannot be point-identified because G^* is unobserved when $R = 0$. The administrative and matched-name diagnostics probe observable margins consistent with the condition; they neither verify it within latent districts nor recover that geography (Section B).

The empirical strategies follow this boundary (Section 5). Province M4 estimates average dynamics over records whose provinces resolve, retaining 98.0 percent of domestic

Deniz births and remaining insulated from missing district identifiers. District M4 and B2 estimate geographic heterogeneity among district-resolved births. Their likelihoods condition on observed district-year denominators and do not reconstruct unresolved births. Aggregate district quantities are more exposed to improving ascertainment than abrupt within-district contrasts, so province and district magnitudes are not treated as interchangeable. The matched-name sensitivities hold records and scale fixed in turn to assess whether differential resolution produces the comparative rank (Sections B and 5.3).

4.3. Outcome measurement, intervention coding, and covariates

The outcome is the annual count of newborns named Deniz in a geographic unit, observed against all births assigned to the same unit and year. The two panel scales require different representations of this pair. At province scale, only 102 of 2,077 cells record no Deniz birth and the median cell contains 10,770 births. The province analysis therefore uses the Jeffreys-smoothed log rate $\log \hat{p}_{jt}$, where $\hat{p}_{jt} = (y_{jt} + 0.5)/(n_{jt} + 1)$. At district scale, 10,051 of 19,434 cells are zero and the median cell contains 736 births. Smoothing would then determine much of the measured outcome in the sparsest cells: it changes a one-count cell by 50 percent. The district analysis instead retains y_{it} as a binomial count with n_{it} as the trial count. This representation preserves the information carried by zero counts and allows their precision to vary with the observed number of births (Sections 5.1 and 5.2).

Five deterministic functions of calendar time instantiate the design of Equation (1). The steps D_t^{71} and D_t^{80} switch on in 1971 and 1980. Three bounded doses rise from zero to one over fixed windows and remain at one thereafter: S_t^{pre} over 1960 to 1970, representing the secular increase described in Section 2; S_t^{71} over 1971 to 1979; and S_t^{80} over 1980 to 1987. The last endpoint uses the end of martial law as an institutional reference for recovery, while emergency powers and political restrictions continued (Commission on Security and Cooperation in Europe, 1988, p. 5). Neither the equal annual increments nor the endpoint measures the actual course of repression or its influence on naming. Because each dose is normalized to the unit interval, its coefficient is the loading associated with a full-unit change in the coded dose rather than a per-year slope; the dynamic propagation of that loading belongs to Section 5.1. The annual implementation, its timing implications, and nearby cap sensitivities are documented in Figure 1 and Sections C.2, 3.3 and D.2.

The political moderator tip is the unit's 1969 TİP vote share in natural percentage points. The fitted measure is centered at the births-weighted mean among units in the top quartile of support at each scale: 4.850 points across provinces and 5.712 across districts. A one-unit change therefore means one percentage point of TİP support. The district political-control battery uses the corresponding 1969 shares for BP, AP, CHP, and MHP. Because BP did not contest every district, non-contests are coded zero. Where BP enters as a focal moderator it carries a ballot-presence indicator and its regime interactions; the symmetric battery, which holds one form across parties, does not (Section E.4). Sections 5.1 and 5.2 state how the fitted interactions and average-unit reconstructions use these measures.

Exposure and secular geographic trends enter differently at the two scales. The province analysis uses the within-province component of lagged log population, a series combining intermittent censuses with electoral rolls and interpolated annually (Section A). The district analysis instead uses the directly observed registry birth denominator, decomposing its log into within-district, between-district-within-province, and between-province components (Mundlak, 1978). Both scales include standardized distance to the coast interacted with calendar year, allowing coastal naming trajectories to differ secularly without assigning that difference to either intervention.

Adjacency enters the analyses in two ways. The province dynamic panel uses a row-standardized queen-contiguity matrix to form the preceding year's mean log rate among neighboring provinces. The district Bayesian model instead uses the 636-node queen graph as the neighborhood structure for its BYM2 fields; it does not include a lagged neighbor outcome. The district graph restores the crossings at the Bosphorus and Dardanelles and retains the three island districts as singleton components. Thus the province term measures lagged spatial propagation, whereas the district graph governs spatial smoothing of heterogeneous responses.

Finally, the geographic identifier is assigned from the recorded birthplace string. It should not be read as later residence: the data cannot follow families who moved after the birth or directly measure transmission through migrant, kinship, or political networks. The panels therefore locate naming at recorded birthplace while leaving the social channel connecting places unobserved. Table 7, in Section A, collects the notation used below, and Section 5 specifies how each empirical strategy uses these measures.

5. Empirical Strategies and Estimation

Three complementary empirical strategies translate the design into distinct comparisons. The province fixed-effects analysis estimates average regime dynamics and their propagation. The district Bayesian analysis estimates geographically heterogeneous responses on the count scale. The matched-name analysis asks whether Deniz is unusual relative to comparable reform-era names. Their scale assignment follows the measurement boundary established in Section 4.2. A district log-linear companion estimates the propagation kernel and its political moderation and provides directional corroboration of the conjunction. Its level magnitudes are not compared with the province analysis and it supplies no counterfactual; it enters the matched-name comparison only as a scale sensitivity, never as a primary engine (Sections B and 5.4).

5.1. Province fixed-effects analysis of average dynamics

The province analysis estimates within-province changes in the naming rate. Province fixed effects absorb additive time-invariant differences across provinces, including persistent differences in baseline naming prevalence, so those differences cannot load directly onto the regime coefficients. They do not remove national events coinciding with the interventions, which remain an identifying limitation of the interrupted-series design. The specification is estimated as a four-rung nested sequence on one sample. M1 contains the interrupted-time-series basis of Equation (1); M2 adds demographic exposure and the coastal trend; M3 adds temporal persistence and lagged spatial propagation; and M4 adds TİP moderation:

$$\begin{aligned}
\log \hat{p}_{jt} = {} & \alpha_j + \beta_w \log N^{W}_{j,t-1} + \beta_c\,(\text{coast}_j \times t) \\
& + \underbrace{\beta_{\text{pre}} S^{\text{pre}}_t + \rho_1 \log \hat{p}_{j,t-1} + \rho_2 \log \hat{p}_{j,t-2} + \lambda\,[\mathbf{W} \log \hat{p}]_{j,t-1}}_{\text{secular trend and spatio-temporal dynamics}} \\
& + \underbrace{\gamma_1 D^{71}_t + \gamma_2 D^{80}_t + \delta_1 S^{71}_t + \delta_2 S^{80}_t}_{\text{two dated steps and two saturating doses}} \\
& + \underbrace{\theta^{\text{pre}}(\text{tip}_j \times S^{\text{pre}}_t) + \theta^{D71}(\text{tip}_j \times D^{71}_t) + \theta^{71}(\text{tip}_j \times S^{71}_t)}_{\text{TIP moderators, diffusion era}} \\
& + \underbrace{\theta^{D80}(\text{tip}_j \times D^{80}_t) + \theta^{80}(\text{tip}_j \times S^{80}_t)}_{\text{TIP moderators, coup era}} + \varepsilon_{jt}.
\end{aligned} \tag{5}$$

With the outcome and covariates defined in Section 4.3, the moderator tip_j is 1969 TİP vote share in natural percentage points, centered at 4.850 points, the births-weighted mean among the top quartile of provinces. Centering changes the reference point rather than the fitted model. The five regime coefficients as fitted describe that reference province. Quantities at the common cross-scale reference of 2.666 points, the mean 1969 district share retained for comparability across the two scales, are reconstructed by applying the corresponding linear combination to each coefficient and its TİP interaction, with covariance propagated through the same transformation. Because the regime basis is a function of calendar time, year fixed effects would absorb the intervention contrasts and are not included.

The dynamic block represents temporal persistence and local geographic dependence in naming rates. Two temporal lags are used because an AR(1) filter leaves Ljung–Box rejections in 7 of 67 detrended province series, while an AR(2) filter leaves 3, approximately the nominal rate. The spatial term uses neighbors' preceding-year rates rather than contemporaneous rates, reducing the simultaneity and reflection concern under temporal ordering. This timing is an identification choice rather than an empirical winner: the offset diagnostic retains 93.8 percent of the local association at one year but does not prefer the lagged term to its contemporaneous counterpart. The two temporal coefficients and λ together form the propagation kernel. Figures 20 and 21 report the full pre-estimation temporal and spatial batteries; Figure 23 shows the dependence remaining after M4 is fitted.

Inference uses Driscoll–Kraay standard errors (Driscoll & Kraay, 1998), which accommodate the remaining cross-sectional and serial dependence without imposing a particular residual graph. The reference distribution is a t with 28 degrees of freedom, one fewer than the 29 fitted years. Two-way clustering by province and year (Colin Cameron et al., 2011) remains in the covariance sensitivity battery, but it is not primary because its raw matrix is not positive definite and its numerical repair materially understates uncertainty for a sparsely identified moderator. Section D reports that diagnostic and the alternative covariance estimates.

In static M1 and M2, dose coefficients give contributions at full dose. In M3 and M4, they enter conditional on the lagged state. The steady-state multiplier $\mathcal{M} = (1 - \rho_1 - \rho_2 - \lambda)^{-1}$ gives the eventual response to sustained input. Finite-horizon responses $\tau_k^{\mathrm{dyn}}(h)$ instead propagate the annual event path through the fitted kernel. We recover these paths from 5,000 coefficient draws using the average-referenced estimates and Driscoll–Kraay covariance. Each gap compares that year's response with the no-shock path; it is neither a steady-state contribution nor a running sum across years.

The dynamic responses are also sensitive to bias in the fitted propagation kernel. With unit fixed effects, lagged outcomes expose the estimator to bias of order $1/T$ (Nickell, 1981). In a dynamic spatial panel, this can affect the spatial lag as well as the temporal terms (Lee & Yu, 2010). Because the long-run multiplier $\mathcal{M}$ is convex in their sum within the stable region, even modest coefficient changes can alter the propagated magnitude and its network share. Section D.6 reports a split-panel jackknife sensitivity that increases both quantities in this application. The short halves and spatial dependence limit its formal interpretation, so it does not establish a conservative bound on the dynamic responses in Section 6.2.

The primary directional claim is a conjunction: the diffusion dose is positive and the coup shift is negative at the common 2.666-point reference, $\delta_1 > 0$ and $\gamma_2 < 0$. An intersection–union test (R. L. Berger, 1982) rejects the complementary null only when both one-sided tests reject, making its p-value the larger of their two one-sided values. A joint Wald test would reject when either coordinate departs from zero and would ignore the prespecified directions, so it cannot test the conjunction. The test concerns the two average main effects only. It does not test the moderators or establish a mechanism, and it remains conditional on the regime coding, dynamic kernel, and covariance estimator. Province-scale TİP support has little spatial structure, so political moderation of propagation is evaluated only in the district companion; the province analysis retains a homogeneous spatial coefficient (Section 6.3).

5.2. District Bayesian analysis of geographic responses

The measurement argument of Section 4.3 assigns the district outcome to a count likelihood. For district i in province $j[i]$ and year t, the analysis models Deniz births as $y_{it} \sim \text{Binomial}(n_{it}, p_{it})$, with the observed births n_{it} entering as the trial count.

An aggregate trajectory cannot distinguish an even nationwide contraction from one concentrated where the name had diffused. The district analysis is designed to estimate that concentration: varying responses identify where the diffusion and coup departures lie, partial pooling stabilizes their estimation in sparse district series, and their joint covariance supplies the comparison. The Bayesian latent-Gaussian formulation regularizes the 636 series jointly and propagates uncertainty through a bounded correlation.

$$
\begin{aligned}
\operatorname{logit}(p_{it}) \;=\;& \underbrace{\alpha + \beta_{\text{pre}} S_t^{\text{pre}} + \gamma_1 D_t^{71} + \delta_1 S_t^{71} + \gamma_2 D_t^{80} + \delta_2 S_t^{80}}_{\text{national regime trajectory}} \\
&+ \underbrace{\beta_w \log n_{it}^{w} + \beta_{bw} \log n_{it}^{bw,p} + \beta_b \log n_{it}^{b,p} + \beta_c\,(\text{coast}_i \times t)}_{\text{exposure decomposition and coastal trend}} \\
&+ \underbrace{\psi\,\text{tip}_i + \boldsymbol{\theta}'(\text{tip}_i \times \text{regime}_t)}_{\text{TIP moderation}} + \underbrace{u_{j[i]} + v_i + s_i^{(0)}}_{\text{province and district intercepts}} \\
&+ \underbrace{s_i^{(71)} S_t^{71} + s_i^{(80\text{D})} D_t^{80} + s_i^{(80\text{S})} S_t^{80}}_{\text{covariance block: where the political signal lives}}
\end{aligned}
\tag{6}
$$

The population blocks use the variables defined in Section 4.3. Population regime coefficients are reconstructed at the common 2.666-point TİP reference within each posterior draw so that the reconstruction carries joint uncertainty. Cumulative counts instead use each district's covariates, fitted response, and observed denominator (Section 3.5). The intercept block combines an exchangeable province term $u_{j[i]}$ with a BYM2 district effect whose unstructured and spatial components are v_i and $s_i^{(0)}$. The final block lets the diffusion, coup, and recovery responses vary by district. Each deviation multiplies its bounded regime variable and describes the district's log-odds departure from the population trajectory when that variable reaches one.

The five specifications form a covariance-by-geography factorial, not a sequential ladder. B0 contains the fixed trajectory and the province and district intercepts, but no varying-response block. B1 adds independent exchangeable responses. B2 adds their unstructured covariance, B3 instead gives the responses independent BYM2 fields, and B4 combines coregionalization with BYM2 geography. All other model blocks, the estimation sample, and the integration strategy remain fixed (Table 1). B2 and B4 are co-primary for different purposes: B2 supplies direct correlation marginals from a trivariate distribution, while B4 represents the covariance structurally and supplies the response surfaces. Their agreement is robustness to alternative representations of the same geographic relationship, not independent replication.

Table 1: District Bayesian specifications by response covariance and geography.

Model	Response covariance	Response geography	Role
B0	none	none	Fixed trajectory and common intercept structure
B1	independent	none	Heterogeneous response benchmark
B2	unstructured trivariate	none	Primary direct response correlations
B3	independent	BYM2	Spatial heterogeneity without response covariance
B4	coregionalized	BYM2	Structural response loadings and spatial localization

The response block specifies which geographic relationships are estimated. A district slope on S^{pre} behaves as a second district intercept after saturation. Slopes on D^{71} and S^{71} compete because both variables remain constant from 1979 onward. The main block is (S^{71}, D^{80}, S^{80}). An alternative diagnostic replaces the varying recovery response with the 1971 step and uses empirical-Bayes estimation. Its allocation of covariance differs, but these joint changes prevent attributing the difference solely to the step. Section D.4 reports the estimates and their implications for distinguishing gradual diffusion from the earlier period's combined response.

At B4 the three responses are chained in the order of the historical periods,

$$s_i^{(71)} \sim \text{BYM2}(\tau_1, \phi_1), \qquad s_i^{(80\text{D})} = \kappa_{21}\, s_i^{(71)} + \xi_i^{(2)}, \qquad s_i^{(80\text{S})} = \kappa_{31}\, s_i^{(71)} + \kappa_{32}\, \xi_i^{(2)} + \xi_i^{(3)}, \quad (7)$$

a triangular linear model of coregionalization (MacNab, 2018). The coup response contains a component shared with the diffusion geography and an innovation of its own; recovery can load on both. Thus κ_{21} carries the diffusion-to-coup connection, while κ_{32} tests whether the coup innovation predicts recovery. With freely estimated loadings and innovation scales, the triangular form represents an unrestricted covariance at a single district. Its additional restriction is spatial: the same loadings connect the BYM2 innovations across the map. B2 imposes no spatial chain and supplies the direct covariance comparison. Table 2 states which question each quantity answers. Reverse-order and other structural checks are reported in Section E.

For exposition, the coregionalization implies the response-scale covariance summary

$$\boldsymbol{\Sigma} = \begin{pmatrix} \tau_1^{-1} & \kappa_{21}\tau_1^{-1} & \kappa_{31}\tau_1^{-1} \\ \cdot & \kappa_{21}^2\tau_1^{-1} + \tau_2^{-1} & \kappa_{21}\kappa_{31}\tau_1^{-1} + \kappa_{32}\tau_2^{-1} \\ \cdot & \cdot & \kappa_{31}^2\tau_1^{-1} + \kappa_{32}^2\tau_2^{-1} + \tau_3^{-1} \end{pmatrix}. \tag{8}$$

This is a coregionalization-scale summary, not the literal empirical covariance of the 636 realized B4 map values, whose BYM2 components may have different spatial shares. The scale-free quantity compared across specifications is ρ, not a loading or an unstandardized covariance.

Table 2: District response questions: B2 total correlations and B4 structural loadings.

Relationship	B2 comparison	B4 structural diagnostic
S^{71} and D^{80}	Primary $\rho(S^{71}, D^{80})$: did contraction concentrate where diffusion occurred?	κ_{21} carries the diffusion-to-coup structural link.
S^{71} and S^{80}	Secondary $\rho(S^{71}, S^{80})$: did diffusion geography predict recovery?	κ_{31} carries the diffusion-to-recovery structural link.
D^{80} and S^{80}	$\rho(D^{80}, S^{80})$ is the clean total no-momentum comparison.	κ_{32} is the direct coup-innovation diagnostic; inherited echo remains in total B4 covariance.

The priors are symmetric in sign and shrink toward simpler latent structures (Section D). Non-TİP population coefficients receive $N(0, 2.5^2)$ priors, while the six TİP terms receive $N(0, 0.10^2)$ priors per percentage point, a scale chosen by prior predictive checking.[3] B4's loadings receive independent $N(0, 1)$ priors. Its BYM2 marginal standard deviations use

[3] The conditional prior predictive calibration holds the non-TİP linear predictor at the logit of the median observed district naming rate and propagates the six TİP coefficients over the realized vote-share values and regime states. Figure 19 compares the prespecified scale with the coefficient-scale-equivariant alternative.

penalized-complexity priors with $\Pr(\sigma > 1) = 0.01$, and the mixing parameters use the prespecified PC prior with median 0.5. B2 places a Wishart prior with identity scale and five degrees of freedom on the three-response precision matrix. None of these priors favors a negative over a positive coupling. The induced-correlation audit and prior sensitivities are reported with the model checks.

Estimation uses integrated nested Laplace approximation (Martins et al., 2013). B2 and B4 use central-composite integration and store configurations for joint latent-field summaries. B2's correlations are direct hyperparameter marginals. B4's implied correlations are propagated from joint hyperparameter draws over the integration design rather than from posterior latent draws. Sensitivity fits estimated by empirical Bayes report an implied correlation as a posterior-mean point without an interval, because only the modal hyperparameter configuration is available. A correlation that is itself a direct hyperparameter retains its marginal interval under either integration strategy.

The district analysis estimates ρ_Δ, evaluates the posterior probability of the average-effect conjunction, and produces a model-projected cumulative gap conditional on its no-shock path. These quantities remain conditional on the likelihood, regime coding, latent structure, and priors; Bayesian estimation does not identify them by itself.

5.3. Matched-name permutation analysis of comparative distinctiveness

The matched-name analysis addresses the generic-trend null left open by the interrupted time series: a modern Turkish name could rise with secular Turkification and later decline without the political meaning attributed to Deniz. The funnel begins with 1,321 names observed in the national series. Restricting eligibility to the 1934 to 1935 language-reform cohort and applying a narrow screen for the martyr family and other leftist revolutionary referents leaves 132 candidates (Nişanyan, 2025). Politically meaningful names from other traditions remain eligible, including Ülkü. Mahalanobis nearest-neighbor matching without replacement selects 40 donors, yielding a final pool of 41 names. Forty donors were chosen for the attainable one-sided resolution, $1/41 = 0.0244$, not for statistical efficiency.

Matching uses four summaries calculated from province-assigned records over 1960 to 1969: mean log rate, linear trend, quadratic curvature, and log total births. The frozen-pool standardized mean differences are 0.192, 0.104, 0.135, and 0.194, respectively, all below 0.25 (Stuart, 2010), and donor ranges bracket Deniz on all four features (Table 3).

Figure 6 shows the broader pre-1971 name universe beside the selected pool.

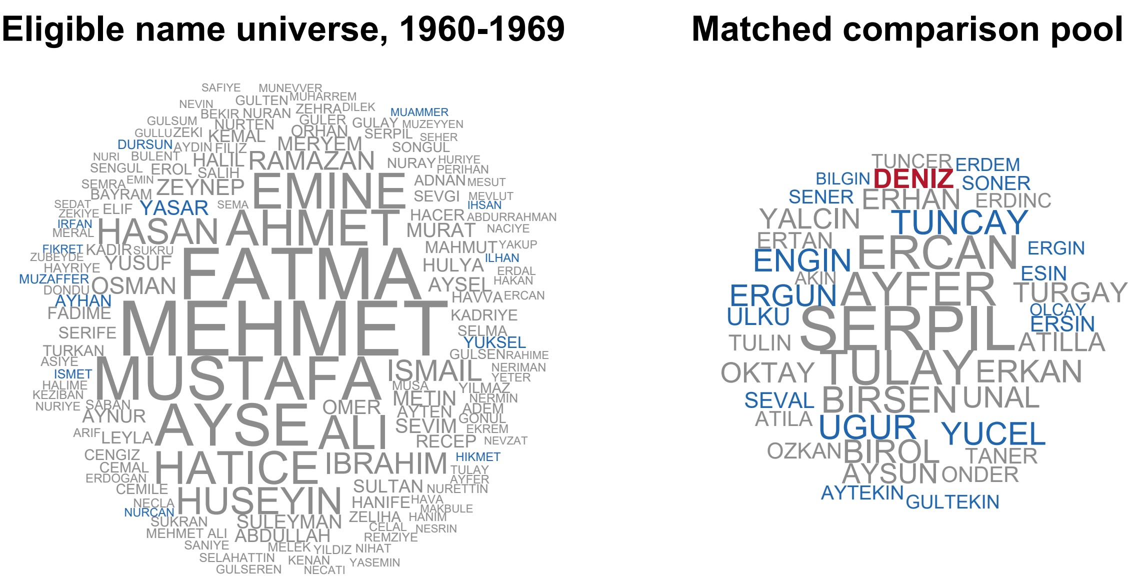

Figure 6: Eligible names (left) and matched pool (right). Size reflects province-assigned frequency, 1960–1969; Deniz is red, unisex names blue, and single-sex names gray.

Table 3: Deniz and matched-name balance, 1960–1969. Standardized differences use the eligible and full cohorts as reference distributions.

Covariate	Deniz	Donor mean	Eligible-cohort SMD	Full-cohort SMD
Pre-period level ($\bar{\ell}^{\text{pre}}$)	-7.369	-7.567	$+0.192$	$+0.167$
Pre-period slope ($\hat{s}^{\text{pre}}$)	$+0.087$	$+0.079$	$+0.104$	$+0.132$
Pre-period curvature ($\hat{c}^{\text{pre}}$)	$+0.004$	$+0.003$	$+0.135$	$+0.172$
Log total births ($\log B^{\text{pre}}$)	$+8.443$	$+8.245$	$+0.194$	$+0.168$

Sex composition is reported rather than matched and is bimodal in the donor pool. Restricting the comparison to unisex names leaves Deniz and 17 donors, so the one-sided floor is $1/18 = 0.056$. A separate registry audit places Deniz's metropolitan share at 0.534 against a donor mean of 0.311, near the top of the pool, but the matching inputs contain neither metropolitan nor *merkez* donor-range statistics. Province M4 avoids district assignment. For district comparisons, the common-record and nonmetropolitan sensitivities in Section B assess whether record selection or geographic aggregation explains the rank difference. The focal B2 metropolitan-exclusion sensitivity remains in Section 6.3. Pool checks appear in Section F.

Each name is fitted under province M4 and district B2, changing only the focal outcome and rebuilding outcome-derived lags. B2 exposes coupling directly without B4's spatial-slope block; the binomial is retained because beta-binomial dispersion is not separately identified from latent heterogeneity. For each engine, the primary joint statistic is a leave-one-out Mahalanobis distance on (S^{71}, D^{80}), with the reference mean and covariance recomputed from the other 40 names. The distance is unsigned and omnibus: it may be driven by one coordinate, whereas prespecified directions apply to component ranks. Reciprocal-condition-number diagnostics accompany the joint ranks. Exhaustive evaluation of all 41 focal-name positions gives $p_{\text{exact}} = (m+1)/41$ without Monte Carlo error, under the exchangeability conditionality stated in Section 3.2 (Imbens & Rubin, 2015; Pesarin & Salmaso, 2010).

The comparison assesses comparative distinctiveness, not a causal magnitude. All 41 fitted couplings are negative under the adjacent diffusion and coup windows; Deniz is distinguished by its rank, not the sign alone. Component and coupling ranks corroborate the prespecified omnibus distance but do not turn it into a conjunction test (Section 3.5). A descriptive name-selection-odds bound is reported only when the unbiased rank is below the conventional threshold; it is not a Rosenbaum matched-pair bound because the design contains no pairs (cf. Rosenbaum, 2020).

5.4. Complementarity and limits of comparison

Table 12 maps estimands to empirical strategies. Province and district analyses project a no-shock path from fitted pre-intervention trajectories; matching uses a realized cohort benchmark. The former depend on extrapolation and model specification, the latter on within-pool exchangeability and observed-feature balance. Shared regime coding means agreement is not independent validation, province and district magnitudes are not directly comparable, and cohort-benchmarked cumulative excess is not a causal effect (Section D.2). Section 6 begins with the comparison across names, then examines the modeled timing and geographic variation.

6. Results

Deniz's naming rate rose substantially more than the matched-name benchmark during the 1970s. It subsequently fell around the 1980 coup, but remained above that benchmark through 1990. We begin with this comparative trajectory, which establishes the scale of Deniz's divergence from other names. The panel models then characterize the changes relative to their projected path and examine how the responses vary across places.

6.1. Deniz relative to comparable names

Deniz's divergence from the matched-name benchmark grew through the 1970s (Figure 7). In 1971, its naming rate exceeded the benchmark by 2.59 births per 10,000, or 27.4 percent. By 1979, the difference had reached 26.24 births per 10,000, or 219.7 percent. The comparison places the expansion of Deniz naming alongside the trajectories of forty names with similar origins and pre-period characteristics.

To construct the benchmark, each comparison name's annual share is divided by its pooled 1960 to 1969 share. The equal-weight mean of these 40 relative paths is then rescaled to Deniz's pooled pre-period share. This preserves the pre-period level without forcing equality in 1970.

Pointwise finite-pool comparisons place Deniz in the upper two positions from 1974 through 1980 and again in 1988. No annual gap exceeds the simultaneous five-percent reference under either a 1971 or 1972 start. After accounting for comparisons across years, the annual analysis therefore does not identify an onset year.

6.1.1 Comparative distinctiveness in the matched-name pool

Deniz also stands out when the modeled growth and coup response are considered jointly. It ranks first among the 41 names in both the province and district analyses, with a conditional pool probability of $p = 1/41 = 0.0244$ in each (Figure 8 and Table 4). Several comparison names share the positive-growth and negative-coup signs. Deniz's distinction therefore concerns its position in the joint distribution of the two responses. The inferential interpretation of this rank remains conditional on exchangeability within the selected pool.

The 82 primary name-specific fits preserve the prespecified vote-share moderator and covariance definitions while replacing only the focal outcome. Deniz's leave-one-out Mahalanobis distance on (S^{71}, D^{80}) is 15.92 in province M4 and 12.53 in district B2.

Table 4: Matched-name ranks and signed components under M4 and B2. Conditional finite-pool probabilities enumerate all 41 focal-name positions.

Analysis	Quantity	Deniz estimate	Rank / 41	Conditional p
Panel A. Primary joint distance				
M4	$D^2(S^{71}, D^{80})$	15.921	1 / 41	0.024
B2	$D^2(S^{71}, D^{80})$	12.528	1 / 41	0.024
Panel B. Signed components				
M4	S^{71}	+0.658	2 / 41	0.049
M4	D^{80}	-0.289	2 / 41	0.049
M4	$S^{71} - D^{80}$	+0.947	1 / 41	0.024
B2	S^{71}	+1.450	1 / 41	0.024
B2	D^{80}	-0.264	4 / 41	0.098
B2	$S^{71} - D^{80}$	+1.714	1 / 41	0.024

The signed components show how Deniz's position varies across dimensions. Under

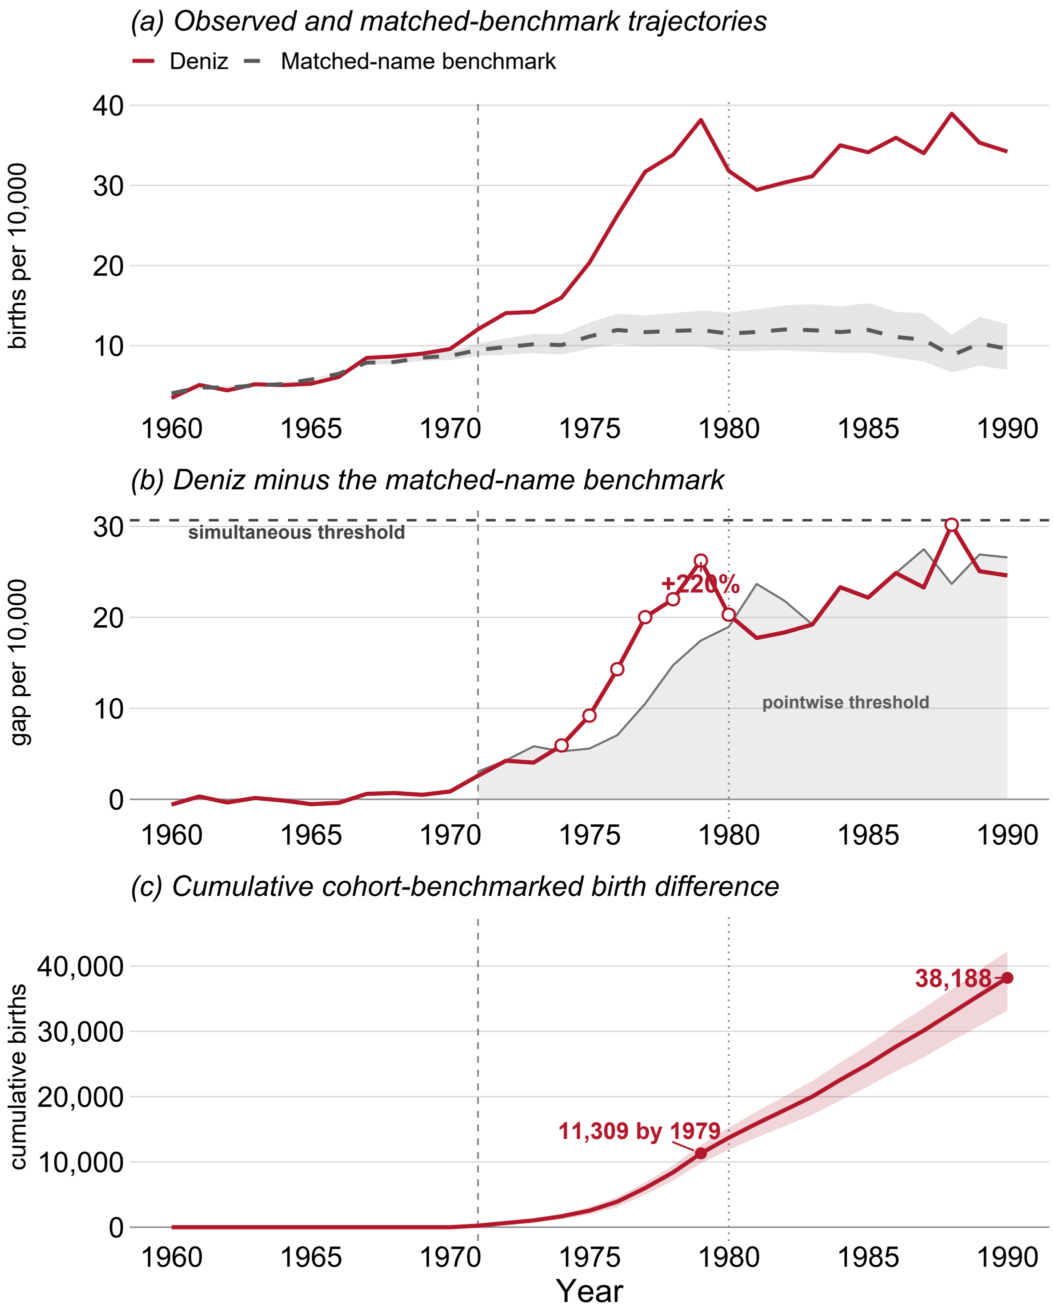

Figure 7: Matched-name benchmark: (a) Deniz and comparison rates per 10,000; (b) annual gaps with pointwise five-percent references; (c) cumulative birth differences. Ribbons show 95 percent donor-composition bands. Circles mark conditional $p \leq 0.05$; the dashed threshold is simultaneous and one-sided.

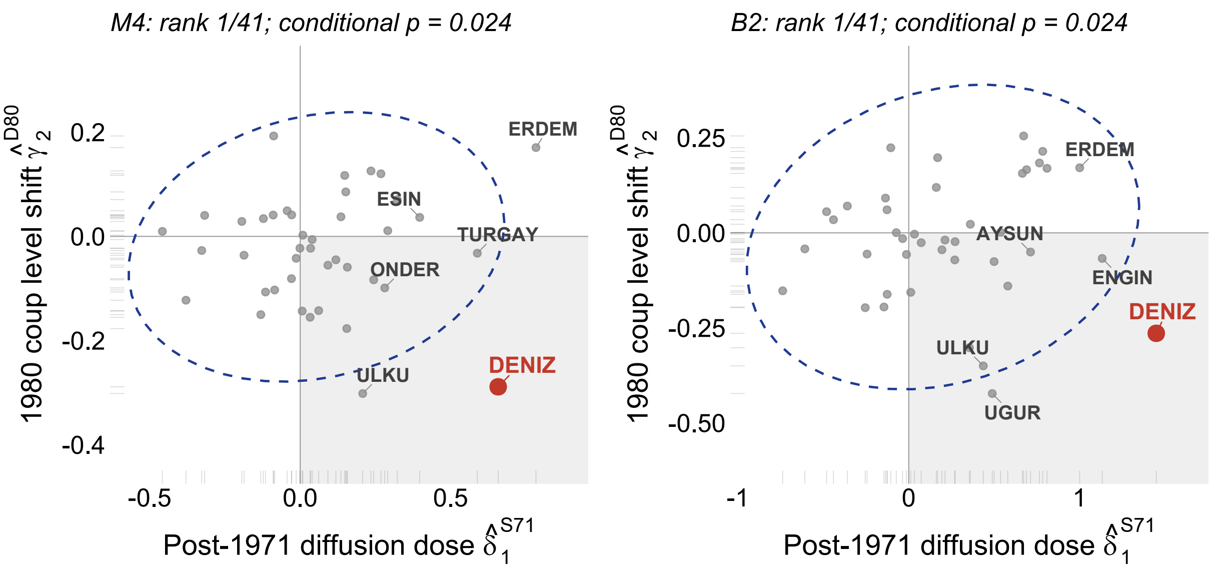

Figure 8: Deniz and 40 matched names on the (S^{71}, D^{80}) plane under province M4 and district B2. Ellipses describe donor distributions; inference uses leave-one-out Mahalanobis scores.

M4, Deniz ranks second on positive diffusion and second on the negative coup coordinate; under B2, it ranks first and fourth, respectively. The signed contrast $S^{71} - D^{80}$ ranks first under both. The complete rank inventory and pre-period diagnostic appear in Table 20, Figure 27, and Section F.

The broader geographic pattern is less distinctive. The full three-correlation comparison gives Deniz a joint distance of 5.92, rank 4 of 41, with $p = 0.098$. The distinction in the joint (S^{71}, D^{80}) comparison does not establish a distinctive three-correlation pattern.

All 41 fitted correlations between diffusion and the coup response are negative in this pool (Section 5.3), so the sign alone does not distinguish Deniz. Its correlation, -0.522, is the most negative value in the pool and narrowly beyond Ergun. The permutation inventory reads each coupling directly off the fitted marginal, so this value differs in the third decimal from the draw-based summary of Table 6.

Deniz's first-place point rank persists across five comparisons: province M4 on all province-resolved records, province M4 restricted to district-resolved records, district M4 on those same records, district M4 outside the six metropolitan provinces, and district

B2. Across 5,000 approximate coefficient-uncertainty recalculations, the corresponding rank-one frequencies are 42.1, 62.6, 89.1, 83.8, and 95.4 percent. These are simulated rank-stability frequencies, not posterior model probabilities.

The differences in stability do not isolate the contribution of finer geography. Restricting province M4 to district-resolved records raises the rank-one frequency by 20.5 percentage points. Fitting district M4 to those same records raises it by a further 26.5 points, but also changes the fixed effects, covariate construction, and leverage (Table 10, Figure 15, and Section B). The district comparisons retain the metropolitan attrition described in Section 4.2.

6.1.2 Cumulative differences on the birth scale

The annual differences accumulate to 38,188 Deniz births above the matched-name benchmark between 1971 and 1990. This cumulative difference continues to increase after the coup because Deniz remains above the benchmark despite the initial contraction in its naming rate. After 1985, the widening gap also reflects a decline among the comparison names. The total describes naming relative to the selected benchmark; it does not count politically motivated parents or births caused by the political events.

The cumulative difference is 11,309 by the end of the diffusion window in 1979, with a central 95 percent donor-resampling composition interval from 9,796 to 12,609. By 1990, the corresponding composition interval is 33,187 to 42,275, and the leave-one-donor range is 37,728 to 39,045. These ranges describe sensitivity to the selected comparison names, not confidence intervals for a causal effect.

The model-projected paths answer a different question and retain their distinct coverage bases in Figure 25 and Section E; Section E.9 reports their cumulative totals beside this cohort-benchmarked count.

6.2. Timing and persistence of the modeled changes

The province models recover the same broad sequence under a different benchmark. Relative to the model-projected path, Deniz naming expands during the 1970s and contracts around the coded 1980 boundary. At the common 2.666-point TİP reference, the diffusion coefficient is positive, $\hat{\delta}_1 = 0.658$ log points $[0.401, 0.914]$, while the 1980 level shift is negative, $\hat{\gamma}_2 = -0.289$ $[-0.373, -0.204]$. These estimates support the two directional predictions within the specified model. They do not locate the onset of the naming response at the sentence or coup itself.

The fitted trajectory shows how these contributions develop through time. By 1979, the complete response is approximately 311 percent above the projected path. It contracts after the 1980 boundary, reaching a post-coup trough of approximately 234 percent above that path in 1982, before rising to 343 percent by 1990. The contraction therefore reduces the earlier divergence without eliminating it. These percentages describe the modeled gap at each year; their magnitude depends on the counterfactual projection and the fitted temporal and spatial dynamics.

Figure 9 also separates the own-province component from the additional response generated by the fitted spatial structure. In 1979, the own-province component is approximately 192 percent above the projected path, while network amplification is approximately 40 percent. The amplification reaches 48 percent in 1980 and 55 percent by 1990. These are multiplicative contributions, so their percentages should not be added. Because all provinces experience the national events, the decomposition describes the model's propagation of a common input without identifying transmission between provinces.

The diffusion coefficient describes the contribution of the fully accumulated dose before dynamic propagation. Its Driscoll–Kraay standard error is 0.125 (one-sided $p = 6.9 \times 10^{-6}$), while the coup coefficient's standard error is 0.041 (one-sided $p = 6.3 \times 10^{-8}$). The intersection–union test requires both signs and takes the larger one-sided value, $p = 6.9 \times 10^{-6}$. The district log-linear companion reproduces both signs (Section 5).

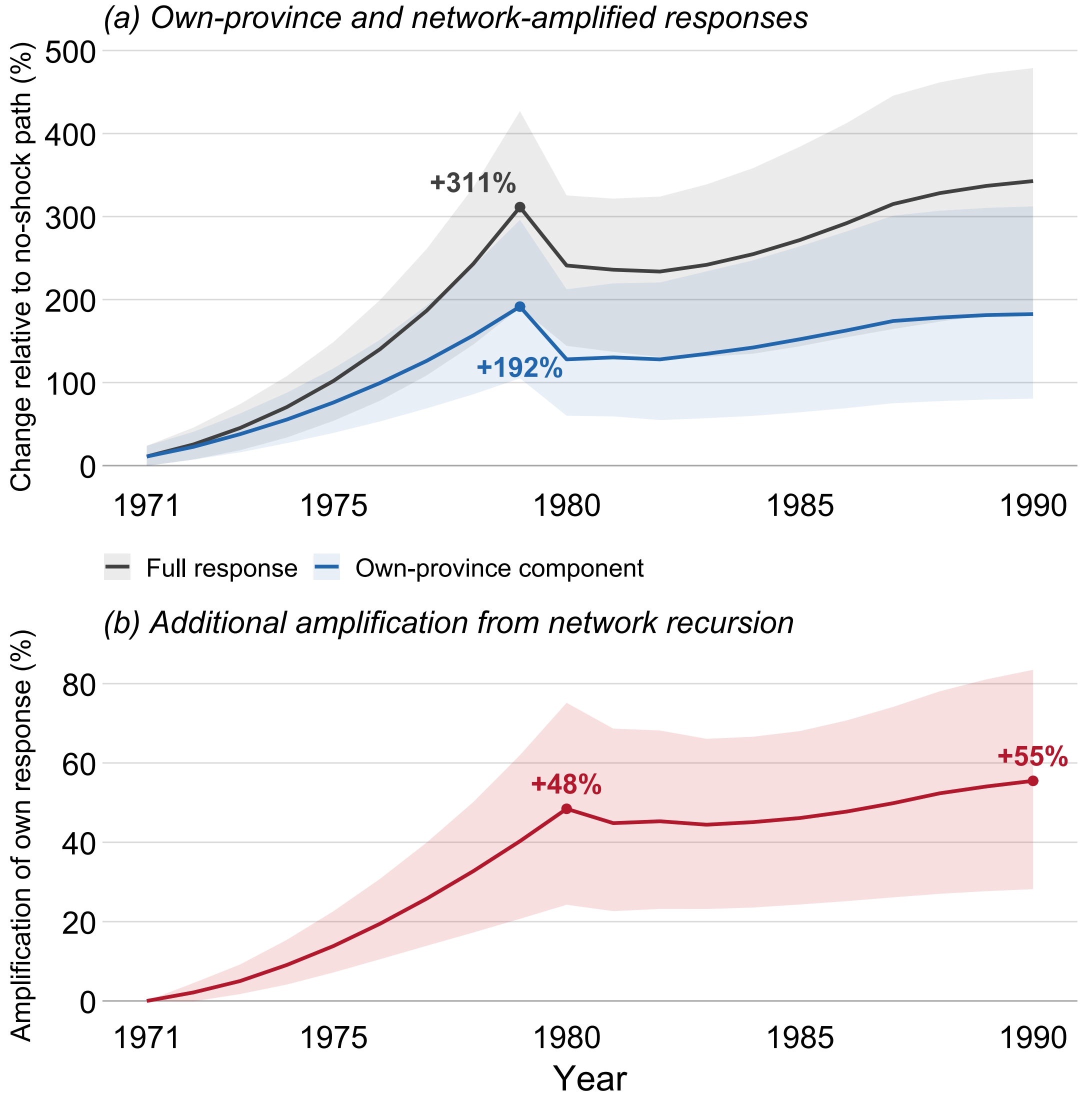

Figure 9: Province M4 responses: (a) own-province and complete percentage changes from the no-shock path; (b) network amplification. Medians and 95 percent bands from 5,000 coefficient draws at the 2.666-point reference.

The specification comparison uses a steady-state scale (Table 5). Multiplying the dynamic diffusion coefficient by 2.242 gives an implied contribution of 1.475 after adjustment to a sustained full dose, rather than the response reached by 1979. The corresponding contribution is 1.393 in M1, 1.498 after adding exposure and coastal controls in M2, and 1.477 in M3. The similar magnitudes show that the estimated expansion persists across these specifications. M1 and M2 have no dynamic propagation, so their dose coefficients already describe the full-dose contribution. M4's coefficient rows use the 4.850-point top-quartile reference; the implied steady-state S^{71} row uses the common 2.666-point reference.

Table 5: Province log-linear specifications, 1962–1990. Driscoll–Kraay standard errors in parentheses.

	M1 ITS basis	M2 + controls	M3 + dynamics	M4 + TİP
Panel A. Coefficients				
S^{pre} (β_{pre})	1.1550*** (0.0471)	1.2990*** (0.0627)	0.6738*** (0.1358)	0.7730*** (0.1151)
D^{71} (γ_1)	0.1634*** (0.0401)	0.1786*** (0.0410)	0.1003* (0.0578)	0.1331** (0.0548)
S^{71} (δ_1)	1.3928*** (0.0434)	1.4984*** (0.0476)	0.6069*** (0.1280)	0.7738*** (0.1641)
D^{80} (γ_2)	−0.3034*** (0.0578)	−0.2948*** (0.0586)	−0.2864*** (0.0408)	−0.3893*** (0.0734)
S^{80} (δ_2)	0.2463*** (0.0473)	0.3136*** (0.0466)	0.2000*** (0.0435)	0.1789*** (0.0352)
$\log \hat{p}_{j,t-1}$ (ρ_1)	—	—	0.1870*** (0.0336)	0.1726*** (0.0329)
$\log \hat{p}_{j,t-2}$ (ρ_2)	—	—	0.1762*** (0.0280)	0.1636*** (0.0292)
$[\mathbf{W} \log \hat{p}]_{j,t-1}$ (λ)	—	—	0.2260*** (0.0571)	0.2178*** (0.0547)
$\log N^{W}_{j,t-1}$ (β_w)	—	−0.6727*** (0.1910)	−0.3319** (0.1593)	−0.3267** (0.1472)
tip $\times$ S^{pre}	—	—	—	0.0304 (0.0210)
tip $\times$ D^{71}	—	—	—	0.0130 (0.0204)
tip $\times$ S^{71}	—	—	—	0.0531* (0.0287)
tip $\times$ D^{80}	—	—	—	−0.0461** (0.0204)
tip $\times$ S^{80}	—	—	—	−0.0112 (0.0130)
Panel B. Fit and steady-state summaries				
Within R^2	0.7748	0.7801	0.8084	0.8113
Steady-state multiplier $\mathcal{M}$	1.000	1.000	2.434	2.242
Implied steady-state S^{71}	1.3928	1.4984	1.4769	1.4752
Province-years	1,943	1,943	1,943	1,943

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p \leq 0.01$.

The recovery coefficient partly offsets the negative coup coefficient. At the common 2.666-point reference, the seven-year recovery dose contributes $+0.203$, giving a coefficient sum of -0.085 $[-0.190, +0.020]$ $(0.051,\ p = 0.107)$. At the 4.850-point top-quartile reference, the same fit gives a coup coefficient of -0.389, a recovery coefficient of $+0.179$, and a sum of -0.210 $[-0.342, -0.079]$ (0.064, $p = 0.003$). The more negative sum at higher TİP support reflects both a larger initial contraction and a smaller recovery coefficient. This contrast belongs to the equal-province-year estimand and attenuates under birth-weighted sensitivities (Table 17 gives the full reference grid).

These sums combine direct contributions before dynamic propagation; they do not measure the remaining naming-rate gap at a particular year. On the same coefficient scale, the 1971 block sums to 0.763 $[0.465, 1.060]$, and the two interruption blocks sum to 0.677 $[0.324, 1.031]$. The response at each year instead incorporates the history of both blocks through the fitted dynamics, as shown in Figure 9.

The coded 1971 step estimates a level shift alongside the gradual diffusion term. It is $+0.163$ $[0.081, 0.246]$ ($p = 0.0003$) in M1 and $+0.105$ $[-0.011, 0.220]$ (0.056, $p = 0.074$) in M4 at the 2.666-point reference. The change occurs when the lag structure enters, not when political moderation is added. This annual step cannot distinguish responses before and after the October sentence, or test whether naming had already responded before 1971. The earlier-response assumption remains unresolved (A2 in Section C.2). Whether suppression was concentrated where diffusion had taken hold is a geographic question.

6.3. Geographic variation and political context

Districts with larger modeled increases during the 1970s tended to experience larger contractions around the coup. Across 636 districts, B2 estimates the response correlation $\rho_\Delta = \rho(S^{71}, D^{80}) = -0.525$ $[-0.667, -0.360]$. B4 gives a closely aligned estimate of -0.524 $[-0.743, -0.263]$. The agreement supports the geographic association across two representations of district responses, although both rely on the same observations and temporal coding (Tables 1 and 6). The summaries use 2,000 direct marginal draws for B2 and 5,000 joint hyperparameter draws for B4.

The allocation of this association within the earlier period is more sensitive to the model. In an existing diagnostic that allows the 1971 step to vary across districts, the correlation between gradual diffusion and the coup response weakens, while the step and combined earlier contribution remain negatively associated with it. This diagnostic also changes the treatment of recovery heterogeneity and uses empirical-Bayes estimation, so it does not isolate the consequence of allowing the step to vary. We therefore distinguish the association with the earlier increase from its attribution specifically to gradual diffusion (Section D.4).

B4 describes how the response fields share geographic variation. Its diffusion-to-coup loading is $\kappa_{21} = -0.280$ $[-0.428, -0.132]$. The shared component accounts for 28.9 percent $[6.9, 55.2]$ of the modeled coup-response variance under the coregionalization parameterization. This is a variance decomposition, not a causal contribution. Recovery also loads on the diffusion field, $\kappa_{31} = -0.256$ $[-0.462, -0.058]$, while the loading on the coup-specific innovation is uncertain, $\kappa_{32} = +0.186$ $[-0.608, +1.057]$. Figure 11 maps the fields on the birth scale. Individually resolved responses are sparse: 61 districts have diffusion-field intervals excluding zero and 47 have such intervals for the coup-specific innovation. The geographic association is therefore assessed through the joint distribution of responses.

Table 6: Posterior response correlations and structural loadings. Asterisks mark 95 percent credible intervals excluding zero; $P(< 0)$ denotes posterior mass below zero.

Quantity	Fit	Posterior mean	95% CrI	$P(< 0)$
$\rho(S^{71}, D^{80})$	B2	-0.525^{*}	$[-0.667, -0.360]$	> 0.999
$\rho(S^{71}, D^{80})$	B4	-0.524^{*}	$[-0.743, -0.263]$	> 0.999
$\rho(S^{71}, S^{80})$	B2	-0.295^{*}	$[-0.490, -0.085]$	0.999
$\rho(D^{80}, S^{80})$	B2	$+0.008$	$[-0.243, +0.263]$	0.482
κ_{21}, diffusion to coup	B4	-0.280^{*}	$[-0.428, -0.132]$	> 0.999
κ_{31}, diffusion to recovery	B4	-0.256^{*}	$[-0.462, -0.058]$	0.994
κ_{32}, coup innovation to recovery	B4	$+0.186$	$[-0.608, +1.057]$	0.341
Echo share of the coup-response field	B4	0.289	$[0.069, 0.552]$	—

Alternative Wishart priors, metropolitan exclusion (-0.448 $[-0.619, -0.247]$), and reverse structural order (-0.507) retain the negative correlation. Section E distinguishes intervals from plug-in checks and reports likelihood, residual, and pre-period qualifications. Figure 10 shows the response correlations.

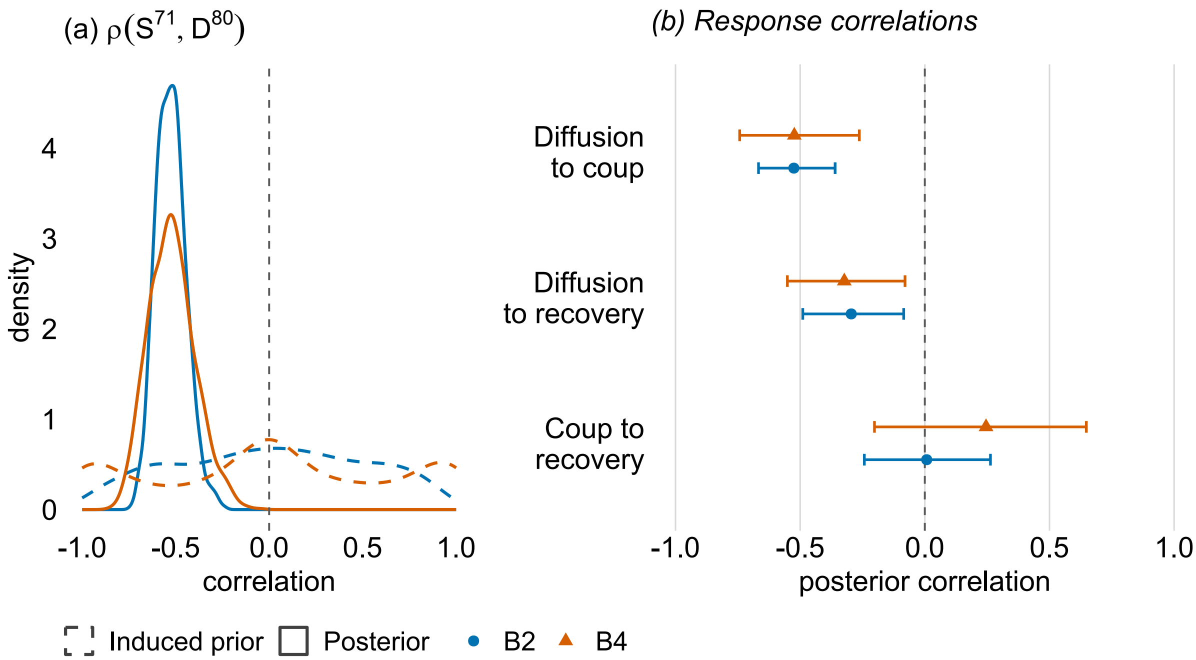

Figure 10: Geographic coupling: (a) primary correlation posteriors with induced priors dashed; (b) all three response correlations. B2 estimates correlations directly; B4 implies them through coregionalization.

Districts with larger modeled diffusion responses also tend to have smaller recovery responses. Under B2, $\rho(S^{71}, S^{80}) = -0.295\ [-0.490, -0.085]$. The correlation between coup and recovery responses is $\rho(D^{80}, S^{80}) = +0.008\ [-0.243, +0.263]$, providing little evidence of an association between those responses. Its uncertainty does not establish that recovery was unrelated to the initial contraction. This total correlation and B4's innovation loading answer different questions: the former concerns the entire coup response, while the latter concerns its component not shared with diffusion.

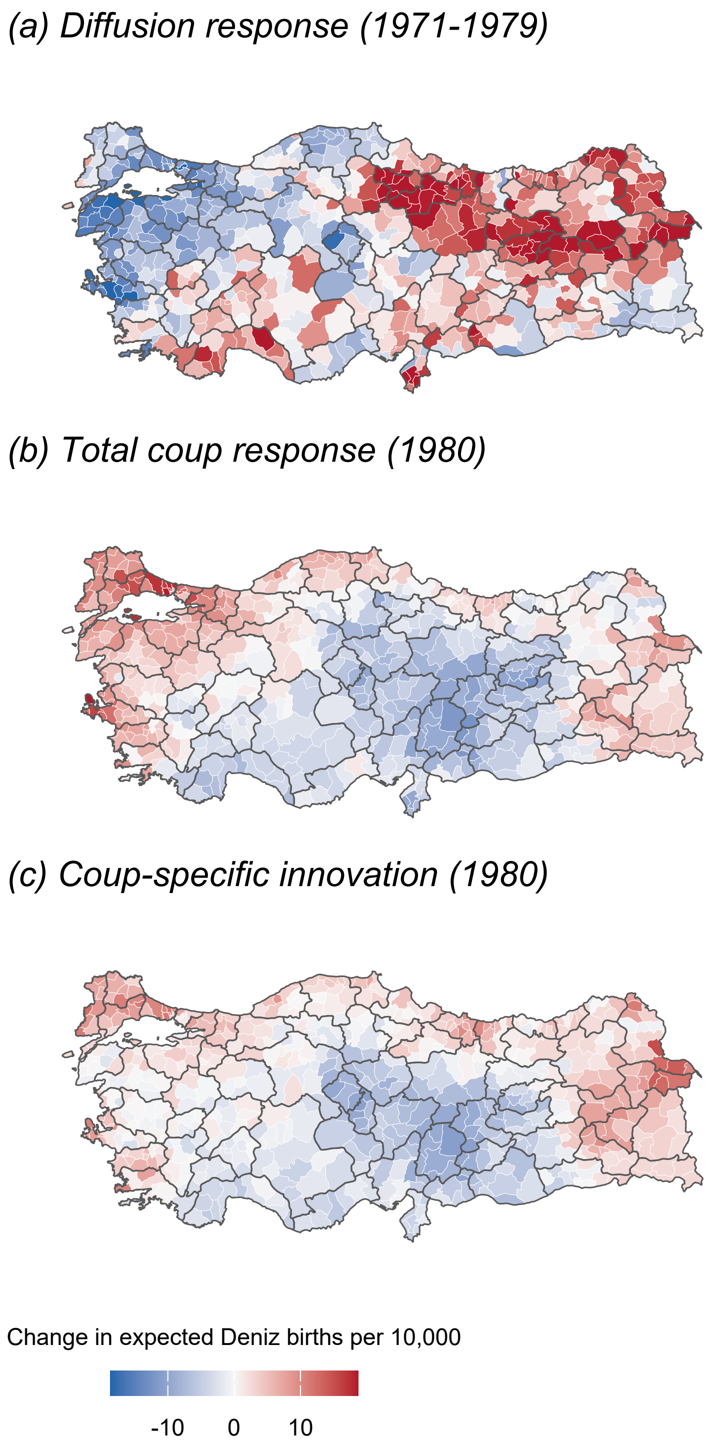

Figure 11: B4 response maps: (a) 1971–1979 diffusion; (b) total coup response; (c) coup-specific innovation. Values are changes in births per 10,000 at the national 1979 rate; these nonlinear translations are nonadditive.

Prior political support provides a further geographic comparison. The modeled coup response is more negative where TİP received greater support in 1969. This association appears in the province model and the Bayesian district analysis, although its province magnitude attenuates under alternative weighting and likelihood choices. Comparisons with other parties also limit a uniquely TİP interpretation. Together, these results connect the contraction to prior political geography without identifying the actors or individual decisions that produced it.

Under equal province-year weighting, province M4 estimates $\theta^{D80} = -0.046$ per percentage point of 1969 TİP support $[-0.088, -0.004]$ (0.020, $p = 0.032$). Birth-weighted least squares and Poisson sensitivities attenuate the point to -0.016 and -0.008. Bayesian district B4 estimates -0.020 per point $[-0.036, -0.004]$ on the log-odds scale. The common negative sign should therefore be distinguished from agreement on magnitude across estimators (Section E.4).

The district log-linear party comparisons further qualify the political interpretation. The recovery interaction with TİP is -0.009 $[-0.019, +0.002]$, while the center-right interaction is $+0.003$ $[+0.0005, +0.006]$. The diffusion comparisons also show associations with center-right support, including an oppositely signed interaction with the lagged neighboring naming rate (Sections 2.3 and E.4). These comparisons do not establish that the geographic variation is unique to TİP.

The province dynamic contrast places the political difference in time. In Figure 12, the naming-rate ratio relative to the projected path is approximately 59 percent higher at the ninetieth than at the tenth percentile of 1969 TİP support in 1979. This median contrast peaks in 1979 and then narrows. Its pointwise 95 percent bands exclude zero from 1975 through 1982, but include zero thereafter. The comparison describes modeled response ratios, rather than a difference in observed naming rates, and remains conditional on equal province-year weighting and the fitted recursion.

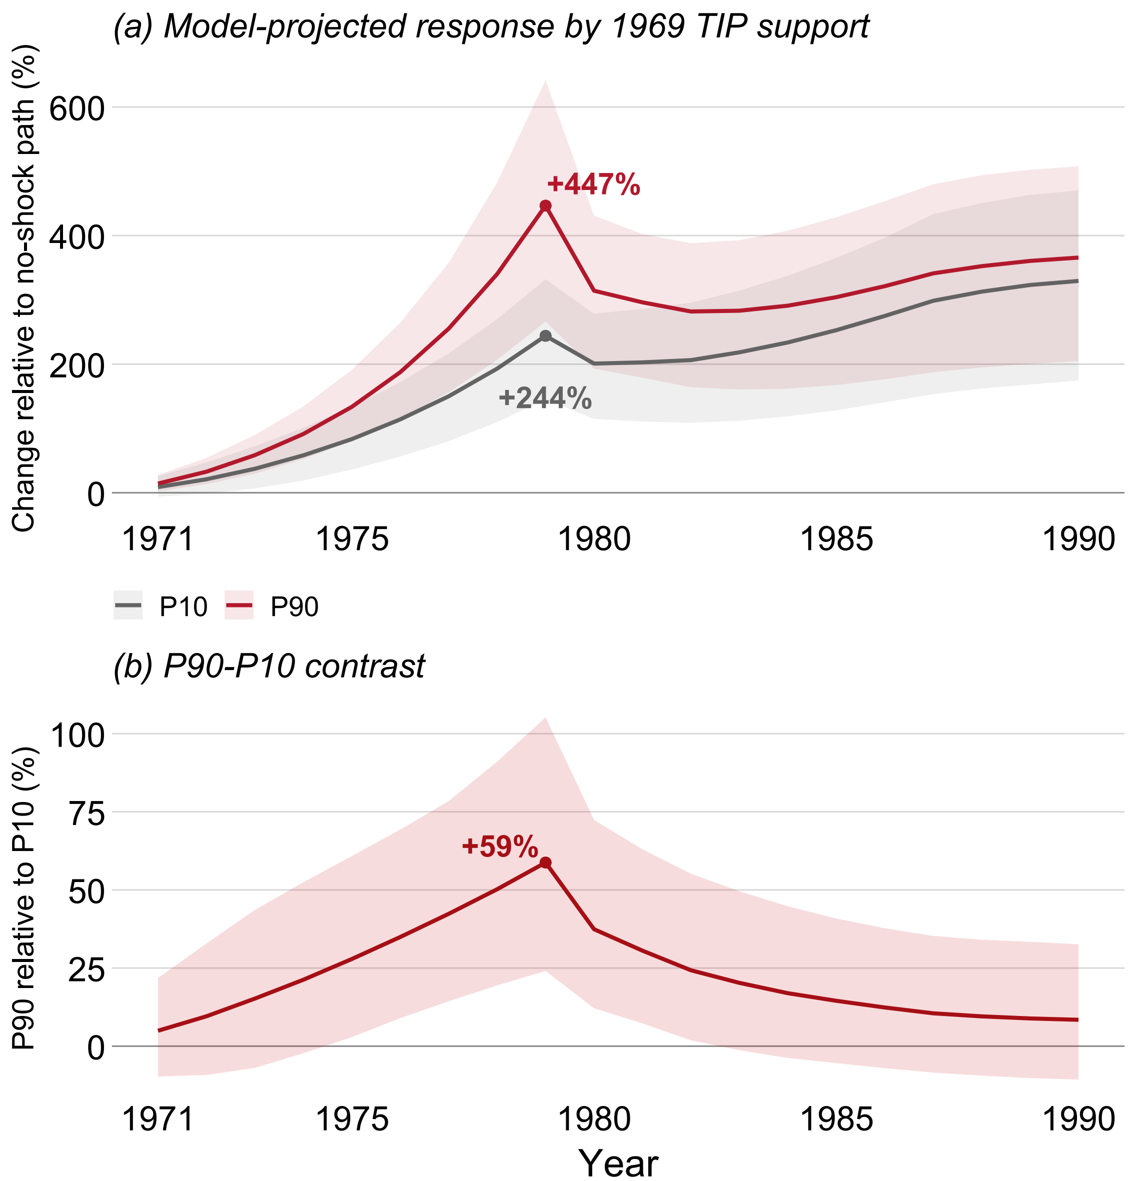

Figure 12: Province M4 responses by TİP support: (a) percentage changes from the no-shock path at the tenth and ninetieth percentiles; (b) their draw-wise contrast, with a 95 percent coefficient-draw band.

6.4. Combined evidence and qualifications

Deniz's trajectory combines an unusually large rise relative to comparable names with a contraction associated with prior political geography. The joint growth and contraction statistic ranks first within the matched pool in both analyses. The province models recover the broad temporal sequence against their projected path, while the district models associate larger contractions with the earlier increase. The contraction is also more negative where prior TİP support was stronger. Together, these comparisons support a political interpretation that does not rest on timing alone.

The contraction reduced Deniz's elevation without eliminating it relative to either benchmark. Its naming rate remained above the matched-name benchmark, and the modeled response remained above the projected path. Persistence of that elevation can therefore coexist with a negative response around the coup. The records leave unresolved when a political response began, why individual parents chose the name, and what produced the geographic concentration.

The interpretation remains conditional on the projection rules, response specifications, and comparability of the selected names. The analyses share registry observations, and the model-based comparisons also share temporal coding, so their agreement is not independent validation (Section 5.4). Attribution specifically to gradual diffusion is sensitive to how the earlier response is divided, and the broader three-correlation pattern is not established as distinctive within the pool. The prediction checklist summarizes these findings and their qualifications (Table 19 and Section E).

7. Discussion

Naming offers evidence about political expression through a choice that retains several possible meanings. Deniz's comparative trajectory and political geography support the interpretation that its political association mattered for adoption (Section 6.1.1 and Table 19). Because the register records naming without asking parents to explain their choice, it allows us to examine behavior alongside research on the divergence between private views and public conduct (Kuran, 1997; Scott, 1990). Yet a family decision becomes part of official records and social interactions. Its ambiguity can coexist with perceived political costs. A contraction in adoption can therefore be consistent with constrained expression while leaving the prevalence of political support unknown.

The geographic evidence helps locate this contraction but leaves its causes open. Larger modeled contractions occurred where the earlier increase was greater and prior TİP support was stronger. Attribution specifically to gradual diffusion depends on the response specification, and the party comparisons do not establish a mechanism unique to TİP (Section 6.3 and Table 6). Targeted repression, anticipation of political costs, detention or emigration, and changes in who has children could contribute to the observed pattern. Shared political and cultural conditions could also shape both periods. The records of births abroad do not measure emigration or establish how departures changed the population having children in Turkey (Figure 5). Distinguishing these explanations requires evidence about household decisions and exposure to repression that this analysis does not observe.

The geographic scope of these inferences also depends on registry coverage. District responses describe births with resolved district identifiers. Extending them to all births within a district requires that resolution does not depend on the child's name within each district and year (Equation (4)). The common-record and nonmetropolitan comparisons retain Deniz's leading comparative rank, but cannot recover the districts of unresolved records (Section B). Improving geographic ascertainment also complicates the interpretation of aggregate district changes and limits direct comparisons with province magnitudes. The size of the registry therefore expands the observable record without resolving its measurement boundaries.

The broader methodological opportunity lies in combining naming trajectories when geographic units share exposure to national political events. A registry can supply comparison names recorded within the same administrative system, but their availability alone does not establish comparability. Research on naming under political pressure and on cultural transmission gives reasons to consider both processes when constructing those comparisons (Algan et al., 2022; Saavedra, 2021). Applications elsewhere require evidence about the focal name's political association, competing cultural influences (J. Berger & Le Mens, 2009; Lieberson, 2000), the timing of relevant events, and recording practices across names and years. The present findings concern Deniz in Turkey from 1960 to 1990. The additional political names provide related comparisons, although their ranks remain descriptive because the comparison pool was selected for Deniz (Section E.10).

Evidence beyond 1990 could examine the persistence of Deniz's trajectory, its geographic association, and its distinctiveness among comparable names. These are separate questions. Continuing the province model would require assumptions about future covariates, parameter stability, and subsequent political or cultural changes (Equation (5)). Saturation of the coded doses by 1987 is therefore insufficient to predict that the observed naming rate approaches equilibrium (Section 6.2). Later acceleration, convergence across districts, or a less extreme comparative rank would qualify claims about persistence. None would by itself rule out a political contribution to the earlier trajectory.

8. Conclusion

Deniz's trajectory supports a political interpretation of naming under authoritarian rule. Its rise during the 1970s and contraction around the 1980 coup were unusually pronounced within the selected cohort of comparable names. The modeled contraction was also greater where the earlier increase was larger and prior TİP support was stronger. Yet Deniz remained above the matched-name benchmark after the contraction. Annual differences from that benchmark accumulate to 38,188 births between 1971 and 1990. This figure summarizes naming relative to the selected comparison names; it does not count politically motivated parents.

To assess this political interpretation against broader cultural influences, we combine projected trajectories with comparisons across names. The panel models estimate departures from a specified continuation of the earlier trajectory and describe their geographic variation. The matched-name analyses assess how distinctive Deniz's modeled response is within its cohort, while the observed comparison trajectories provide a separate benchmark for naming counts. Together, these comparisons make the political interpretation assessable without an unexposed geographic comparison. Their interpretation remains conditional on the projection rules, comparability of the selected names, and registry coverage.

Substantively, the paper extends the study of everyday political practices (Scott, 1985) to a family decision whose social meaning need not be exclusive. A name can express identification with a cause while remaining available for family, cultural, or aesthetic reasons. The Deniz findings are consistent with political conditions shaping the appeal of such a choice, while leaving individual intentions and the mechanisms producing the contraction unresolved. An administrative record of everyday decisions can therefore preserve traces of political history without containing explicit declarations of political allegiance.

References

Abramitzky, R., Boustan, L., & Eriksson, K. (2020). Do Immigrants Assimilate More Slowly Today Than in the Past? *American Economic Review: Insights*, *2*(1), 125–41. https://doi.org/10.1257/AERI.20190079

Algan, Y., Malgouyres, C., Mayer, T., & Thoenig, M. (2022). The Economic Incentives of Cultural Transmission: Spatial Evidence from Naming Patterns Across France. *The Economic Journal*, *132*(642), 437–470. https://doi.org/10.1093/EJ/UEAB058

Bai, J., & Perron, P. (1998). Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica*, *66*(1), 47–78. https://doi.org/10.2307/2998540

Bai, J., & Perron, P. (2003). Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics*, *18*(1), 1–22. https://doi.org/10.1002/jae.659

Baicker, K., & Svoronos, T. (2019, July). *Testing the Validity of the Single Interrupted Time Series Design* (NBER Working Paper No. 26080). National Bureau of Economic Research. Cambridge, MA. https://doi.org/10.3386/w26080

Berger, J., & Le Mens, G. (2009). How Adoption Speed Affects the Abandonment of Cultural Tastes. *Proceedings of the National Academy of Sciences*, *106*(20), 8146–8150. https://doi.org/10.1073/pnas.0812647106

Berger, R. L. (1982). Multiparameter hypothesis testing and acceptance sampling. *Technometrics*, *24*(4), 295–300. https://doi.org/10.1080/00401706.1982.10487790

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination. *American Economic Review*, *94*(4), 991–1013. https://doi.org/10.1257/0002828042002561

Central Intelligence Agency. (1978, December). *Turkey Under Martial Law* (tech. rep. No. RP M 78–10506) (Memorandum dated 27 December 1978, reproduced in Foreign Relations of the United States, 1977–1980, Volume XXI, document 126). Central Intelligence Agency. https://history.state.gov/historicaldocuments/frus1977-80v21/d126

Colin Cameron, A., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business and Economic Statistics*, *29*(2), 238–249. https://doi.org/10.1198/jbes.2010.07136

Commission on Security and Cooperation in Europe. (1988, July). *The State of Human Rights in Turkey* (tech. rep.) (Contemporary account of the lifting of martial law, the continuing state of emergency in 1987, the December 1985 friendly settlement before the European Commission of Human Rights, and the 6 September 1987 referendum). U.S. Commission on Security and Cooperation in Europe (Helsinki Commission). Washington, DC. https://www.csce.gov/wp-content/uploads/2016/07/1988-07-hr-in-turkey.pdf

Dhaene, G., & Jochmans, K. (2015). Split-panel jackknife estimation of fixed-effect models. *The Review of Economic Studies*, *82*(3), 991–1030. https://doi.org/10.1093/restud/rdv007

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *The Review of Economics and Statistics*, *80*(4), 549–560. https://doi.org/10.1162/003465398557825

Fryer, R. G., & Levitt, S. D. (2004). The Causes and Consequences of Distinctively Black Names. *The Quarterly Journal of Economics*, *119*(3), 767–805. https://doi.org/10.1162/0033553041502180

Hausman, C., & Rapson, D. S. (2018). Regression Discontinuity in Time: Considerations for Empirical Applications. *Annual Review of Resource Economics*, *10*(Volume 10, 2018), 533–552. https://doi.org/10.1146/annurev-resource-121517-033306

Hazlett, C., & Xu, Y. (2018). Trajectory Balancing: A General Reweighting Approach to Causal Inference With Time-Series Cross-Sectional Data. *SSRN Electronic Journal*. https://doi.org/10.2139/SSRN.3214231

Human Rights Watch. (1999). *Violations of Free Expression in Turkey* [Documents Penal Code Articles 141, 142 and 163, repealed by the Anti-Terror Law (Law No. 3713) of 12 April 1991, and the tightening of the southeastern emergency regime under Decree-Law No. 413 (1990)]. https://www.hrw.org/reports/pdfs/t/turkey/turkey2916.pdf

Imai, K., & Kim, I. S. (2019). When Should We Use Unit Fixed Effects Regression Models for Causal Inference with Longitudinal Data? *American Journal of Political Science*, *63*(2), 467–490. https://doi.org/10.1111/ajps.12417

Imbens, G. W., & Rubin, D. B. (2015, January). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press. https://doi.org/10.1017/CBO9781139025751

Kuran, T. (1997). *Private truths, public lies : the social consequences of preference falsification*. Harvard University Press. https://www.hup.harvard.edu/books/9780674707580

Lee, L.-f., & Yu, J. (2010). Some recent developments in spatial panel data models. *Regional Science and Urban Economics*, *40*(5), 255–271. https://doi.org/10.1016/j.regsciurbeco.2009.09.002

Li, K. (2017). Fixed-effects dynamic spatial panel data models and impulse response analysis. *Journal of Econometrics*, *198*(1), 102–121. https://doi.org/10.1016/j.jeconom.2017.02.001

Lieberson, S. (2000). *A Matter of Taste*. Yale University Press.

Lopez Bernal, J., Cummins, S., & Gasparrini, A. (2017). Interrupted time series regression for the evaluation of public health interventions: a tutorial [See the corrigendum (Lopez Bernal, Cummins, & Gasparrini, 2021): the equation printed at p. 351 is erroneous, and Supplementary Appendix 5, to which the text defers models (d)–(f), was never published with the article]. *International Journal of Epidemiology*, *46*(1), 348–355. https://doi.org/10.1093/ije/dyw098

Lopez Bernal, J., Cummins, S., & Gasparrini, A. (2021). Corrigendum to: Interrupted time series regression for the evaluation of public health interventions: a tutorial [First published online 3 September 2020. Corrects the segmented-regression equation to $Y_t = \beta_0 + \beta_1 T + \beta_2 X_t + \beta_3 (T - T_0) X_t$, and supplies the missing Supplementary Appendix 5 as Supplementary Appendix 4, in which Table 4.3 specifies a bounded ramp-then-hold intervention regressor. PubMed indexes this item as 49(4):1414 (2020); the publisher's version of record is 50(3):1045]. *International Journal of Epidemiology*, *50*(3), 1045. https://doi.org/10.1093/ije/dyaa118

MacNab, Y. C. (2018). Some Recent Work on Multivariate Gaussian Markov Random Fields. *TEST*, *27*(3), 497–541. https://doi.org/10.1007/s11749-018-0605-3

Martins, T. G., Simpson, D., Lindgren, F., & Rue, H. (2013). Bayesian Computing with INLA: New Features. *Computational Statistics & Data Analysis*, *67*, 68–83.

McDowall, D., McCleary, R., & Bartos, B. J. (2019, October). *Interrupted Time Series Analysis*. Oxford University Press. https://doi.org/10.1093/oso/9780190943943.001.0001

Menchetti, F., Cipollini, F., & Mealli, F. (2023). Combining counterfactual outcomes and ARIMA models for policy evaluation. *The Econometrics Journal*, *26*(1), 1–24. https://doi.org/10.1093/ectj/utac024

Mundlak, Y. (1978). On the Pooling of Time Series and Cross Section Data. *Econometrica*, *46*(1), 69–85. https://doi.org/10.2307/1913646

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, *49*(6), 1417–1426. https://doi.org/10.2307/1911408

Nişanyan, S. (2025). Nişanyan Adlar: Türkiye Kişi Adları Sözlüğü [Online database, accessed May 2026]. https://www.nisanyanadlar.com

Öztan, G. G., & Bezci, E. B. (2015). Türkiye'de Olağanüstü Hal: Devlet Aklı, Askerler ve Siviller [Dates the OHAL regime to the National Security Council decision of 1 March 1984, when martial law was lifted in the first thirteen provinces and replaced by a state of emergency]. *Mülkiye Dergisi*, *39*(1), 159–186.

Pesarin, F., & Salmaso, L. (2010, March). *Permutation Tests for Complex Data: Theory, Applications and Software*. Wiley. https://doi.org/10.1002/9780470689516

Riebler, A., Sørbye, S. H., Simpson, D., Rue, H., Lawson, A. B., Lee, D., & MacNab, Y. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, *25*(4), 1145–1165. https://doi.org/10.1177/0962280216660421

Rosenbaum, P. R. (2020). *Design of Observational Studies*. Springer International Publishing. https://doi.org/10.1007/978-3-030-46405-9

Saavedra, M. (2021). Kenji or Kenneth? Pearl Harbor and Japanese-American assimilation. *Journal of Economic Behavior & Organization*, *185*, 602–624. https://doi.org/10.1016/j.jebo.2021.03.014

Sayari, S. (2010). Political Violence and Terrorism in Turkey, 1976–80: A Retrospective Analysis. *Terrorism and Political Violence*, *22*(2), 198–215. https://doi.org/10.1080/09546550903574438

Scott, J. C. (1985). *Weapons of the Weak: Everyday Forms of Peasant Resistance*. Yale University Press.

Scott, J. C. (1990). *Domination and the Arts of Resistance: Hidden Transcripts*. Yale University Press.

Simpson, D., Rue, H., Riebler, A., Martins, T. G., & Sørbye, S. H. (2017). Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors. *Statistical Science*, *32*(1), 1–28. https://doi.org/10.1214/16-STS576

Stuart, E. A. (2010). Matching Methods for Causal Inference: A Review and a Look Forward. *Statistical Science*, *25*(1), 1–21. https://doi.org/10.1214/09-STS313

Türköz, M. (2017). *Naming and Nation-building in Turkey: The 1934 Surname Law*. Palgrave Macmillan. https://doi.org/10.1057/978-1-137-56656-0

A. Data Construction and Cleaning

Section 4.1 reports the record counts from which the analysis begins. This appendix collects the notation used throughout the paper and then summarizes the cleaning behind those counts: the normalization of name strings, the separation of births abroad, the resolution of birthplaces against the constant 1961 boundary set, and the construction of the population denominators Section 4.3 uses as exposure. The full phase-by-phase documentation, including every correction file, accompanies the replication materials.

A.1. Notation

The indicator function is $\mathbb{1}\{\cdot\}$, and $\mathbf{W}$ is the row-standardized queen contiguity matrix. Hats on observed proportions denote Jeffreys smoothing; hats on fitted paths denote model estimates. The notation $\rho(\cdot,\cdot)$ denotes response correlations, while ρ_1 and ρ_2 are temporal lag coefficients. Regime coefficients give contributions at full dose in M1 and M2. In dynamic models, they are conditional on the lagged state; propagated annual responses and steady-state contributions are reported separately (Section 5.1).

Table 7: Notation and variable definitions.

Symbol	Meaning
Indices, dimensions, and measurement	
i, j, t	District, province, and year. 636 districts, 67 provinces, 31 years from 1960 to 1990; 19,434 observed district-years and 2,077 province-years.
y_{it}	Count of newborns registered as Deniz in cell (i,t).
n_{it}	All births registered in cell (i,t); the denominator, and the binomial trial count in the district analysis.
$\hat{p}$	Jeffreys posterior-mean rate, $(y+0.5)/(n+1)$ under a Beta(0.5, 0.5) prior; the province analysis outcome is $\log \hat{p}$.
ω	Pre-switch unresolved-to-resolved Deniz-rate ratio in the six metropolitan offices, estimated over 1984 to 1986; $\omega = 1.256$ (Section B).
Regime doses (Equation (1)); each spans 0 to 1 and is constant once saturated	
S_t^{pre}	Secular pre-trend, accumulating 1960 to 1970.
D_t^{71}	Step at the 1971 shock: 0 before, 1 from 1971 on.
S_t^{71}	Diffusion dose, accumulating 1971 to 1979.

continued overleaf

Table 7, continued

Symbol	Meaning
D_t^{80}	Step at the 1980 coup: 0 before, 1 from 1980 on.
S_t^{80}	Post-coup *recovery* dose, accumulating 1980 to 1987.
Population-level coefficients (both analyses)	
β_{pre}	Coefficient on the secular pre-trend.
γ_1, γ_2	Level shifts at the 1971 and 1980 steps.
δ_1, δ_2	Coefficients on the diffusion and recovery doses.
ρ_1, ρ_2	First and second temporal lags of the outcome.
λ	Spatial lag: coefficient on $[\mathbf{W} \log \hat{p}]_{\cdot,t-1}$.
$\mathcal{M}$	Long-run multiplier, $(1 - \rho_1 - \rho_2 - \lambda)^{-1}$.
$\beta_w, \beta_{bw}, \beta_b$	Mundlak within, between-district-within-province, and between-province components of log exposure.
β_c	Coastal secular trend, distance to coast $\times$ year.
Political moderators	
tip_i, tip_j	Unit share of the 1969 vote for the Turkish Workers' Party, time-invariant; indexed by district in the Bayesian analysis and by province in the log-linear analysis. Natural percentage points in every analysis, centered at the births-weighted mean share among the units in that scale's top quartile of support, 4.850 points across provinces and 5.712 across districts (Sections 5.1 and 5.2). Moderation coefficients therefore read per percentage point at both scales, and the reading at the common cross-scale reference is recovered as $\beta + \theta(\bar{x} - c)$ for centering reference c and $\bar{x} = 2.666$, the mean 1969 district share.
ap_i, ap_j	Center-right share of the same election on the same scaling; the prespecified negative control.

continued overleaf

Table 7, continued

Symbol	Meaning
$\theta^{\text{pre}}, \ldots, \theta^{80}$	TİP moderation coefficients, one on each of the five doses: $\theta^{\text{pre}}, \theta^{D71}, \theta^{71}, \theta^{D80}, \theta^{80}$.
λ^{TIP}	Coefficient on tip_i interacted with the spatial lag; district scale only.
ψ	Main effect of tip_i where it is estimable; absorbed by unit fixed effects in the log-linear analyses.
Latent objects, district Bayesian analysis	
$u_{j[i]}$, v_i, $s_i^{(0)}$	Exchangeable province intercept, and the unstructured and spatially structured components of the district BYM2 intercept.
$s_i^{(71)}$, $s_i^{(80\text{D})}$, $s_i^{(80\text{S})}$	District-varying diffusion, coup, and recovery responses.
$\xi_i^{(2)}$, $\xi_i^{(3)}$	Coup and recovery innovation fields: what each regime did that the preceding fields do not explain.
$\kappa_{21}, \kappa_{31}, \kappa_{32}$	Coregionalization loadings linking the fields; κ_{21} carries the diffusion-to-coup coupling.
$\rho(\cdot, \cdot)$	Correlation between two district-varying responses; the estimand of the geography.
$\tau_\bullet$, $\phi_\bullet$	Precision and BYM2 mixing of each field.

A.2. Cleaning pipeline summary

Five deterministic phases carry the raw extract to the analysis panels, and each phase writes the correction files that make its assignments reproducible. The deduplication pass removes 1,013,640 exact duplicate records, 3.22 percent of the extract. Because duplication is strongly non-random by province, this step protects the spatial results in particular. String normalization reduces 167,035 unique birthplace spellings to a matchable form. Variant recovery folds 2,520 records carrying Deniz spelling variants and 23 approved concatenated forms into the canonical indicator. A wider variant set that adds 774 disputed fused forms serves only as a sensitivity definition. The foreign-birth cascade classifies 459,941 records, 1.5 percent, as births abroad through country extraction, a curated dictionary, 403 reviewed patterns, and a 4,519-entry corrections lookup, while sparing Turkish place names that resemble foreign ones. Domestic birthplaces then resolve against the constant 1961 hierarchy through exact matching over 96,190 unique values, 3,562 unconditional manual corrections covering 1.86 million records, and a bounded fuzzy pass. Entries whose district cannot be established are left unassigned rather than guessed, and 181 districts created after 1961 map to their parent units.

The exposure series combines census counts, observed for 617 of the 636 districts, with a registered-voter fallback for the remaining 19, rescaled by the census-to-voter ratio observed at district level in 1965 and interpolated to annual frequency. The 1969 electoral returns are observed for 634 historical districts, with the two later boundary districts inheriting their parent's shares. After every correction pass, the aggregated province-year birth totals are re-certified against the raw extract cell by cell with zero discrepancy.

B. Missing District Identifiers: Mechanism, Assumption, and Sensitivity

The province and district panels resolve different fractions of the same registrations. Section 4.2 states the identifying condition and assigns the two scales to different questions. This appendix documents where resolution fails, which implications of that condition are observable, and whether record selection can explain the matched-name result.

B.1. Record loss and measurement scope

The principal loss occurs after province assignment. District identifiers are absent from 8,332,924 of 28,394,417 province-resolved births, or 29.3 percent. Deniz has lower district retention than births generally: the district panel contains 38,340 of 65,732 province-

Table 8: Registry resolution into province and district panels. Retention uses raw Deniz records; analysis coverage uses domestic Deniz births.

Stage	Records	Deniz	Deniz lost	Deniz kept (%)
all records	30,455,765	68,956	—	100.0
domestic	29,995,824	67,094	1,862	97.3
domestic, analysis window	29,995,824	67,094	0	97.3
province resolved	28,394,417	65,732	1,362	95.3
district resolved (panel)	20,061,493	38,340	27,392	55.6

The pattern-mixture profile uses the pre-switch enrichment ratio $\omega = 1.256$, estimated from 1984 to 1986 metropolitan records. The largest single-year attrition step in the pooled metropolitan series falls in 1987 (38.1 % $\rightarrow$ 25.2 %); the practice itself ended office by office, not on that date.

resolved Deniz births. By contrast, only 282 of 19,716 possible district-year cells contain no district-resolved domestic birth. Bounds for those empty cells address a small grid-completion problem, not the millions of records with known provinces and latent districts.

The district-to-province ascertainment ratio rises from 0.52 in 1960 to 0.99 in 1990. Its log change is -0.052 over 1970 to 1971 and $+0.019$ over 1979 to 1980, compared with $+0.169$ over 1971 to 1979 and $+0.175$ over 1980 to 1987. Improving resolution is therefore more closely aligned with the accumulating doses than with the abrupt steps. This pattern cautions against comparing aggregate district and province magnitudes. It does not by itself establish bias in a district-specific naming rate. Metropolitan offices changed their recording practices at different dates; 1987 marks the largest pooled movement rather than a common switch.

B.2. Mechanism and identifying assumption

Metropolitan registration offices often entered a bare province in the birthplace field rather than a district. Deniz was concentrated in these urban records. Conditional on a district string being present, however, Deniz and other births resolve at essentially the same rate. Place names containing the string Deniz also resolve, which rules out a name-string matching defect. The loss arises from what the office recorded, not from how the cleaning pipeline processed the string.

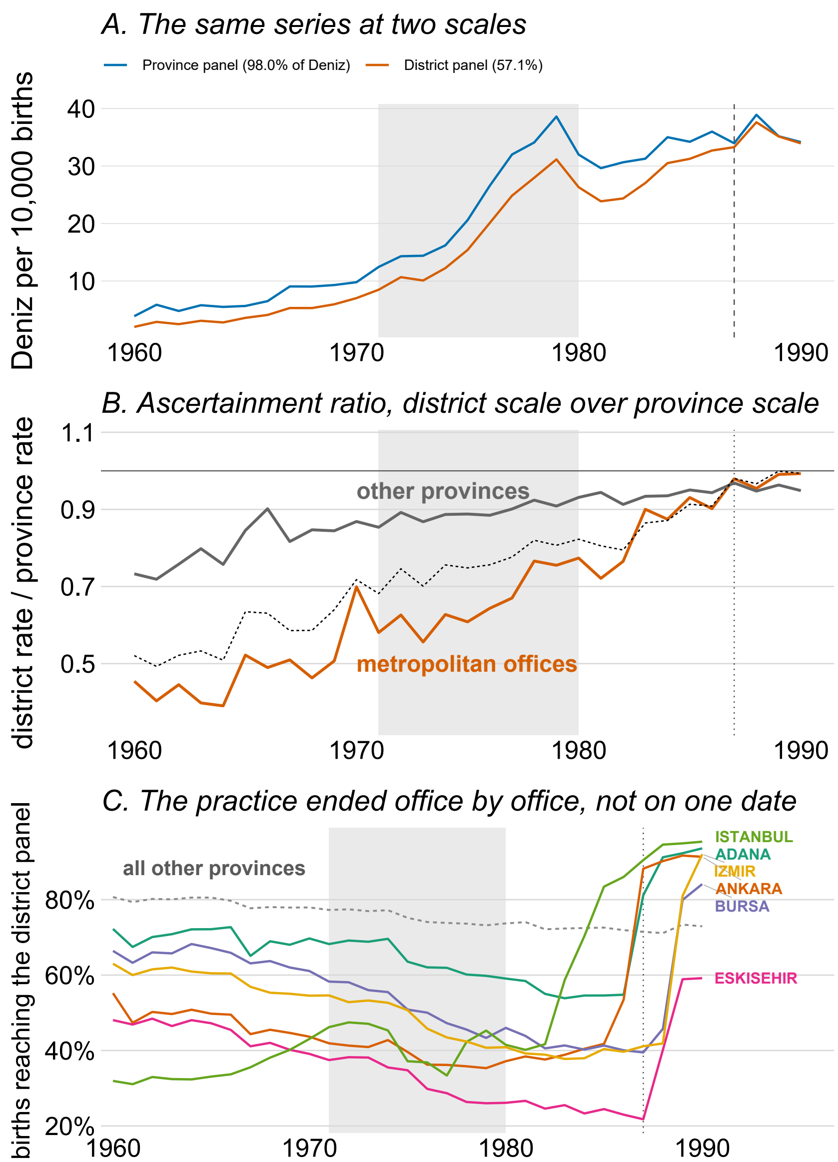

Figure 13: District resolution: (a) Deniz rates at both geographic scales; (b) district-to-province rate ratios, with the national ratio dashed; (c) district-identifier coverage by metropolitan office.

The district analyses extend to all births only under the name-neutral proportional-thinning condition defined in Equation (4). This condition is not ordinary missingness at random on observed data because true district G^* is latent when a record does not resolve. Within an actual district-year, it requires district resolution to be independent of the given name after conditioning on true geography and year. If it holds, name counts and birth denominators are thinned together in expectation. The district likelihood then loses information but does not mechanically change the expected conditional naming rate. It does not recover unresolved births or their geography.

The matched-name diagnostics probe two observable margins of this condition. Table 9 and Figure 14 first compares observed retention with a name-free prediction constructed from each name's province-year composition. The prediction for Deniz is 66.3 percent, eight percentage points above its observed 58.3 percent retention. This is the third-largest absolute residual in the 41-name pool, so province-year composition does not fully explain Deniz's under-resolution. Unobserved within-province urban sorting remains plausible.

Table 9: District-resolution diagnostics. Composition ranks order absolute residuals from largest; office-switch ranks order absolute errors from smallest. Reference reports the conditional finite-pool probability $m/41$.

Diagnostic	Deniz	Position / 41	Reference
Province-year composition	-8.0 pp retention residual	3 / 41	0.073
Office-switch mixture	-3.3% prediction error	4 / 41	0.098
Pre-switch enrichment	$\omega = 1.256$	—	—

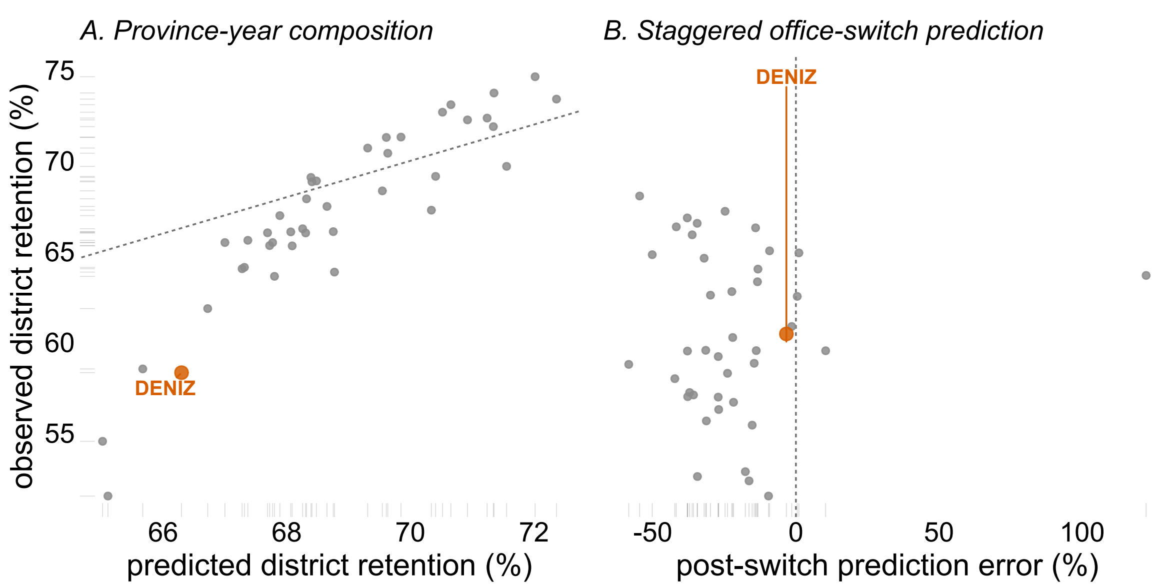

Figure 14: District-resolution diagnostics: (a) observed retention against province-year composition predictions; (b) post-switch rate prediction errors. Deniz is highlighted.

The second diagnostic uses staggered changes in metropolitan recording practice. Pre-switch resolved and unresolved rates combined with the observed coverage gain predict the subsequent resolved rate without a fitted parameter. Deniz's prediction error is -3.3 percent, the fourth-smallest absolute error in the pool. The same pre-switch records yield an unresolved-to-resolved rate ratio of $\omega = 1.256$. This result supports name-neutral release at the administrative margin, but the office changes occur late in the panel and do not verify the condition during the 1971 diffusion window. No observed stratum identifies the latent districts of province-only records, so assigning them to *merkez* would manufacture geography.

B.3. Common-record and geographic-scale sensitivity

The matched-name sensitivity separates record restriction from the bundled district-scale specification. The first frame uses all province-resolved records in province M4. The second restricts province M4 to district-resolved records reaggregated to the same 67 provinces. The third holds those resolved records fixed but estimates district M4 at their observed district scale. A fourth M4 analysis excludes districts in the six metropolitan provinces, changing the target population to locations outside the offices that generated most missingness. District B2 provides an estimator-family comparison on the district-resolved frame. Rank stability is evaluated over 5,000 coefficient draws combining name-specific marginal uncertainty with the corresponding focal-fit correlation template. The frequencies are diagnostics, not posterior model probabilities.

Table 10: Matched-name sensitivity across record and geographic restrictions. Frequencies report rank stability over 5,000 coefficient draws, not posterior model probabilities.

Analysis	S^{71}	D^{80}	D^2	Rank	$P(r=1)$	$P(r \leq 2)$
Province M4: all province-resolved records	+0.658	-0.289	15.92	1 / 41	42.1%	82.3%
Province M4: district-resolved records	+0.700	-0.288	14.36	1 / 41	62.6%	88.1%
District M4: district-resolved records	+0.458	-0.206	20.19	1 / 41	89.1%	97.2%
District M4: nonmetropolitan records	+0.435	-0.196	18.44	1 / 41	83.8%	94.8%
District B2: district-resolved records	+1.450	-0.264	12.53	1 / 41	95.4%	99.9%

All five point estimates rank Deniz first of 41 on the prespecified two-coordinate Mahalanobis statistic. Restricting province M4 to district-resolved records increases the simulated rank-one frequency from 42.1 to 62.6 percent. Moving those same records to district

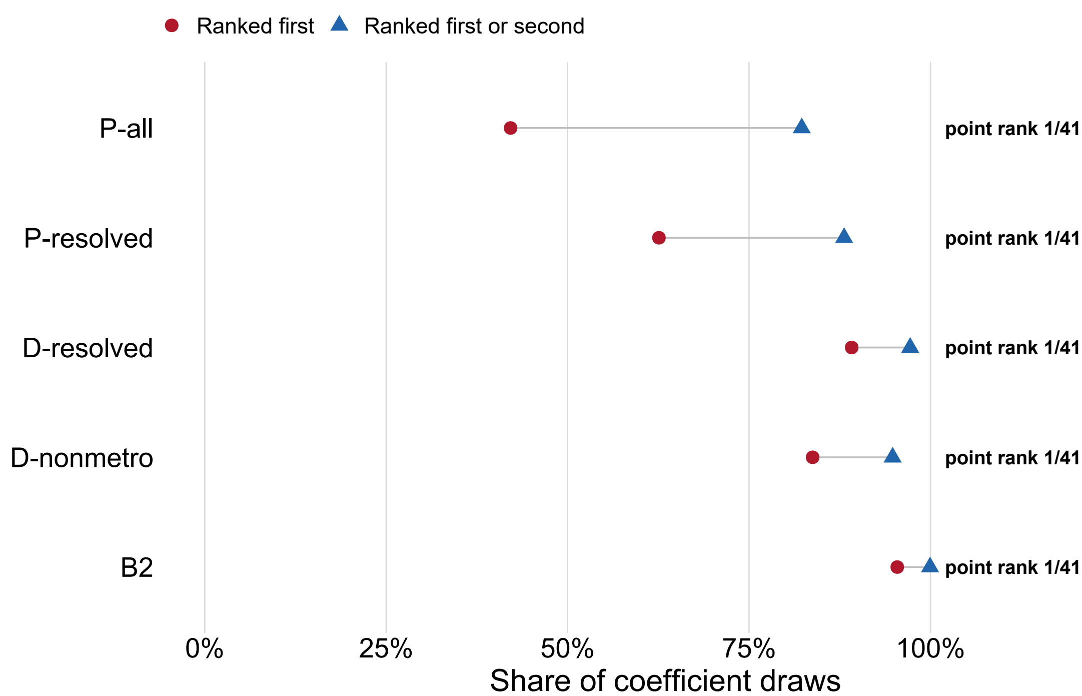

Figure 15: Matched-name rank stability across record and geographic restrictions. Points show frequencies of ranks one and two-or-better; labels report point-estimate ranks.

scale raises it to 89.1 percent. Excluding metropolitan provinces lowers the district value to 83.8 percent, while B2 reaches 95.4 percent. The first-place point rank is unchanged, while its simulated stability varies across these choices. The comparison using the same records at district scale also changes the fixed effects, covariate construction, and leverage. It therefore does not isolate the contribution of geographic resolution.

Greater simulated stability in the district fits must also be read alongside their selective coverage. Every district analysis uses the 57.1 percent of domestic Deniz births whose districts resolve. Deniz's metropolitan share of 0.534 stands at the top of the pool, its retention residual of -8.0 percentage points is the third largest of 41, and the office-switch evidence for name-neutral release comes from 1984 to 1986 rather than the diffusion window. The point rank remains first in all five observed frames, including the nonmetropolitan one. These comparisons do not establish name-neutral thinning inside the latent districts during the 1970s.

The same facts motivate the strongest surviving rival account. The most metropolitan name in the pool could ride an urban naming fashion that accelerated fastest through the 1970s and reversed most sharply around 1980. Urbanity was not a matching feature, so the donors do not bracket Deniz on it. Three observations answer the rival without eliminating it. The nonmetropolitan district frame retains the first rank at 83.8 percent stability, so the political signal does not live only in the metropolitan offices. The donor pool is drawn from the same secular reform cohort, so its realized paths embed the era's modern-naming fashion. Finally, a fashion account must still reproduce the conjunction: acceleration during the early 1970s, contraction during the repression period surrounding the coup, and a deeper contraction where the left had polled strongest. These comparisons constrain the rival account without eliminating it or locating either change at the precise date of an individual event.

B.4. Focal sensitivity and inferential limits

Two focal district-M4 checks vary the treatment of incomplete resolution without assigning any province-only birth to a district. A pattern-mixture profile multiplies each observed district rate by $1 + (\omega - 1)u_{pt}$, where u_{pt} is the unresolved share in its province-year. Across $\omega \in \{1, 1.128, 1.256, 1.50\}$, S^{71} moves from 0.458 to 0.466 and D^{80} from -0.206 to -0.203. A separate specification adds province-year resolution and its regime interactions (four in total). Its average-reference estimates are $S^{71} = 0.435$ and $D^{80} = -0.189$.

Table 11: District-M4 resolution sensitivities by unresolved-to-resolved rate ratio ω and coverage controls. Driscoll–Kraay standard errors in parentheses.

Sensitivity	S^{71}	D^{80}
1.000	+0.458 (0.062)	-0.206 (0.034)
1.128	+0.461 (0.063)	-0.205 (0.035)
1.256	+0.463 (0.064)	-0.205 (0.035)
1.500	+0.466 (0.065)	-0.203 (0.037)
Coverage controls	+0.435 (0.061)	-0.189 (0.035)

These exercises preserve the focal signs but do not identify the missing districts.

The pattern-mixture profile imposes a common within-province enrichment factor. The coverage-control specification conditions on a path that may itself be contemporaneous with the interventions and is therefore a sensitivity rather than the preferred model. The defensible scope is consequently narrow. Province M4 remains primary for average dynamics because missing district identifiers do not remove records from their known provinces. District M4 and B2 provide geographically richer corroboration among resolved births. Extending their conditional rates to all births requires Equation (4), while their aggregate magnitudes remain exposed to changing ascertainment.

C. Estimands and Empirical Strategies

Section 3.5 defines the four estimands on which the argument rests. This appendix maps each estimand to its design comparison, empirical strategy, and reporting location. Quantities carried by different comparisons are not interchangeable, even when they share notation, because measurement constraints give the strategies different scales and counterfactual benchmarks (Section 4.2).

C.1. The design beside difference-in-differences

Figure 4 places the two counterfactual logics beside one another. Both rely on an unobserved counterfactual: difference-in-differences borrows it from contemporaneous controls, whereas an interrupted series projects the treated series from its own past. The distinction is less absolute than it first appears. Two-way fixed effects still depend on no-carryover and sequential-ignorability conditions (Imai & Kim, 2019, pp. 469–476); designs that condition on pre-treatment outcomes assume a functional relation in those outcomes (Hazlett & Xu, 2018, pp. 9–10). Where a shock reaches every unit, a within-series counterfactual is therefore a formalized design choice rather than a second-best regression (Menchetti et al., 2023, p. 2).

Table 12: Estimands, design comparisons, and reporting locations.

Estimand	Design comparison	Analysis	Reported in
$\tau_k^{\text{dyn}}(h)$: propagated log-rate contrast, also reported as a percentage	Episode versus no-shock path at the common TİP reference, pooled over 67 provinces	P, province log-linear panel	Section 6.2
E_T^M, E_T^P, E_T^B: cumulative differences in birth counts	Observed Deniz minus realized cohort benchmark (M); fitted factual minus no-shock counts on each resolved population (P and B)	M, P, and B	Section 6.1.2, Section E.9
ρ_Δ: geographic coupling of heterogeneous responses	Cross-district covariation of latent responses within one name, over 636 districts	B, district Bayesian multilevel analysis	Section 6.3
p_{exact}: comparative position in the pool	Cross-name variation over the frozen pool of 41, conditional on within-pool exchangeability; all name positions enumerated	M, matched-pool analysis using statistics from P and B	Section 6.1.1

C.2. The assumption ledger

The assumptions in Table 13 are the design-specific face of consistency, positivity, sequential ignorability, and conditional stationarity (Menchetti et al., 2023, pp. 5–7). They govern the causal and temporal comparisons. They do not subsume the separate measurement condition for incomplete district identifiers. That condition is defined in Equation (4), and its mechanism and sensitivity evidence appear in Section B.

Table 13: Design assumptions, threats, and corresponding checks.

Assumption	*Statement*	*Principal threat*	*Design answer*
A1. Pre-trend extrapolability	The untreated path is recoverable from pre-1971 history.	Alternative pre-trend rules fit the 1960s similarly but diverge when projected; changing polynomial order can move a nominal interval off its predecessor (Imbens & Rubin, 2015, p. 336).	Model-based analyses retain the assumption, holding the secular component constant after 1970. The cohort benchmark instead uses the realized donor path and does not extrapolate Deniz's pretrend (Section 3.2).
A2. No earlier response	The modeled changes do not begin before their coded boundaries.	The early break interval extends into 1970; the later break precedes the 1980 coup.	Annual coding includes the March 1971 point estimate but cannot exclude a response in the baseline. The November 1979 break limits attribution of the contraction to the coup alone (Section 2.1).
A3. No differential concurrent shocks	Other changes around the political episodes moved donor names too.	Registry changes, urbanization, or a name-specific fashion could move Deniz differently.	Donors share the administrative system, but equal responses to recording changes remain an assumption. The comparative trajectories and coverage checks in Section 4.2 constrain rival accounts without excluding differential recording or cultural influences.

Table 13: The assumption ledger *(continued)*.

Assumption	*Statement*	*Principal threat*	*Design answer*
A4. Intervention exogeneity	The historical events did not respond to naming.	Event timing could be selected by the behavior under study.	A tribunal sentence and a coup amid escalating violence (Sayari, 2010) are not plausible responses to birth registrations. This does not establish when naming responded or isolate an effect of either event from the surrounding political episode.
A5. Known dose form	The intervention ramps and termini are correctly specified.	Results may turn on saturation or on an unbounded forcing.	The equal-increment ramp is a maintained form (Lopez Bernal et al., 2021, Supplementary Appendix 4, Table 4.3); bounded regressors permit equilibrium (Li, 2017, Assumption F, p. 6). Historical justification is needed for the windows (Lopez Bernal et al., 2017, p. 351): 1987 marks an institutional transition, not a measured end to repression. The coding audit assesses sensitivity (Section D.2); Section 3.3 reports collinearity.

Table 13: The assumption ledger *(continued)*.

Assumption	*Statement*	*Principal threat*	*Design answer*
A6. No substitution across names	The interventions did not move donor paths through a fixed naming budget.	Substitution toward Deniz after 1971 and away after 1980 could alter the donor benchmark and mimic part of the rise and contraction.	The donor jackknife assesses dependence on individual donors. It does not rule out collective displacement across the pool, whose magnitude requires separate assessment.

D. Model Detail

Sections 5.1 and 5.2 state the two estimating specifications and leave their implementation detail to this appendix. We document the covariance choice for the province panel, the coding-invariance audit behind the dose windows, the district Bayesian likelihood and priors, the membership of the covariance block, the provenance of the model-comparison criteria, and the split-panel jackknife that bounds dynamic-panel bias in the propagation kernel.

D.1. Covariance inference for the province panel

The province analysis uses Driscoll–Kraay standard errors with a t_{28} reference, allowing cross-sectional and serial dependence without requiring a fixed residual graph. The raw two-way province-and-year covariance is not positive semidefinite. Its eigenvalue repair leaves point estimates unchanged but can compress uncertainty, most visibly for the pre-window TİP moderator, which appears significant only under the repaired matrix. Province-only, year-only, heteroskedasticity-robust, and repaired two-way estimates therefore form a sensitivity battery. Driscoll–Kraay remains the prespecified covariance because it is positive definite and does not depend on that numerical repair (Section 5.1).

D.2. Coding invariance

All three empirical strategies share the regime dating and dose coding, so agreement among them cannot validate the coding itself (Section 5.4). The coding audit moves each dose terminus through the range admitted by the institutional record, considers an uncapped variant and a static segmented variant, and recomputes the prespecified quantities at every cell. These checks are sensitivity analyses, not new identification designs: they show how the fitted trajectories depend on the chosen exposure window and cannot distinguish a politically dated disruption from an unmeasured, contemporaneous national change.

The dated keyings are the audited cells with the most institutional content, and we report those cells here rather than the full grid, whose uncapped and static segmented variants are summarized by the separation they buy in Section 3.3. The audit engine re-estimates the M4 specification on its own frame, so its baseline coefficients differ slightly from the average-referenced values of Section 6.2. The within-engine contrasts are the evidence. When the first boundary moves from 1971 to the execution year of 1972 and the baseline extends through 1971, the engine's diffusion dose falls from 0.584 to 0.507 while the coup shift holds at -0.285 against -0.286, the model-projected cumulative gap falls from 23,273 to 21,159 births, and Deniz retains the first joint rank in the matched pool under both keyings, with the diffusion component slipping from second to third.

The alternative second boundary is 1979, the first full year of the martial law declared in December 1978. This specification also ends diffusion in 1978 and extends recovery from 1979 through 1990, whereas the main specification caps recovery in 1987. Under these combined changes, the downward shift weakens from -0.286 to -0.153, the joint distance falls from 13.40 to 8.12, and the joint rank moves from first to second of 41, at $p = 0.049$. The signed signature retains the first rank. The positive diffusion loading and negative shift persist across these specifications, but their magnitudes and comparative ranks depend on the coding. Stronger results under the 1971 and 1980 boundaries do not validate those dates, and the existing 1979 comparison does not isolate the contribution of the starting date from that of the recovery window.

D.3. District Bayesian specification and priors

The district likelihood is binomial with the count of Deniz registrations as the outcome and all registered births as the denominator. Every B0 through B4 specification includes a global intercept, the fixed-effect trajectory in Equation (6), an exchangeable province intercept $u_{j[i]}$, and a district BYM2 intercept whose unstructured and spatial parts are v_i and $s_i^{(0)}$ (Table 7). Thus B0 is a full two-intercept baseline,

$$\alpha + \boldsymbol{x}'_{it}\boldsymbol{\beta} + u_{j[i]} + v_i + s_i^{(0)},$$

with no response-slope block. B1 and B2 add nonspatial response slopes; B3 and B4 instead add district response geography, as summarized in Table 1. The fixed trajectory and denominator are held constant across these fits.

Non-TİP population coefficients have $N(0, 2.5^2)$ priors, the global intercept has $N(0, 5^2)$, and the six TİP terms have the prespecified $N(0, 0.10^2)$ prior per vote-share point. The province intercept has a penalized-complexity prior $\Pr(\sigma_v > 1) = 0.01$. District-intercept, B1, B3, and B4 innovation standard deviations share the same penalized-complexity scale prior; each BYM2 mixing fraction has $\mathrm{PC}(0.5, 0.5)$ (Riebler et al., 2016; Simpson et al., 2017). B2 uses the prespecified three-dimensional identity-scale Wishart precision prior with five degrees of freedom. B4 uses standard-normal priors on its free triangular loadings. The response-slope priors are therefore distinct from, and do not replace, the two intercept components.

The district graph is queen contiguity with two ferry links joining four components.

BYM2 fields use the graph scale, sum-to-zero constraints, and component adjustment required by that disconnected graph. The restricted-spatial-regression check residualizes the centered TİP moderator against the spatial basis; it is a diagnostic for confounding between that moderator and the graph, not a new estimand or a bound on the TİP interaction. No current alternative-graph refit is available, so the queen graph and ferry links remain maintained assumptions rather than empirically verified invariances.

All B0 through B4 fits share the binomial likelihood, adaptive strategy, and CCD hyperparameter integration used for the primary analysis. Only B2 and B4 retain INLA configurations for their joint summaries. B2 direct correlations are summarized from direct marginal draws. B4 correlations and variance shares are reconstructed from joint hyperparameter draws, which propagate the loading and scale uncertainty; the published structural summaries use 5,000 such draws. Empirical-Bayes sensitivity refits are labeled as plug-in quantities and are not presented as posterior intervals.

D.4. Covariance-block membership

The prespecified covariance block contains the diffusion response S^{71}, the coup response D^{80}, and the recovery response S^{80}. B2 estimates their total correlations; B4 represents the same three responses through spatial coregionalization. The fixed effects in both models also include the 1971 step.

The available alternative uses a trivariate block (S^{71}, D^{71}, D^{80}), replacing the varying recovery response with the varying 1971 step. The fixed recovery term remains. It therefore changes both step and recovery heterogeneity, rather than adding a fourth response to B2. In this empirical-Bayes diagnostic, $\rho(S^{71}, D^{71}) = -0.391$, $\rho(S^{71}, D^{80}) = -0.133$, and $\rho(D^{71}, D^{80}) = -0.524$. The combined earlier contribution has correlation $\rho(S^{71} + D^{71}, D^{80}) = -0.588$ with the coup response.

The pairwise quantities are empirical-Bayes posterior-mean summaries; the combined contrast is calculated from the posterior-mean covariance matrix. They are not directly comparable to the primary B2 posterior summary. Within the alternative fit, the step and combined earlier contribution are more negatively associated with the coup response than the gradual component. This qualifies attribution specifically to gradual diffusion, but the joint specification changes do not isolate the effect of permitting the step to vary. The main analysis retains B2 and B4; Section 6.3 states this qualification alongside their geographic results.

D.5. Extraction and diagnostic provenance

The model-comparison criteria describe fit within the specified family, not a post-treatment outcome definition. Relative to B0, the factorial comparisons yield DIC and WAIC changes of 1,037.8 and 1,038.5 for independent nonspatial response slopes (B1); relative to B1, the correlation block improves them by 21.0 and 24.4 (B2), while independent response geography improves them by 132.3 and 137.1 (B3). Relative to B3, coregionalization improves them by 8.6 on both criteria (B4). Because the covariance and geography dimensions are crossed, B2 and B3 are not stages of one sequential ladder.

The diagnostics in Section E retain their provenance. The binomial-versus-beta-binomial comparison is a likelihood check, the residual quantile comparison is a calibration check, and the pre-period BYM2 comparison is a sparse, model-based description rather than an identification test. No diagnostic is treated as evidence of actor intent or as an independent replication of the primary correlation.

D.6. Dynamic-panel bias and the propagation kernel

Fixed-effects estimation with lagged outcomes can produce bias of order $1/T$ (Nickell, 1981). In a dynamic spatial panel, bias can also affect the spatial lag (Lee & Yu, 2010). As Section 5.1 explains, the multiplier $\mathcal{M}$ is convex in $\rho_1 + \rho_2 + \lambda$ within the stable region. Changes in this kernel therefore affect both the response magnitude and its division between own-province persistence and network amplification. Examining only the temporal terms would leave the spatial contribution unassessed.

A split-panel jackknife assesses sensitivity to the fitted kernel. The specification is re-estimated on the two chronological halves, retaining predetermined lags. Twice the full-sample estimate minus the mean of the halves is designed to remove the leading $1/T$ term under the conditions of Dhaene and Jochmans (2015). This exercise adjusts only the three dynamic terms, because neither shock basis moves in both halves. Table 14 reports the resulting kernel and downstream quantities. The propagation sum rises from 0.554 to 0.764, the multiplier from 2.242 to 4.243, and the network share of the steady-state response from 31.0 to 49.1 percent. Holding the direct loading fixed, the steady-state diffusion contribution at the common 2.666-point reference would rise from 1.475 to 2.79 log points.

Every reported quantity in Section 6 retains the original fitted kernel. Halves of fourteen and fifteen years estimate the dynamics noisily and do not identify both shock bases. Moreover, dependence across spatial units limits direct application of the jackknife's formal guarantees. The exercise yields greater propagation in this implementation; it does not establish the direction of the original estimates' bias or bound the true magnitudes in Section 6.2.

Table 14: Split-panel jackknife for the M4 propagation kernel, 1962–1990. Network share is the off-diagonal proportion of the steady-state response.

	Fitted	First half	Second half	Corrected
ρ_1	0.173	0.025	0.156	0.255
ρ_2	0.164	0.037	0.130	0.244
λ	0.218	0.144	0.196	0.266
Propagation sum $\rho_1 + \rho_2 + \lambda$	0.554	—	—	0.764
Long-run multiplier $\mathcal{M}$	2.242	—	—	4.243
Own-province component	1.546	—	—	2.158
Network component	0.696	—	—	2.086
Network share of $\mathcal{M}$	0.310	—	—	0.491

E. Additional Results

This appendix collects supplemental descriptive analyses and model checks. It reports the persistence of recorded naming, the detailed political geography, the hotspot classification under false-discovery-rate control (Section 2.3), the district political-control comparison (Section 2.4), the party-specific moderation battery for the coup arm (Section E.4), the prior predictive calibration of the TİP moderation block, the province dynamic specification and residual diagnostics behind M4 (Section 5.1), the Bayesian sensitivity and qualified model-projection checks referred to in Section 6.3, the shock contrasts across the TİP distribution (Section E.8), the model-projected cumulative totals (Section E.9), the martyr-satellite replication (Section E.10), and the complete prediction checklist (Section E.11).

E.1. Persistence of recorded Deniz naming

Persistence of recorded Deniz naming reflects both prevalence and the number of births observed. Figure 16 shades each district by the number of years between 1970 and 1990 with at least one Deniz registration. Persistence concentrates along the southern coast and across the central and northeastern interior, but it partly reflects exposure: a large district registering Deniz at the national rate rarely misses a year, whereas a small district at the same rate misses many.

The map also understates the practice in metropolitan areas. Some metropolitan offices recorded a newborn's province but left the district field empty, and Deniz is disproportionately urban. The resulting attrition thins observed metropolitan district counts. Section 4.2 measures the coverage problem, while Section B documents the recording practice and explains why no defensible allocation rule can restore the missing district labels.

E.2. Political concordance after false-discovery-rate control

The classifications use local Moran's I on empirical Bayes-smoothed district rates under queen contiguity, with rank-based conditional-permutation p-values from 9,999 draws. The adjusted map applies Benjamini–Hochberg false-discovery-rate control at five percent across 636 district tests.

At the nominal threshold, 61 of 75 Deniz hotspots lie at or above the median TİP vote share, compared with 42.4 percent of the 85 coldspots.

False-discovery-rate control reduces the nominal level classification from 75 to 22 hotspots and from 85 to 24 coldspots. It strengthens rather than weakens the political concordance: 20 of the 22 surviving hotspots lie above the district median of the 1969 TİP vote share, compared with 7 of the 24 surviving coldspots. The same procedure applied to the post-1971 acceleration surface leaves none of its 25 nominal clusters, so the robust local clustering belongs to naming prevalence rather than to the subsequent response.

Table 15: Shares of FDR-robust hotspots and coldspots above each party's district vote-share median.

1969 vote share	Hotspots (%)	Coldspots (%)
TİP	90.9	29.2
Birlik Partisi	86.4	12.5
CHP	59.1	37.5
MHP	50.0	45.8
Adalet Partisi	36.4	41.7

The FDR-robust classification forms four geographic blocs rather than a simple divide between left and right. Tunceli holds 6 of the 20 hotspots above the TİP median and anchors an eastern concentration of nine with Erzincan, Malatya, and Sivas. İstanbul holds five, the İzmir coast three, and a Kırşehir–Nevşehir Cappadocian cluster three. The negative pole describes districts with lower left support. The electoral data do not measure hostility to the left or support for either intervention.

The comparison limits a TİP-specific interpretation. Birlik Partisi nearly matches the

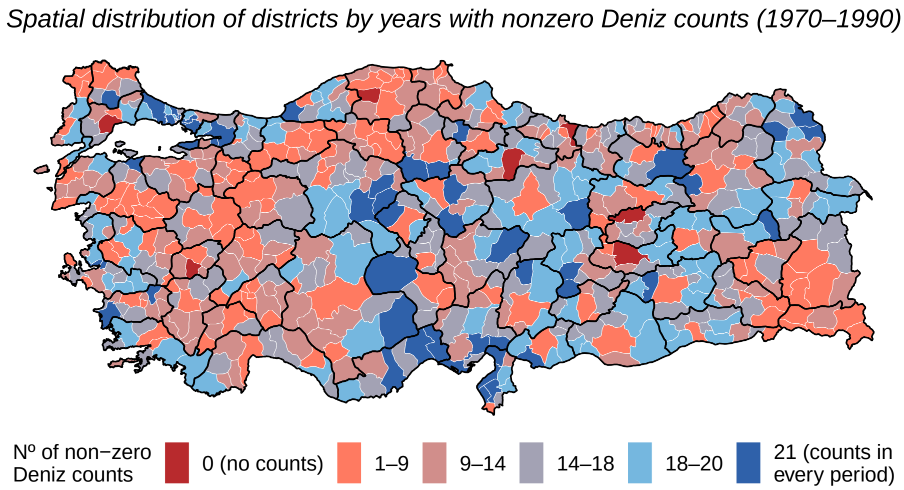

Figure 16: Years with at least one Deniz registration by district, 1970–1990. Persistence reflects birth volume as well as naming prevalence.

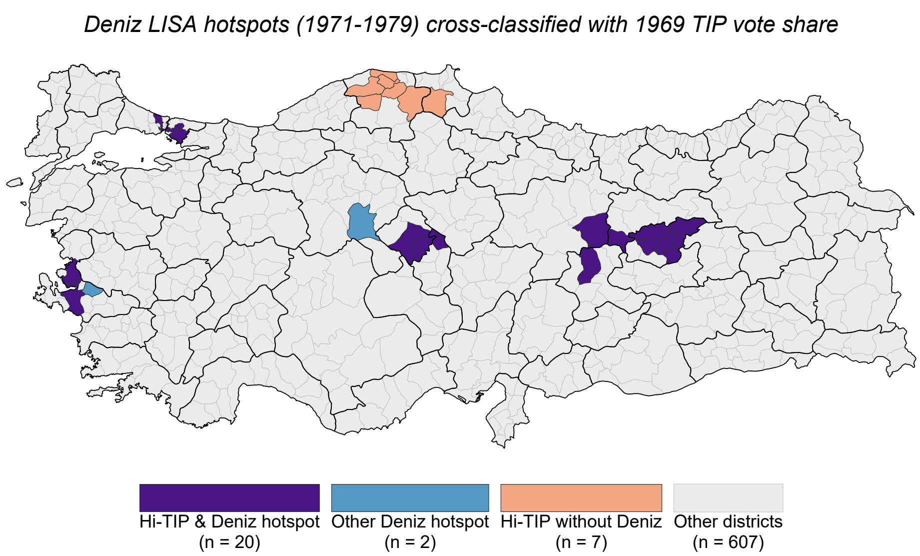

Figure 17: Deniz prevalence clusters, 1971–1979, by 1969 TİP support after five-percent false-discovery-rate control. Province borders are black.

TİP hotspot concordance and separates the coldspot pole more sharply. The level map therefore captures a broader political geography rather than a mechanism unique to the party used as the moderator.

TİP support is itself spatially clustered. Across the 636-district analysis geography under queen contiguity, its 1969 vote share has a global Moran's I of 0.399, with a vanishing p-value. Because this measure is a time-invariant cross-section, the statistic describes the baseline political map rather than a response in either regime window. It motivates a district spatial analysis but does not uniquely privilege TİP, since the other 1969 party shares are also clustered.[4]

E.3. The diffusion moderator against the center-right control

The center-right control enters the prespecified battery at district scale only; the province refits of Section E.4 are exploratory. On the diffusion dose, neither moderator resolves: the center-right interaction is $+0.0018$ $[-0.0038, +0.0074]$, $p = 0.512$, against TİP's $+0.0019$ $[-0.024, +0.028]$, $p = 0.883$. The contrast lies on the kernel channel. The center-right interaction with the neighbors' lagged rate is -0.0037 $[-0.0051, -0.0023]$, $p = 6.1 \times 10^{-6}$, against TİP's $+0.0129$ $[+0.0031, +0.0227]$, $p = 0.012$; jointly, the two remain -0.0031 $[-0.0047, -0.0015]$ and $+0.0098$ $[+0.0005, +0.0192]$, respectively. The control is better resolved in that joint fit, with $|t| = 4.02$ against 2.15. Model-free Spearman gradients point in the same directions, $+0.075$ ($p = 0.06$) for TİP and -0.190 ($p = 1.5 \times 10^{-6}$) for the center-right share; the producer reports no intervals for these gradients. TİP is also less clustered, with Moran's $I = 0.399$ against 0.545. These contrasts relate modeled propagation to a broader political geography without identifying a mechanism specific to left support. Section E.4 extends the comparison to the other 1969 parties and examines the coup association.

[4] Global Moran's I is 0.616 for the Alevi Birlik Partisi, 0.545 for the Adalet Partisi, 0.302 for the CHP, and 0.263 for the MHP, with vanishing p-values. Accordingly, Section 6.3 uses these vote shares in alternative-party comparisons.

E.4. Party-specific moderation of the coup response

A moderation claim about the left is only as strong as its behavior beside the other parties, so the coup arm receives the same symmetric treatment the diffusion channel received above. At province scale, exploratory refits of M4 swap the moderator while holding the specification fixed: TİP enters in percentage points, the center-right AP share on the same scaling, and Birlik Partisi as centered ballot presence, because the Alevi party stood in 298 of the 636 districts and its share is zero elsewhere. On the coup shift, the TİP interaction is -0.046 $[-0.088, -0.004]$ and the BP presence interaction -0.104 $[-0.169, -0.039]$, while the AP interaction is $+0.001$ $[-0.006, +0.008]$. Fitted jointly, the TİP and BP terms remain resolved at -0.045 $[-0.081, -0.009]$ and -0.073 $[-0.131, -0.015]$ and the AP term remains null; the joint event blocks reject for TİP ($p = 0.0007$), AP ($p = 0.006$), and BP ($p = 0.050$), and the AP and BP additions together reject at $p = 1.4 \times 10^{-5}$ (Figure 18).

At district scale, the log-linear battery swaps all five 1969 moderators through an identical six-interaction form on the same sample, each in percentage points centered at its own top-quartile mean. No coup-shift interaction resolves individually at this scale: TİP is -0.004 $[-0.023, +0.014]$, BP -0.012 $[-0.025, +0.002]$, AP -0.001, CHP $+0.001$, and MHP -0.003, while the prespecified Bayesian moderator retains its resolved slope, -0.020 $[-0.036, -0.004]$ (Section 6.3). BP also carries the only resolved diffusion interaction in the battery, $+0.020$ $[+0.005, +0.035]$. Because the battery holds the form identical across parties, its BP rows carry no ballot-presence indicator and therefore partly measure ballot access rather than support intensity.

Together, the reported comparisons are consistent with a contraction associated with left and Alevi political geography, but the evidence depends on scale and specification. The district log-linear coup interactions remain unresolved, and the Bayesian TİP association does not establish a mechanism unique to that party. These exploratory refits qualify the moderator's interpretation; they do not replace the prespecified TİP specification, and they carry the same agency caveat as every geographic result in Section 6.3.

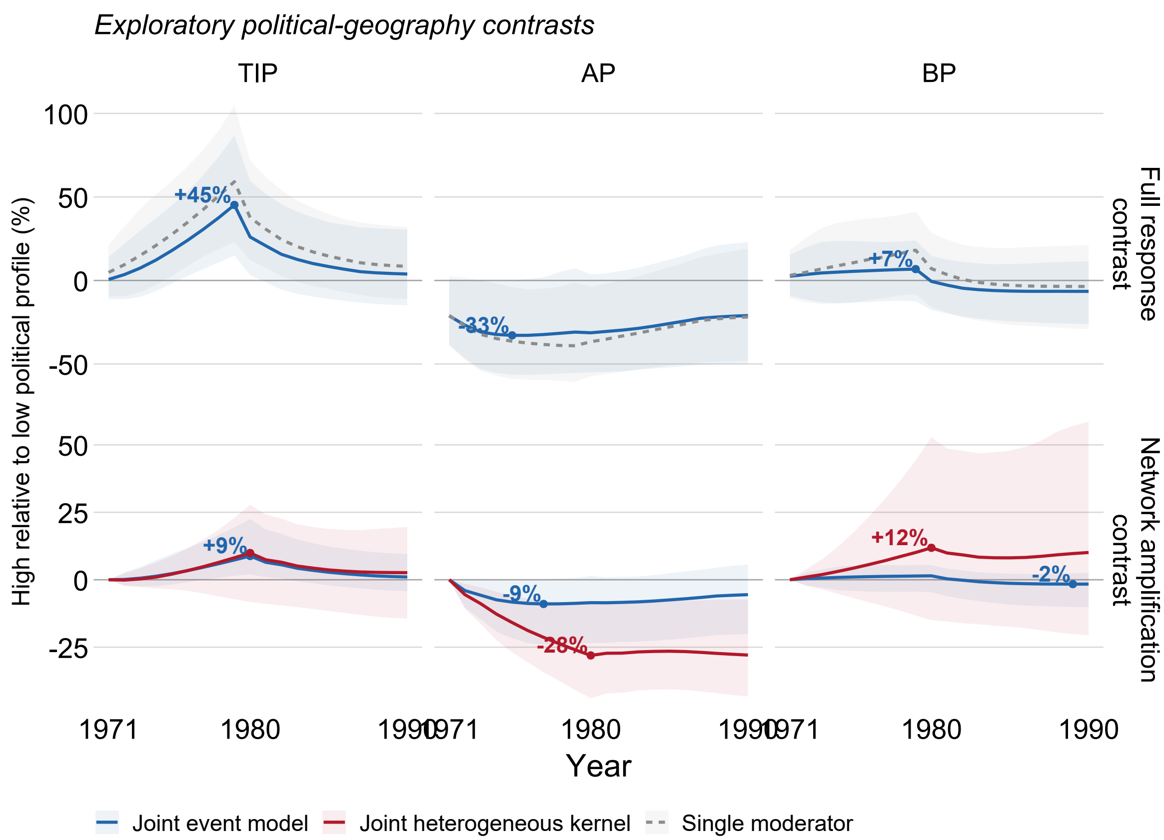

Figure 18: Exploratory province political contrasts, high versus low moderator values in percent. Rows show full responses and network amplification under the three specifications; bands are 95 percent coefficient-draw intervals.

E.5. Prior predictive calibration of TİP moderation

The TİP prior is calibrated on the outcome scale rather than selected from a generic default. The check takes 100,000 joint draws of all six moderation coefficients, evaluates them over the realized district-year vote shares and regime states, and holds the remaining linear predictor at the logit of the median observed district naming rate. Under the coefficient-scale-equivariant standard deviation of $2.5/3.001 = 0.833$ per percentage point, 15.8 percent of draws over the empirical design and 41.5 percent at the highest TİP district exceed the maximum observed naming rate. The prespecified standard deviation of 0.10 reduces those shares to 0.03 and 4.8 percent, respectively. It therefore limits implausible outcome predictions without fixing the sign of any moderation coefficient.

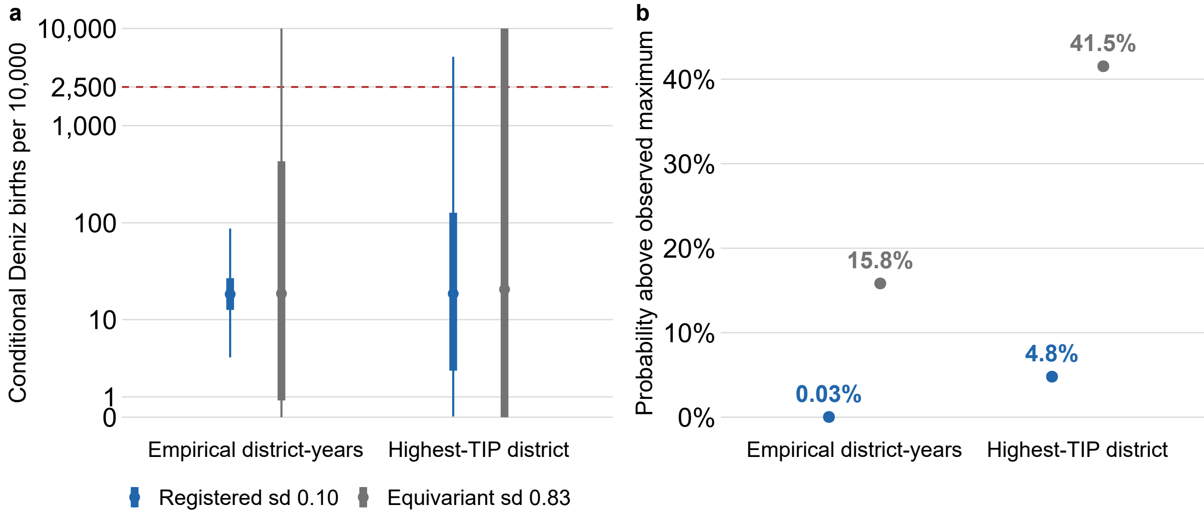

Figure 19: TİP prior predictions: (a) median, interquartile range, and 95 percent interval on a pseudo-log rate axis; (b) probability of exceeding the maximum observed district rate, dashed in (a).

E.6. Province dynamic specification and residual diagnostics

The temporal battery motivates the prespecified second-order recursion before M4 is fitted. The median partial autocorrelation remains positive through lag two, while province-specific AR(2) filtering reduces the Ljung–Box rejection share from 34.3 to 4.5 percent. This does not imply that one common AR(2) kernel removes every province's serial dependence; it explains why the fitted model carries two lags rather than one. The ACF and PACF summaries report province medians, interquartile ranges, and pooled means; the Ljung–Box comparison uses province-specific AR(1) and AR(2) filters.

The spatial battery distinguishes geographic association from national temporal persistence. Mean Moran's I falls from 0.163 in the raw rate to 0.041 after province-specific AR(1) filtering, but contemporaneous neighbor correlation remains 0.246 against a permuted baseline of 0.037. Similar excess correlation at temporal offsets one and two does not establish direction. The disjoint second-order-neighbor correlation is about half the first-order value, supporting the prespecified first-order graph as a parsimonious representation.

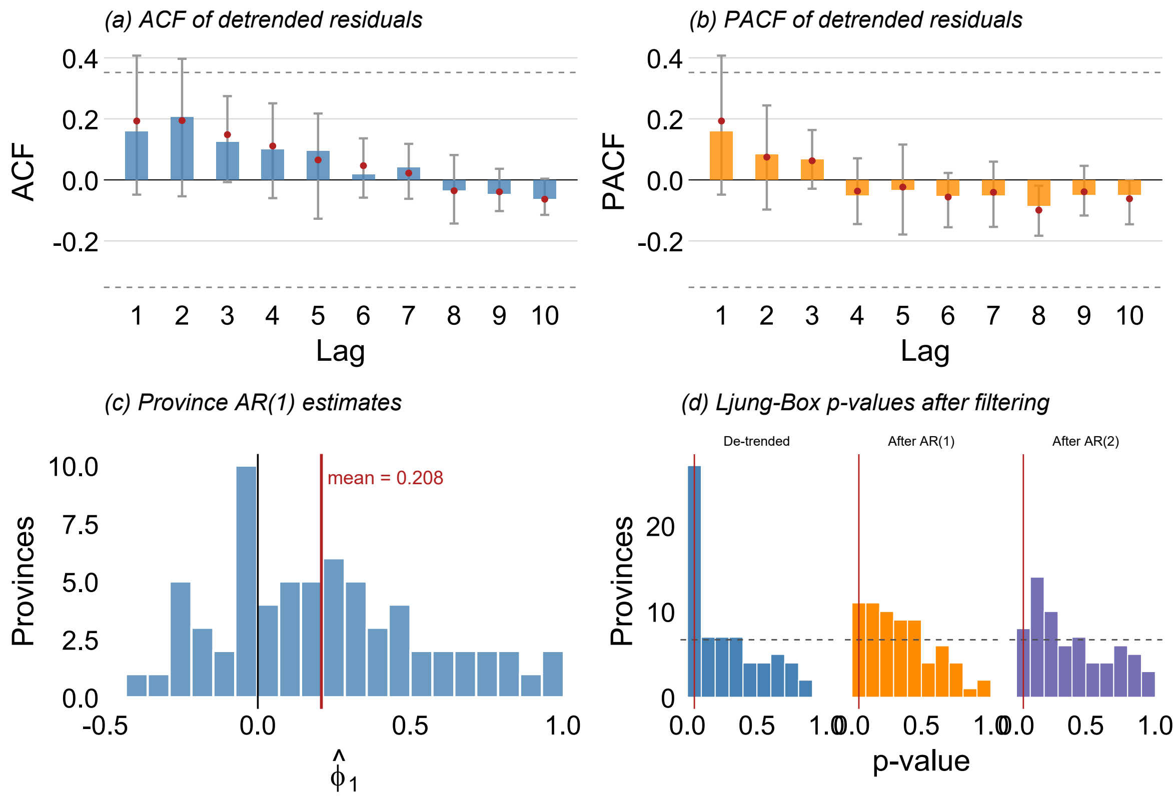

Figure 20: Province temporal diagnostics after removing fixed effects and the secular trend: (a,b) ACF and PACF; (c) AR(1) estimates; (d) Ljung–Box p-values before and after AR filtering. Dashed counts in (d) show the uniform-null reference.

The localized impulse diagnostic in Figure 22 asks whether the fitted M4 transition produces delayed, attenuating responses with graph distance. For interpretation, the localized input is scaled as a one-time doubling of the origin's naming rate. The average recipient response then peaks after one year among adjacent provinces, at about 3.13 percent; after two years at distance two, at about 0.168 percent; and after five years at distance three or more, at about 0.00261 percent. This ordered attenuation is consistent with a stable propagation kernel. The exercise applies an unobserved one-province input, so it diagnoses the model's internal mechanics rather than estimating localized causal spillovers.

The remaining diagnostic asks whether M4 leaves systematic temporal or spatial dependence in the residuals. Figure 23 reports both margins, the residual autocorrelation across provinces and the annual corrected Moran's I, with the permutation-flagged years marked.

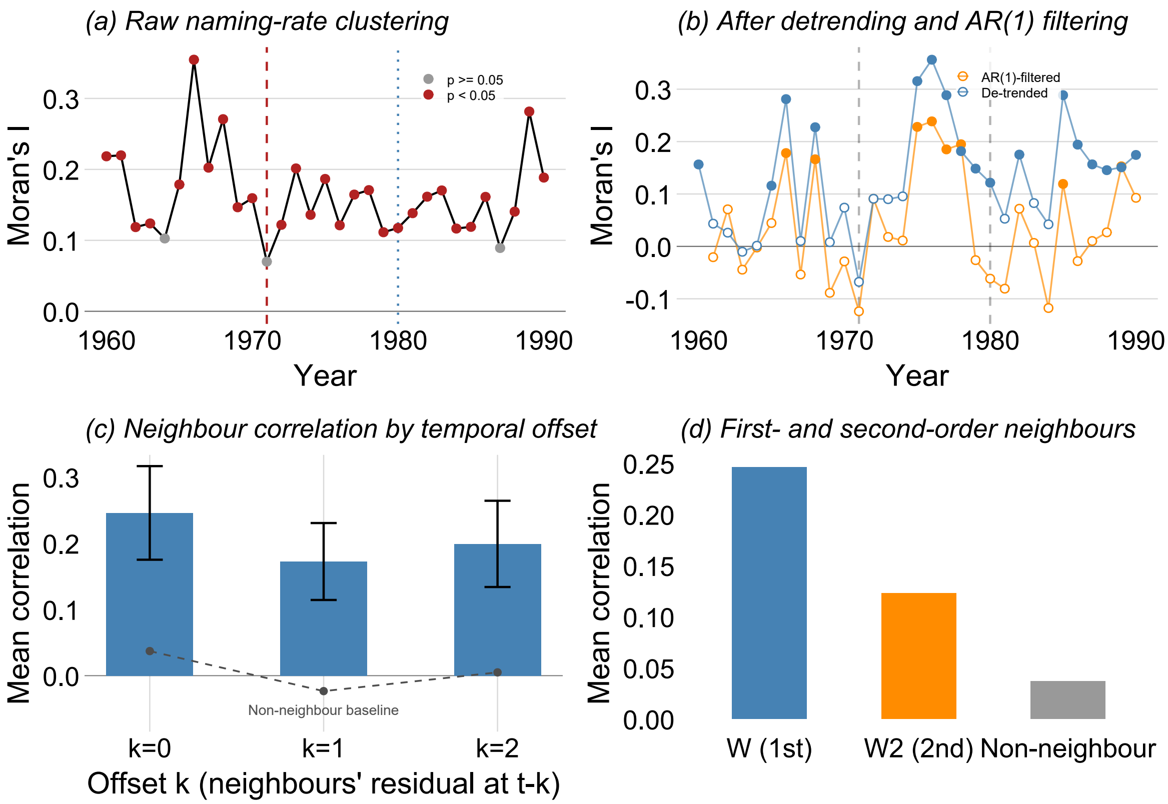

Figure 21: Province spatial diagnostics: (a) raw annual Moran's I; (b) detrended and AR(1)-filtered values; (c) neighbor correlations by temporal offset against label permutations; (d) correlations by graph distance.

E.7. Bayesian model diagnostics and sensitivity

The sensitivity evidence is deliberately separated by its inferential status. B2 prior and metropolitan checks retain direct posterior intervals. The two widening-prior B4 fits propagate joint hyperparameter uncertainty over the CCD design because wider priors could increase uncertainty without moving the mode. Both retain the registered interval width. Other B4 sensitivity fits use empirical-Bayes plug-ins and report points. The checks change one modeling choice at a time and do not provide independent causal validation of the primary correlation.

The model-projected paths in Figure 25 use the common 1962 to 1990 support; M4 loses the first two years to AR(2) initialization. The projection zeroes the regime steps and

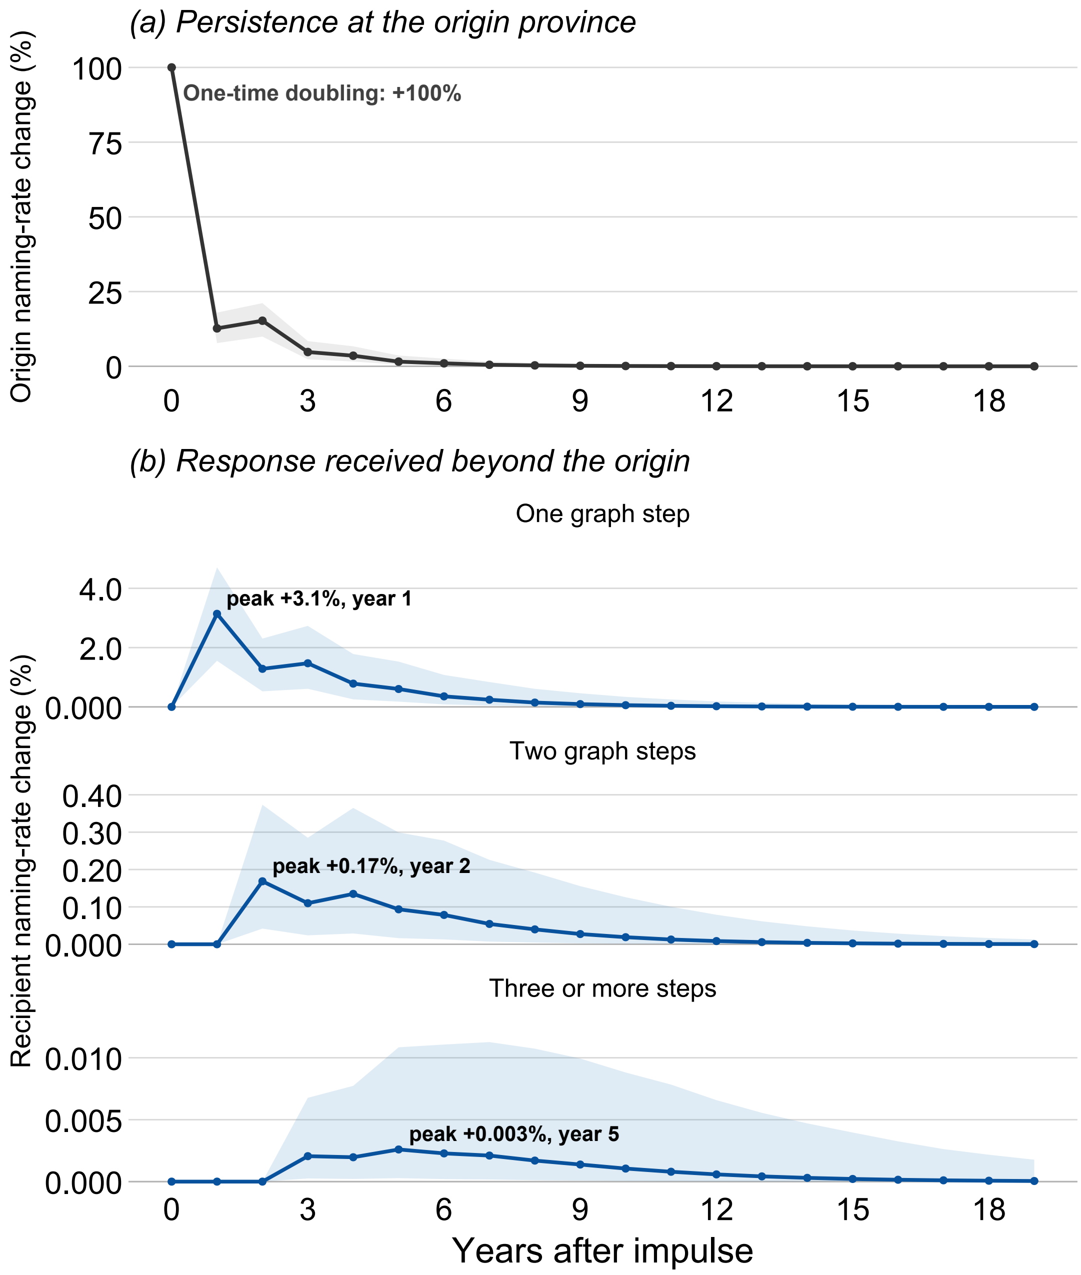

Figure 22: M4 response to a one-time doubling at each origin: (a) origin response; (b) recipients by graph distance, on separate scales. Lines average over origins; bands are 95 percent coefficient-draw intervals.

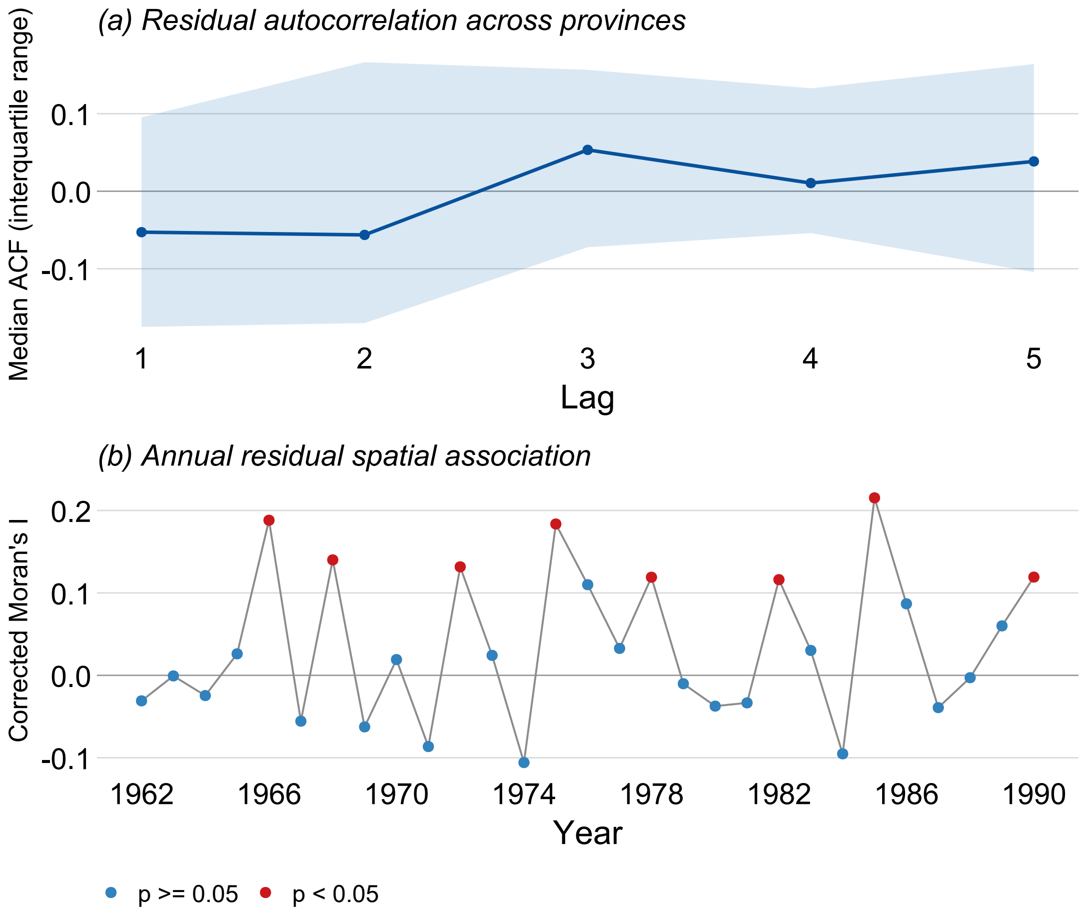

Figure 23: M4 residual diagnostics: (a) autocorrelation; (b) annual corrected Moran's I. Red points have permutation $p < 0.05$.

post-1971 doses and holds the secular component at its 1970 plateau. These conditional paths do not provide an independent counterfactual test.

E.8. Shock contrasts across the TİP distribution

Table 17 evaluates the fitted M4 arms over the 1969 TİP distribution, from the tenth percentile to the ninety-fifth, without refitting. The immediate coup coefficient is negative with an interval excluding zero at every grid point. The post-coup coefficient sum has an interval containing zero at the tenth percentile, is -0.047 at the median, and has an interval below zero only from the top-quartile reference upward. As Section 6.2 explains, this sum combines the direct coup shift and full-dose recovery contribution before propa-

Table 16: Sensitivity checks for diffusion–coup coupling. EB denotes empirical-Bayes plug-in estimates.

Check	Result	Status
Wishart prior (B2)	-0.523 under 4 and 8 degrees of freedom; both intervals exclude zero	Direct posterior sensitivity
Coregionalized field and loading priors (B4)	Tightening either prior gives -0.540 and -0.528; widening either gives -0.520 $[-0.744, -0.263]$ and -0.520 $[-0.743, -0.264]$	Tightening: EB plug-in. Widening: full posterior, because a wider prior is the case that could inflate uncertainty
Beta-binomial likelihood (B4)	-0.555; dispersion 6.8×10^{-5} at the boundary	EB plug-in; no evidence of useful overdispersion
Residual quantiles	RQR Kolmogorov–Smirnov $p = 0.582$ (binomial) and $p = 0.072$ (beta-binomial)	Calibration comparison
Metropolitan exclusion (B2)	-0.448 $[-0.619, -0.247]$	Direct posterior sensitivity
Reverse structural order (B4)	-0.507	EB plug-in sensitivity
Restricted spatial regression (B4)	-0.525	EB plug-in; not a bound on TİP moderation
Predictive fit and residual dependence	B4 versus B2 LGOCV: -1.1599 versus -1.1663; binomial mean annual residual Moran's $I = 0.0034$	Predictive and residual diagnostic
Pre-period response geography	ΔWAIC $= -0.36$ for BYM2 versus iid	Sparse, modeled diagnostic; not an identification test
Graph sensitivity	No current alternative-graph refit is available; the queen graph and ferry links remain assumptions	Limitation reported, not an invariance check

gation; it is not the remaining naming-rate gap. The common 2.666-point reference is the district mean. The 4.850-point reference is the birth-weighted mean of the top quartile of provinces. The table orders columns by vote share because the district mean and province percentiles use different weights.

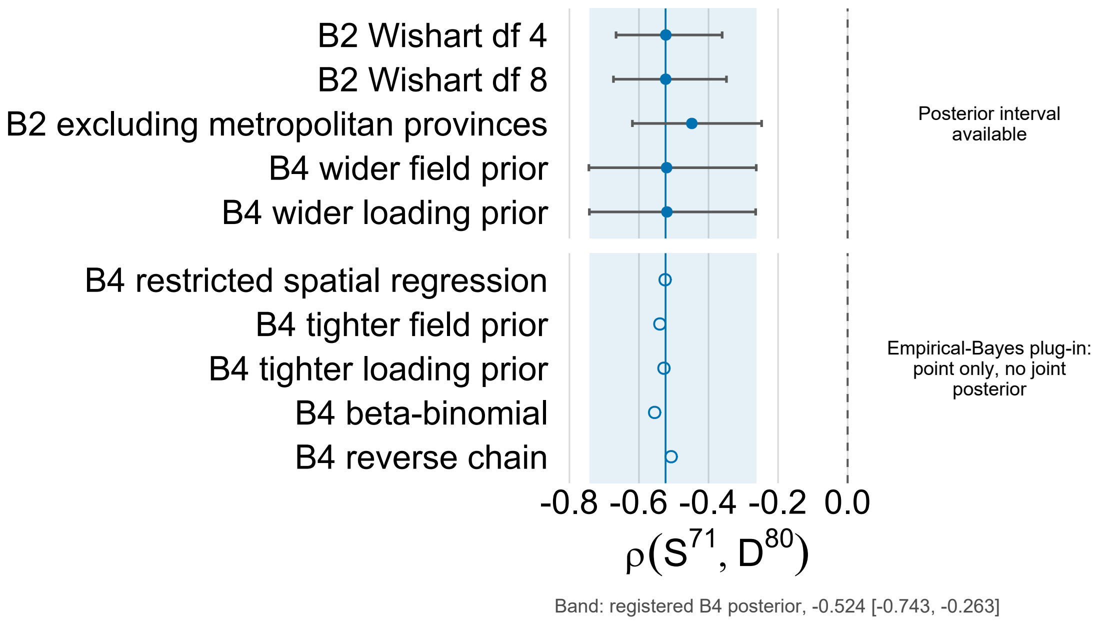

Figure 24: Diffusion–coup coupling across sensitivity fits. Bars show 95 percent credible intervals; isolated points are empirical-Bayes plug-ins. Shading and the solid line mark the registered B4 interval and point estimate.

Table 17: Province M4 shock contrasts by 1969 TİP support (percentage points). Driscoll–Kraay standard errors in parentheses.

Contrast TIP (pp)	p10 1.20	p25 1.68	median 2.00	p75 2.58	mean 2.67	Q4 ref. 4.85	p90 4.91	p95 6.23
D^{71} impact (γ_1)	0.086	0.092	0.096	0.104*	0.105*	0.133**	0.134**	0.151**
	(0.074)	(0.067)	(0.063)	(0.057)	(0.056)	(0.055)	(0.055)	(0.070)
S^{71} diffusion (δ_1)	0.580***	0.606***	0.623***	0.653***	0.658***	0.774***	0.777***	0.847***
	(0.112)	(0.115)	(0.118)	(0.124)	(0.125)	(0.164)	(0.165)	(0.195)
1971-block coefficient sum ($\gamma_1 + \delta_1$)	0.666***	0.697***	0.719***	0.757***	0.763***	0.907***	0.911***	0.998***
	(0.131)	(0.135)	(0.138)	(0.144)	(0.145)	(0.173)	(0.174)	(0.194)
D^{80} impact (γ_2)	−0.221***	−0.243***	−0.258***	−0.285***	−0.289***	−0.389***	−0.392***	−0.453***
	(0.038)	(0.036)	(0.037)	(0.040)	(0.041)	(0.073)	(0.075)	(0.099)
S^{80} recovery (δ_2)	0.220***	0.215***	0.211***	0.204***	0.203***	0.179***	0.178***	0.163***
	(0.054)	(0.049)	(0.047)	(0.042)	(0.041)	(0.035)	(0.035)	(0.042)
Post-coup coefficient sum ($\gamma_2 + \delta_2$)	−0.001	−0.029	−0.047	−0.080	−0.085	−0.210***	−0.214***	−0.289***
	(0.055)	(0.053)	(0.052)	(0.051)	(0.051)	(0.064)	(0.065)	(0.080)

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p \leq 0.01$, two-sided on 28 degrees of freedom.

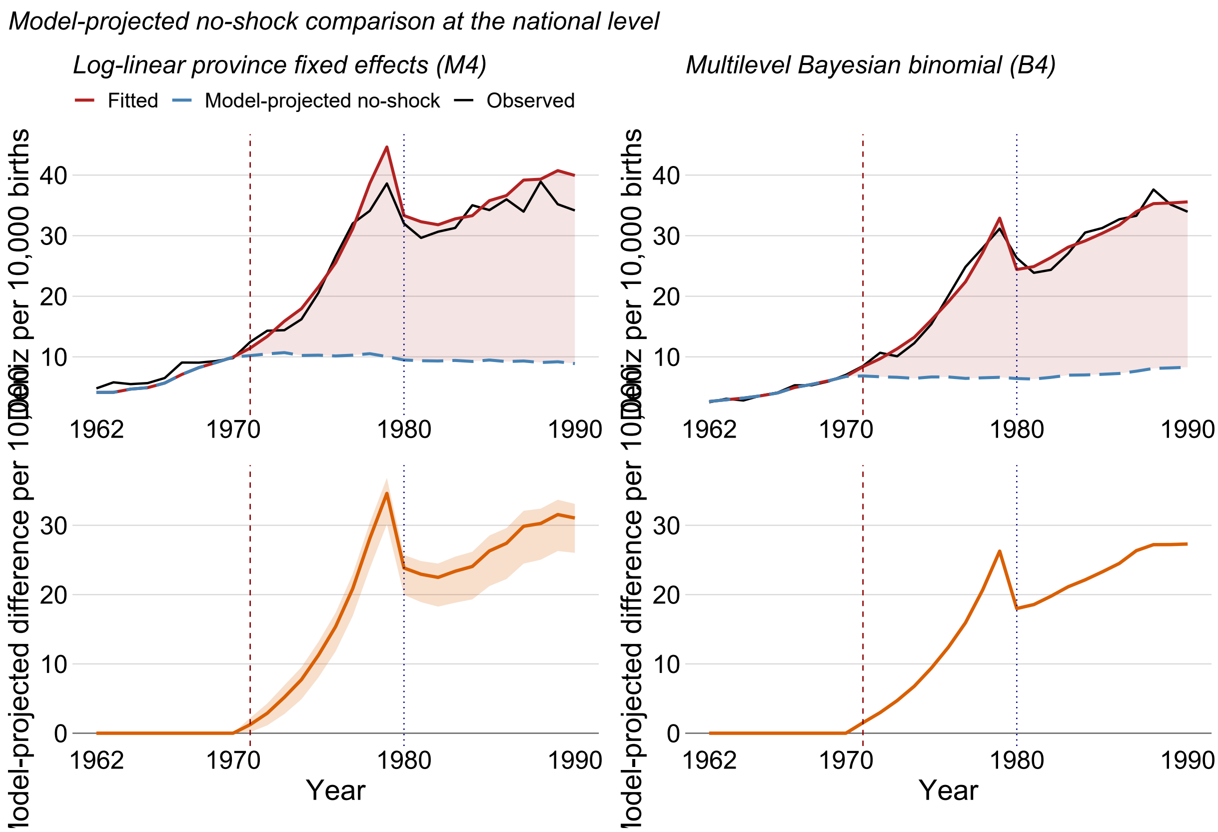

Figure 25: Observed, fitted, and no-shock national Deniz rates per 10,000 births, 1962–1990, with fitted-minus-projected differences below. B4 uses a posterior-mean plug-in path.

E.9. Model-projected cumulative totals

The cohort-benchmarked count of Section 6.1.2 is the headline cumulative quantity because its benchmark is realized rather than projected. The two model-based totals answer the companion question of the fitted difference associated with the shocks against each model's no-shock projection. They sit on different coverage bases, so none of the three numbers bounds another. The province linear model, on the 98 percent province-resolved base, attributes 44,115 births over 1971 to 1990, with a 95 percent interval from 36,047 to 48,022 under 5,000 coefficient draws. The district Bayesian model, on the 57 percent district-resolved base, attributes 26,130 births, with a 95 percent interval from 25,192 to 26,992 over joint posterior summaries retained at the CCD configurations, splitting into 6,871 $[6{,}530, 7{,}210]$ during the diffusion window and 19,259 $[18{,}566, 19{,}886]$ after the coup.

The posterior-mean plug-in path shown in Figure 25 accumulates to 25,929, 0.8 percent below the joint summary. Roughly three quarters of the modeled district excess accrues after the coup, consistent with continued elevation above the fitted no-shock path despite contraction.

Three caveats govern the comparison. The model totals inherit the pre-trend extrapolation of assumption A1, which the cohort benchmark avoids. The district total sits on the resolved sample, so changing ascertainment and composition complicate its interpretation and comparison with the other accounts. Improving ascertainment alone does not establish a direction of bias. Finally, the totals are not competing estimates of one number: the province and cohort quantities count against different benchmarks on the national and province-resolved bases, while the district quantity counts only resolved births. No model-based total bounds the cohort-benchmarked estimand in either direction, and none is a causal magnitude.

E.10. The martyr satellites

The martyr family supplies a staggered replication that the frozen pool cannot evaluate exactly. The pool was matched to Deniz's pre-1971 trajectory, so it is not an exchangeable reference for the other names and no exact p-value is quoted; the ranks below are descriptive positions in the same 41-name frame. Mahir Çayan and Ulaş Bardakçı were killed in March and February 1972, and their names are keyed to 1972. Devrim, the word for revolution, carries no single martyr event and is shown under the 1971 keying (Table 18).

Table 18: Martyr-name province estimates with Driscoll–Kraay standard errors. Ranks within the Deniz-matched pool are descriptive, without exact p-values.

	Keying	S dose	D^{80}	Rank S	Rank D^{80}	Joint rank	D^2
Mahir	1972	0.499 (0.108)	-0.289 (0.060)	3	2	1	12.40
Ulaş	1972	1.008 (0.183)	-0.622 (0.188)	1	1	1	57.26
Devrim	1971	-0.158 (0.369)	-0.661 (0.288)	33	1	1	40.69

The satellites extend the comparative pattern while qualifying its interpretation. Every name in this group, Deniz included, holds a first- or second-position coup shift, and the joint distance ranks first for all three satellites. Devrim, whose revolutionary meaning predates the 1971 events, has no resolved positive diffusion response. Its estimate does not establish the absence of a response, but distinguishes the reported evidence for expansion across these names. The monthly intervals are compatible with several events rather than precisely aligned with individual deaths (Figure 26). Mahir's first point estimate is June 1971, with an interval from March 1971 to August 1973 that includes his death in March 1972. Ulaş's is December 1971, with an interval from November 1971 to April 1972 that includes the February 1972 death. Both point estimates precede those deaths. Deniz's first interval excludes the sentence and execution (Section 2.1), while Devrim's first break falls in 1967 without a corresponding martyr event.

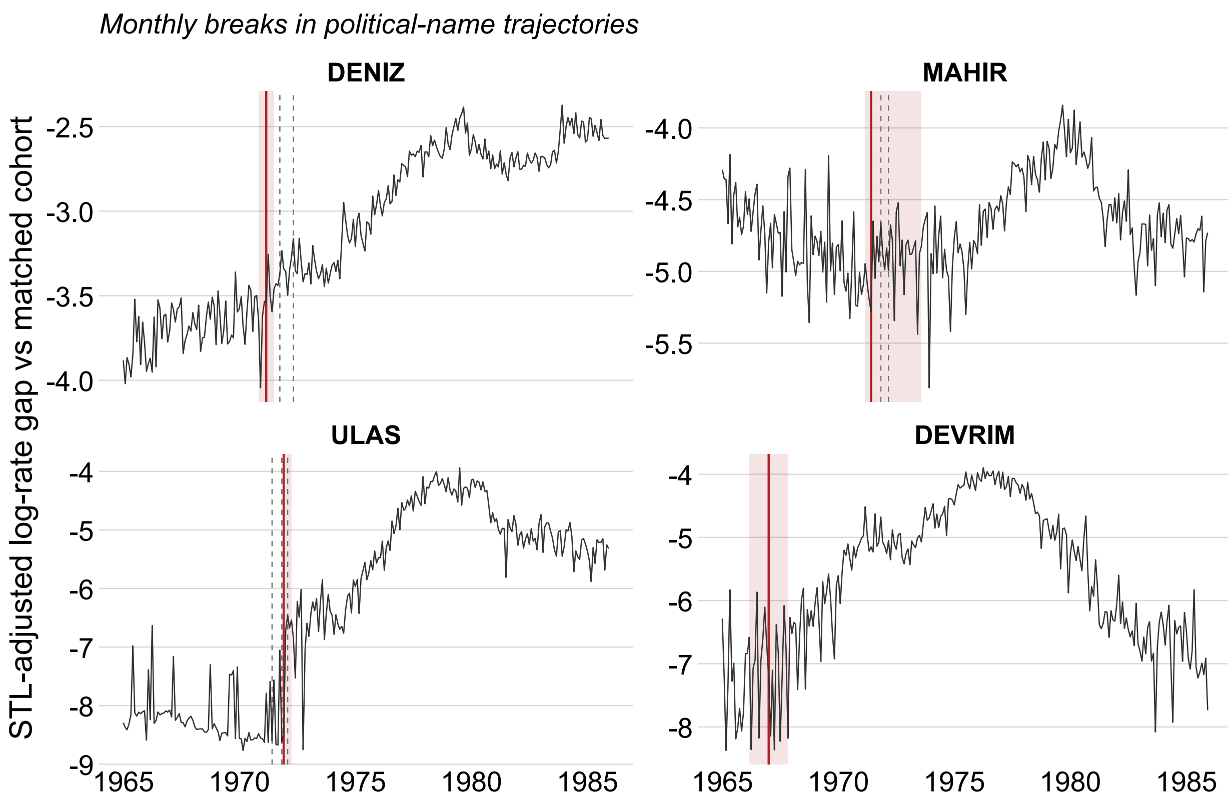

Figure 26: STL-adjusted log-rate gaps between each focal name and the matched cohort. Solid lines and bands mark the first estimated break and its uncertainty; dashed lines mark political events.

E.11. Complete prediction checklist

The checklist distinguishes the four observable implications of Section 2.2 from related geographic and descriptive comparisons. Expansion, contraction and subsequent recovery, geographic spread, and comparative distinctiveness appear in that order. Geographic spread remains unresolved: the maps describe naming prevalence, and fitted propagation illustrates model dynamics without identifying transmission or expansion into new locations. The response association provides related geographic evidence; its allocation specifically to gradual diffusion depends on the response specification. The directional conjunction test evaluates expansion and contraction in the province model, while the joint distance assesses extremity within the matched pool. These results remain conditional on their respective designs and do not jointly identify a single causal mechanism.

The post-coup coefficient sum qualifies recovery by combining its direct full-dose contribution with the coup shift. It is not the remaining naming-rate gap after propagation.

The final row records a descriptive comparison with naming prevalence just before the coup. Its correlation with the estimated coup response is -0.054, without a posterior interval. This comparison does not establish that prevalence is unrelated to contraction.

Table 19: Predictions and related comparisons, with reporting locations and qualifications. P denotes province M4, B the district Bayesian analysis, and M the matched-name comparison.

Prediction	Evidence in	Result	Verdict
Expansion after 1971	Section 6.2 (P), Section 6.3 (B)	Direct full-dose loading $\delta_1 = 0.658$ (0.125); steady-state contribution 1.475; 1979 response $+311.4$ percent $[190.1, 427.1]$	Positive under the specified projection
Contraction at the coded 1980 boundary	Section 6.2 (P)	$\gamma_2 = -0.289$ (0.041), $p = 6.3 \times 10^{-8}$; intersection–union $p = 6.9 \times 10^{-6}$ with expansion	Holds under annual coding
Post-coup coefficient sum	Section 6.2 (P)	$\gamma_2 + \delta_2 = -0.085$ $[-0.190, +0.020]$ at the 2.666-point reference and -0.210 $[-0.342, -0.079]$ at the top-quartile reference	Interval includes zero at 2.666 points; lies below zero at the top-quartile reference
Geographic spread	Sections E and 6.2	Fitted propagation and descriptive prevalence maps do not directly test expansion into new locations or transmission	Unresolved
Comparative distinctiveness	Section 6.1.1 (M)	The prespecified two-coordinate D^2 ranks 1/41 in both specifications; the cohort count is descriptive and carries no exact rank	Holds conditionally within the frozen pool
Concentrated geography	Section 6.3 (B)	$\rho(S^{71}, D^{80}) = -0.525$ in B2 and -0.524 in B4; $\kappa_{21} = -0.280$ $[-0.428, -0.132]$	Negative association; allocation depends on response specification
Covariance pattern (secondary geographic claim)	Section 6.1.1 (M)	Three-correlation $D^2 = 5.92$, rank 4/41, $p = 0.098$	Distinctiveness not established
No coup-response momentum	Section 6.3 (B)	$\rho(D^{80}, S^{80}) = +0.008$ $[-0.243, +0.263]$; $\kappa_{32} = +0.186$ $[-0.608, +1.057]$	Unresolved
Prevalence rival	Section E.11 (B)	Eve-of-coup prevalence correlation -0.054; no posterior interval is produced	Descriptive comparison only

F. Permutation Record

This appendix preserves the design diagnostics, rank sensitivity, and the complete inference inventory behind Section 6.1.1. The frozen 41-name pool is evaluated with symmetric leave-one-out statistics for the model-based ranks. The natural-scale benchmark remains comparative accounting rather than another permutation statistic.

F.1. Pre-period comparability

The display assesses observed trajectory comparability, not balance on unmeasured composition or geography. Matching uses province-assigned records through 1969; 1970 is included only in the plot.

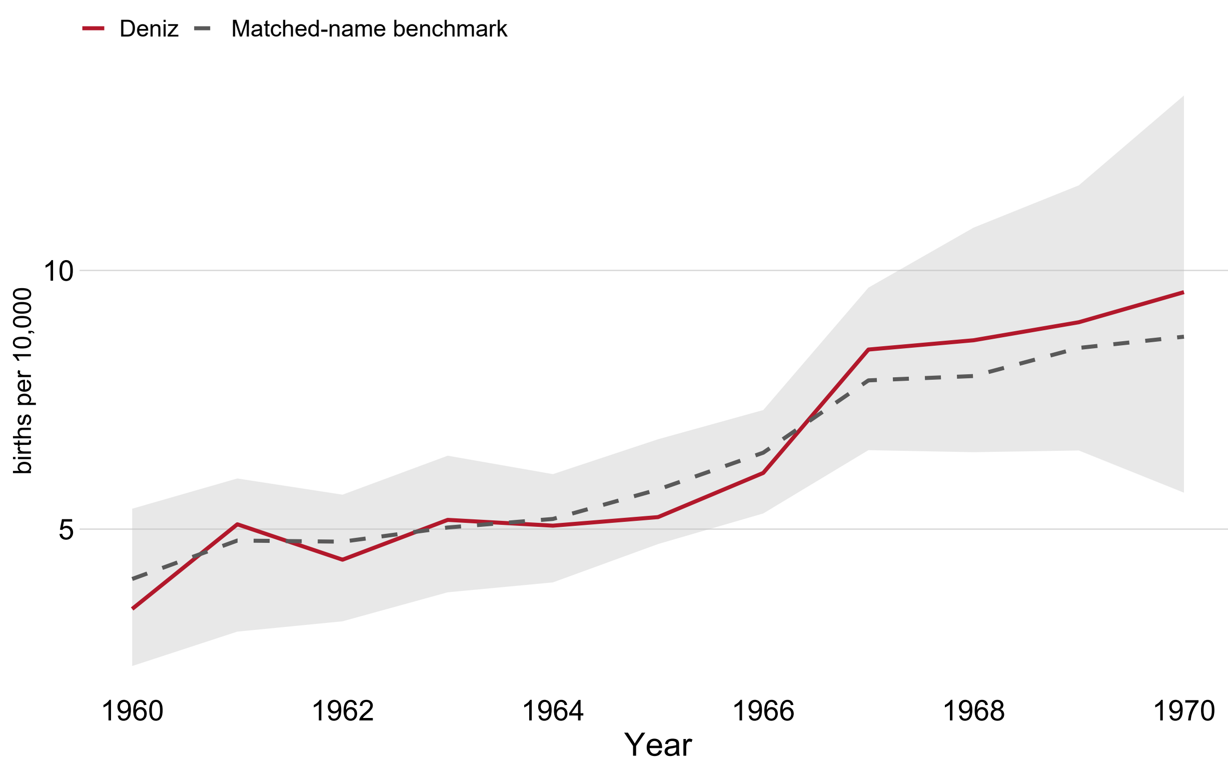

Figure 27: Deniz against the equal-weight mean and range of 40 comparison names, 1960–1970. Matching uses 1960–1969; 1970 is shown for context.

F.2. Coefficient uncertainty, record coverage, and geographic scale

The reported ranks treat fitted coefficients as fixed summaries. The approximate sensitivity recalculates them under name-specific marginal uncertainty and a shared within-model correlation template. Because incomplete district resolution is inseparable from the scale comparison, the integrated five-specification results and their limits are reported in Table 10, Figure 15, and Section B. This appendix retains the complete rank inventory and negative evidence below.

F.3. Complete rank inventory and negative evidence

The complete inventory separates primary tests from signed components, diagnostics, negative controls, and the restricted-pool sensitivity. The primary quantity throughout is the joint (S^{71}, D^{80}) magnitude that carries the political signal hypothesis; the coupling and covariance rows are transparency diagnostics of the secondary geographic-structure claim. All 41 B2 diffusion-to-coup correlations are negative under the adjacent event windows, so their sign is nondiagnostic. Deniz has the largest negative magnitude, but the broader three-correlation distance ranks 4 of 41 with $p = 0.098$. Restricting the pool to Deniz and 17 unisex comparison names yields rank 2 of 18 under M4 and rank 1 of 18 under B2. Their conditional probabilities are 0.111 and 0.056, respectively; both analyses have the same attainable floor, $1/18 = 0.056$.

Table 20: Complete matched-name rank inventory: primary tests, signed components, diagnostics, negative controls, and unisex-pool sensitivity.

Analysis	Quantity	Prediction	Deniz estimate	Rank / pool	Conditional p	Role
M4	$D^2(S^{71}, D^{80})$	Joint political signature	15.921	1 / 41	0.024	Primary
District M4	$D^2(S^{71}, D^{80})$	Joint political signature	20.190	1 / 41	0.024	Sensitivity
B2	$D^2(S^{71}, D^{80})$	Joint political signature	12.528	1 / 41	0.024	Primary
B2	$D^2(\rho_{12}, \rho_{13}, \rho_{23})$	Broader covariance pattern	5.922	4 / 41	0.098	Secondary
M4	S^{71}	Larger	+0.658	2 / 41	0.049	Component
M4	D^{80}	Smaller	-0.289	2 / 41	0.049	Component
M4	$S^{71} - D^{80}$	Larger	+0.947	1 / 41	0.024	Component
M4	TIP $\times$ S^{71}	Larger	+0.053	2 / 41	0.049	Diagnostic
M4	TIP $\times$ D^{71}	Larger	+0.013	17 / 41	0.415	Diagnostic
M4	TIP $\times$ D^{80}	Smaller	-0.046	1 / 41	0.024	Diagnostic
M4	TIP $\times$ S^{80}	Smaller	-0.011	16 / 41	0.390	Diagnostic
B2	$\rho(S^{71}, D^{80})$	Smaller	-0.522	1 / 41	0.024	Diagnostic
B2	$\rho(S^{71}, S^{80})$	Smaller	-0.298	5 / 41	0.122	Diagnostic
B2	S^{71}	Larger	+1.450	1 / 41	0.024	Component
B2	D^{80}	Smaller	-0.264	4 / 41	0.098	Component
B2	$S^{71} - D^{80}$	Larger	+1.714	1 / 41	0.024	Component
B2	TIP $\times$ S^{71}	Larger	+0.027	2 / 41	0.049	Diagnostic

Analysis	Quantity	Prediction	Deniz estimate	Rank / pool	Conditional p	Role
B2	$\text{TIP} \times D^{71}$	Larger	+0.024	5 / 41	0.122	Diagnostic
B2	$\text{TIP} \times D^{80}$	Smaller	-0.021	2 / 41	0.049	Diagnostic
B2	$\text{TIP} \times S^{80}$	Smaller	-0.030	7 / 41	0.171	Diagnostic
M4	$\text{TIP} \times S^{pre}$	No separation	+0.030	Not applicable	Not applicable	Negative control
B2	$\text{TIP} \times S^{pre}$	No separation	+0.009	Not applicable	Not applicable	Negative control
B2	$\rho(D^{80}, S^{80})$	No separation	+0.011	Not applicable	Not applicable	Negative control
M4	$D^2(S^{71}, D^{80})$	Joint political signature	11.706	2 / 18	0.111	Sensitivity
B2	$D^2(S^{71}, D^{80})$	Joint political signature	8.393	1 / 18	0.056	Sensitivity

An alternative equal-weight diagnostic based on national Jeffreys-adjusted log rates does not reproduce the primary result. Deniz's post-1971 RMSPE ranks 7 of 41 with $p = 0.171$. Its post-to-pre RMSPE ratio ranks first but is sensitive to the small pre-period denominator. This negative result shows that not every plausible trajectory score is confirmatory. Its full exhibit and construction remain with the replication materials.