When Rigid Blocs Crack: Elite-Coordinated Voter Switching in an Identity-Based Party System

Boris Gorelik *

Posted Date: 10 March 2026

doi: 10.20944/preprints202511.2148.v2

Keywords: bloc voting; electoral volatility; party loyalty; elite coordination; ecological inference; ultra-orthodox politics; voter transitions; identity politics

 doi:10.20944/preprints202511.2148.v2

Article

When Rigid Blocs Crack: Elite-Coordinated Voter Switching in an Identity-Based Party System

Boris Gorelik

Azrieli College of Engineering, Jerusalem, Israel; boris@gorelik.net

Abstract

When tightly-knit communities suddenly show electoral volatility, does it signal weakening group identity, or does it reveal something deeper? This question matters wherever centralized authority structures shape bloc voting. Conventional wisdom interprets such shifts as boundary erosion. This paper presents evidence for the opposite. Drawing on an extreme-case design, I exploit a natural experiment—Israel's 2019–2022 political deadlock—to track voter transitions within ultra-Orthodox communities, where ethnically distinct subgroups maintain near-total political separation despite shared religious practice. Using ecological inference on ballot-box data from five population centers across six elections (2019–2022), I find exceptionally high baseline party loyalty (90–95%), a dramatic disruption during the March 2020—March 2021 transition when switching surged to 12–19%, and a swift return to high loyalty within 13 months—though the shifted voters remained with their new parties. The synchronized switch of voting loyalty across geographically dispersed cities, occurring without residential mobility, is consistent with elite-mediated bloc realignment rather than emerging voter independence. Paradoxically, the capacity for mass switching may reflect stronger, not weaker, institutional control. These findings challenge how scholars of party–voter linkages interpret electoral volatility in identity-based voting blocs: apparent instability may reflect disciplined coordination, and what looks like boundary erosion may actually reveal institutional strength operating through collective action.

Keywords: bloc voting; electoral volatility; party loyalty; elite coordination; ecological inference; ultra-orthodox politics; voter transitions; identity politics

Introduction

"ושמרת לעשות ככל אשר יורוך" "And you shall do according to all that they instruct you" (Deuteronomy 17:10)

Across many democracies, communities exhibit cohesive voting patterns that resemble durable "voting blocs" anchored in strong group identities. Classic cleavage theories and the literature on ethnoreligious mobilization argue that such identities can cement stable partisan alignments over long periods (Lipset & Rokkan, 1967; Curiel & Zeedan, 2024), particularly when parties maintain strong organizational linkages to civil society that constrain voter availability to competitors (Martin, de Lange, and van der Brug 2022). From U.S. evangelicals (Campbell et al., 2011) to European regional party loyalists to African ethnic voting coalitions (Horowitz, 1985), these blocs appear remarkably stable over time. Yet a central question for political behavior remains: are these loyal voting blocs truly immobile, or can their rigid boundaries temporarily yield under stress? This question functions as a stress test of bloc discipline theories, exploiting a rare quasi-natural experiment in Israel that allows us to observe how deeply internalized political loyalties respond to sudden systemic shocks.

Israel's ultra-Orthodox (*Haredi*, plural *Haredim*) Jews provide an extreme-case test of this question (Seawright and Gerring 2008). They combine centralized rabbinic authority with dense organizational networks of yeshivas and synagogues. Repeated elections over a compressed period supply ballot-box data that make the Haredi sector uniquely suited to reveal mechanisms of elite-coordinated volatility that are theorized across many settings but rarely directly observable.

 doi:10.20944/preprints202511.2148.v2

The Haredi sector is widely portrayed as one of the most disciplined electoral blocs in Israel, with high turnout rates and reliably pro-clerical voting (Freedman 2020; Leon 2014; Malach and Cahaner 2025; Dödtmann 2022). Haredi political culture traditionally frames voting not as an individual right but as a collective religious duty dictated by spiritual authorities. As Haredi politician Tzuriel Krispel reflected after losing his reelection bid as mayor: "The Haredi public does not have the right to choose, only the right to vote" (Blum 2007). The expectation is that voters follow rabbinic directives rather than evaluate candidates independently. Yet perceived unity can mask underlying volatility. Simas and Lothamer (2025) find that even strong party identifiers may temporarily defect following intraparty conflict.

Israel's 2019–2022 political deadlock was a protracted crisis in which four national elections were held between April 2019 and March 2021 amid repeated failures to form a governing coalition. This crisis allows a peek into a theoretical puzzle applicable beyond Israel: can highly disciplined voting blocs harbor latent realignment pressures that surface under specific temporal and local conditions?

The Haredi sector is also internally divided along broad ethnic lines. Ashkenazi (European-origin) and Sephardi/Mizrahi (Middle Eastern/North African-origin) streams, though internally diverse, are each primarily represented by separate political parties: United Torah Judaism (UTJ) and Shas, respectively. Ashkenazi Haredim are represented by UTJ, whose electorate is almost entirely ultra-Orthodox Ashkenazi. In contrast, Sephardi/Mizrahi Haredim predominantly vote for Shas, though Shas's voter base extends well beyond the strictly ultra-Orthodox community. Both parties occupy a distinctive niche in Israel's multidimensional party space, combining right-wing positions on the Israeli–Palestinian conflict with left-wing economics (Zur and Bakker 2025). The vast majority of Shas's electoral support has come from non-ultra-Orthodox voters, reflecting the party's broad appeal (Keren-Kratz 2025). Leon (2014) characterizes Sephardi ultra-Orthodoxy as "strict ideology, liquid identity," suggesting more permeable boundaries than Ashkenazi streams. Malach and Cahaner (2025) find that UTJ behaves as a "sectarian party with dynamic fringes," drawing ~95% of its potential core support but subject to modest flows at the edges. This ethnic cleavage within a religiously unified population—which resists binary classification, as both parties fall on a continuum of ethnicization rather than being purely "ethnic" or "non-ethnic" (Hamza and Maeda 2025)—provides a unique opportunity to study how internal boundaries within cohesive blocs respond to electoral shocks. Analyzing transitions between these ethnically-aligned parties reveals whether ethnic-political boundaries can temporarily weaken during crises without triggering permanent realignment, a pattern with broad implications for understanding identity-based voting blocs worldwide.

The mechanism at work—centralized institutional authority channeling bloc votes through organizational networks—is not unique to the Haredi case. In Italy, Lanzara et al. (2024) show that the identity of the bishop in office explained significant variation in Christian Democracy vote share across dioceses, with episcopal replacements producing swings of up to three percentage points through parish-level networks. In the United States, pastoral endorsements and voter guides distributed through evangelical church networks shape electoral behavior in ways that parallel rabbinic directives (Campbell, Green, and Layman 2011). In sub-Saharan Africa, ethnic brokers and traditional chiefs coordinate bloc voting through patronage networks that span dispersed communities (Horowitz 1985). The Haredi case provides an unusually clean empirical window into this mechanism because repeated elections over a compressed period allow tracking of both disruption and recovery—a temporal resolution rarely available in other settings.

An anomaly in residential segregation data (Gorelik 2025) first suggested that such a disruption had occurred but failed to provide an explanation. This paper presents this empirical puzzle, develops a theoretical framework to interpret it, and then uses ecological inference to estimate voter transition matrices across successive elections. The analysis reveals a dramatic but temporary collapse in voter loyalty to Shas during the March 2020–March 2021 transition, with substantial city-level variation in disruption magnitude, followed by rapid recovery that suggests temporary stress

 doi:10.20944/preprints202511.2148.v2

fractures rather than permanent realignment. I argue that this pattern carries implications for any democracy featuring identity-based voting blocs.

The remainder of the paper proceeds as follows. The Empirical Puzzle section presents the anomaly that motivated this study. The Theoretical Framework section develops the interpretive lens. The Methods section describes the data, ecological inference model, and sample restrictions. The Results section presents transition estimates at the national and city levels. The Conclusions section develops the implications and situates the findings within the broader study of identity-based voting.

The Empirical Puzzle

This study was sparked by an anomaly discovered in my previous research on Haredi residential segregation (Gorelik 2025). That study documented persistent spatial separation between Ashkenazi and Sephardi Haredim across Israel's cities. However, Ashdod, a southern coastal city, exhibited a sudden, unexplained drop in residential segregation between the March 2020 and March 2021 elections, followed by a plateau (Figure 1). The short period between these elections excluded major demographic shifts or new housing developments as explanations. The anomaly was puzzling: if people did not physically move, what changed? One plausible answer: politics. Residential segregation changes slowly, but voting behavior can shift rapidly. If Sephardi voters temporarily switched from Shas to UTJ (or vice versa), this would create apparent "integration" in segregation indices, which rely on party votes as proxies for ethnicity, without any residential movement. While this proxy is theoretically grounded and performs well under stable conditions, the Ashdod episode suggests it is particularly fragile in highly disciplined populations where centralized authority can rapidly redirect voting behavior.

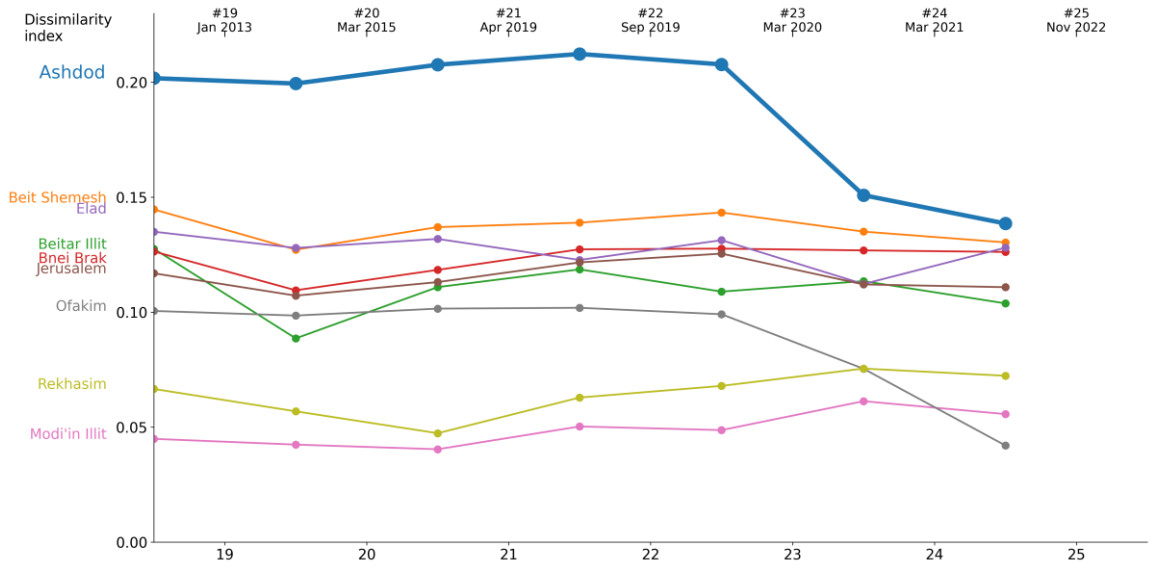

Figure 1. Ashdod's segregation anomaly was electoral, not demographic. The dissimilarity index tracks spatial segregation between Ashkenazi and Sephardi Haredim within each city. Ashdod's sharp dip between Knesset 23–24 (March 2020–March 2021) coincides with the voter transition disruption analyzed in this study. Reconstructed from Gorelik, 2025.

The puzzle deepens when we examine Shas voter loyalty across the full sequence of elections. As I will show later in this paper, Shas retained its voters at remarkably high rates—98%, 99%, 98%—across successive election pairs during 2019–2020. Then, in the March 2020 to March 2021 transition (Knesset 23→24), loyalty collapsed to approximately 74% at the national level, with even steeper drops in some cities (Ashdod fell to roughly 65%). Yet by the following transition (Knesset 24→25, March 2021 to November 2022), Shas loyalty rebounded to 97%. This dramatic 98→99→98→74→97

 doi:10.20944/preprints202511.2148.v2

trajectory—a sudden crack in an otherwise rigid pattern, followed by swift recovery—is not easily explained by gradual dealignment, demographic change, or shifting policy preferences. The synchronized timing across geographically dispersed cities further suggests a coordinated mechanism more than independent local dynamics.

Was the dramatic loyalty collapse evidence of a broader, under-the-radar pattern of elite-coordinated voter transitions within this ostensibly rigid bloc? The political deadlock of 2019–2022, in which four national elections were held amid repeated failures to form a governing coalition, provides a natural experiment for testing precisely this question.

Theoretical Framework

Three strands of research each capture part of the pattern observed in this study but none accounts for the full trajectory of stability, sudden disruption, and rapid recovery. Classic cleavage theory (Lipset and Rokkan 1967; Bartolini and Mair 1990) predicts that identity-based voting blocs remain durable so long as the social networks and organizational structures that encapsulate them persist. This explains the Haredi sector's high baseline loyalty but cannot account for the sudden disruption: if encapsulation is intact, why do loyalty rates plummet? Electoral volatility typologies (Pedersen 1979; Mainwaring and Zoco 2007) distinguish within-bloc from between-bloc volatility and recognize that elite-driven supply-side changes can produce temporary swings. Mainwaring, Gervasoni, and España-Najera (2017) further refine this distinction by separating within-system volatility (shifts among existing parties) from extra-system volatility (shifts to new parties), showing that the two types have different causes and consequences. Yet neither typology specifies the conditions under which such swings will reverse rather than crystallize into permanent realignment. Punctuated equilibrium models (Baumgartner and Jones 1993) capture the rhythm of long stasis interrupted by sudden change, yet they originated in policy studies and do not address the institutional mechanisms, specifically centralized spiritual authority, that can both trigger and terminate electoral disruptions in identity-based blocs.

This paper develops an interpretive framework, "rigidity with stress fractures," that draws on these perspectives to specify the conditions under which identity-based voting blocs experience reversible disruption rather than permanent realignment. The framework proposes that such blocs will exhibit high stability punctuated by sudden, reversible disruptions when three scope conditions are met: (i) a centralized institutional authority, such as rabbinic leadership, exercises strong influence over electoral behavior; (ii) that authority operates through organizational networks, such as yeshiva systems, with sufficient reach to produce geographically synchronized shifts; and (iii) the triggering crisis is exogenous and temporary rather than structural, so that withdrawal of the disrupting directive can restore the prior equilibrium. Under such conditions, synchronized volatility across geographically dispersed communities is consistent with coordinated elite action operating within an otherwise stable institutional framework, rather than indicating genuine erosion of group boundaries. The framework thus aims to distinguish temporary stress fractures—reversible disruptions driven by elite coordination—from both the permanent stability predicted by frozen-cleavage models and the gradual erosion assumed by dealignment theories.

If the framework applies, three observable patterns should distinguish elite-coordinated disruption from ordinary electoral volatility:

1. **Geographic synchronization.** Switching should appear simultaneously across geographically dispersed cities within the same electoral cycle, rather than emerging at different times in different localities—because the coordinating authority operates through national institutional networks.
2. **Within-bloc direction.** Switching should occur between parties within the identity bloc rather than toward external parties—because elite directives channel voters through intra-communal networks, not toward outsiders.

 doi:10.20944/preprints202511.2148.v2

3. **Rapid, synchronized recovery.** Once the coordinating directive is withdrawn, retention rates should return to baseline within one electoral cycle—the leak stops as fast as it starts, even if individual voters who already switched do not return.

The Results section evaluates these expectations against the estimated transition matrices.

Methods

Data Sources and Scope

This study uses official election results from the Central Elections Committee of Israel for Knesset elections 19 through 25, spanning from the January 2013 election (Knesset 19) through the November 2022 election (Knesset 25). These datasets provide polling-station-level counts of registered voters, valid votes, and votes for each political party. The analysis focuses on cities and towns with a significant ultra-Orthodox (Haredi) presence, following the same selection framework as in my previous study on intra-Haredi ethnic segregation (Gorelik, 2025). Throughout this paper, "country-level" refers to aggregated estimates across all ballot boxes that meet the inclusion criteria (defined below) nationwide, while "city-level" refers to included boxes within each city. All data and replication code are available at https://github.com/bgbg/segregation_paper.

In Israel, voters are assigned to polling stations by residential address. Typical ballot boxes contain 536–650 registered voters (interquartile range from recent elections). I identify Haredi ballot boxes using a 75% threshold: boxes where combined votes for Shas and UTJ exceed three-quarters of all votes cast, maintaining consistency with Gorelik (2025). To ensure reliable statistical estimation, I include only cities with at least five qualifying boxes.

Data were obtained directly from the Central Elections Committee's official online repository, supplemented by previously digitized archives for earlier elections (via Dr. Keren-Kratz, personal communication, 2024). The resulting dataset includes polling-station-level results for all eligible cities, with a focus on those exhibiting consistent Haredi voting patterns.

Ballot-Box Alignment and Analytical Categories

Ballot boxes were matched between consecutive elections by city and ballot-box ID. Where multiple boxes shared the same citybox pair across elections, their results were aggregated. Mismatched boxes that could not be aligned between consecutive elections were excluded from that specific transition analysis, but remained available for analysis in subsequent election pairs where they could be matched. This ensured that all included boxes represented consistent geographical and demographic units across each election pair while maximizing data retention.

Abstention was defined as the difference between the number of registered voters in a ballot box and the number of valid votes cast. This captures both non-voters and those who submitted invalid or blank ballots, following common practice in aggregate electoral analysis (Romero et al. 2020; King 1997).

Four electoral categories were modeled: (1) Shas (Sephardic ultra-Orthodox party), (2) UTJ (United Torah Judaism, the Ashkenazi ultra-Orthodox party alliance), (3) Other parties (all non-Haredi political lists), and (4) Abstention (registered but did not cast a valid vote). All non-Haredi parties are combined into a single "Other" category because our focus is intra-Haredi dynamics; distinguishing among individual parties would add dimensionality without illuminating Haredi behavior. The restriction to homogeneous Haredi ballot boxes (≥75% Shas+UTJ votes) ensures that nearly all Shas and UTJ votes in this sample originate from Haredi populations, isolating within-sector dynamics from broader societal trends.

Model Overview

The estimation framework follows a hierarchical Bayesian ecological inference model, designed to infer unobserved voter transition matrices between consecutive elections from aggregate data. This

 doi:10.20944/preprints202511.2148.v2

problem is inherently challenging because individual ballots are secret and only aggregate counts are available. Thus, I adopt a classic ecological inference scenario (Robinson 1950; King 1997). The hierarchical structure pools information across localities, improving precision in small samples while maintaining flexibility across elections. Full model specifications and diagnostics are provided in Appendix B.

For each election pair $(t, t+1)$, the model infers the probability that a voter from category i at election t votes for category j at election t+1. Individual transitions are unobserved, but aggregate vote counts provide marginal constraints. Building on Goodman's (1953) ecological regression framework and subsequent Bayesian extensions (King 1997; Rosen et al. 2001), I model voter transitions probabilistically using a hierarchical Bayesian approach that generalizes the linear model to address several key challenges: unobserved individual transitions, spatial heterogeneity in voting patterns, high dimensionality of transition matrices, and overdispersion in vote counts (Wakefield 2004; Forcina and Pellegrino 2019).

Important conceptual distinction: The transition probabilities represent the *probability of voting behavior* at each election, not the *movement of individual voters* between parties. When a party's retention rate increases, this indicates that leakage to other parties stopped, a return to a new steady state, rather than the literal return of specific voters who previously defected.

Each election pair is modeled using three levels of hierarchy: (1) National transition matrix ($M^{country}$) represents the baseline transition probabilities for the entire electorate, (2) Shared deviation pattern (D)a low-rank structure capturing how cities collectively deviate from national trends, and (3) City-specific scalar (δ_c)a single latent variable scaling the deviation pattern for each city.

This hierarchical structure avoids the "uniform swing" assumption of traditional ecological inference while managing dimensionality. I impose a rank-1 structure on city-level deviations: each city's deviation is captured by a single scalar δ_c that scales a shared deviation pattern D, reflecting that electoral shifts are often driven by one dominant cleavage rather than completely unique city-by-city patterns (Brown and Payne 1986; Puig and Ginebra 2015).

City-level transition matrices are constructed as:

$$M^{(c)} = \text{softmax}(Z^{country} + \delta_c D)$$

where $Z^{country}$ are national logits transformed via the softmax function. This yields a compact yet flexible structure that captures both overall national shifts and localized deviations, while enabling partial pooling: cities with limited data shrink toward the national pattern, while those with strong evidence of deviation can diverge (Gelman and Hill 2007).

Priors and Likelihood

Transition logits are modeled using logistic-normal priors, avoiding the independence assumptions of Dirichlet distributions (Wakefield 2004; Glynn and Wakefield 2010). I add a diagonal bias term B to same-party transitions to reflect voter inertia, the well-documented tendency to stick with previous choices (Campbell, Green, and Layman 2011; Clarke et al. 2004). The unusually short intervals between Israeli elections during this period (5.3 to 19 months) further strengthen this inertia prior. Full prior specifications, likelihood details, and overdispersion handling are provided in Appendix A.

Temporal Extension, Validation, and Implementation

The model is applied sequentially across subsequent election pairs from the January 2013 election (Knesset 19) through the November 2022 election (Knesset 25). Posterior means of $Z^{country}$, D, and B from one transition are used as prior means for the next, allowing gradual temporal evolution while maintaining continuity across election cycles.

Country-level transition matrix parameters converged well across all election pairs (Gelman-Rubin $\hat{R} < 1.01$, effective sample sizes exceeding 6,500, no divergent transitions). Posterior

 doi:10.20944/preprints202511.2148.v2

predictive checks confirmed good fit; full convergence diagnostics, including city-deviation parameters, are reported in Appendix B. Models were implemented in PyMC 5.0 using four chains with 3,000 draws and 5,000 tuning steps per chain (target acceptance 0.99, seed 42). Full diagnostics and code are available at https://github.com/bgbg/segregation_paper.

Corpus Analysis of Haredi Media

To investigate the mechanisms behind the observed electoral disruption, I collected a corpus of Haredi news articles and forum discussions from the two dominant ultra-Orthodox news websites in Israel: Behadrey Haredim (bhol.co.il) and Kikar HaShabbat (kikar.co.il). The corpus was assembled by systematically scraping all available content published between March 2020 and March 2021, yielding approximately 58,000 items: 25,000 news articles from Behadrey Haredim, 26,000 news articles from Kikar HaShabbat, and 7,000 forum discussion threads from the Behadrey Haredim forums.

Each item was stored as a structured record containing publication date, title, full text, author, and source URL. Forum threads additionally preserved individual posts with per-post dates and authors, capturing the temporal development of community discussions.

I searched the corpus systematically using Hebrew-language keyword queries organized around the research hypotheses: rabbinic voting instructions, intra-Haredi party friction, leadership crises, and COVID-related political upheaval. The initial automated filtering reduced the corpus to several hundred candidate articles, which were then read and evaluated manually. Each claim presented in the "What Caused the 23-24 Disruption?" section is supported by verbatim citations from the original sources, with full source metadata (publication date, outlet, and URL) provided in the Supplementary Materials.

Note on notation: Throughout this paper, election transitions are denoted using the arrow notation (→) to indicate the direction from one election to the next. For example, "23→24" refers to the transition from the March 2020 election (Knesset 23) to the March 2021 election (Knesset 24). This notation emphasizes that I am examining voter flows between consecutive elections rather than the elections themselves in isolation.

The following section presents the estimated transition patterns at national and city levels.

Results

The analysis confirms high baseline loyalty within both Shas and UTJ, yet reveals significant temporal and geographic variation in voter transitions. Most strikingly, a dramatic but temporary disruption in the March 2020–March 2021 transition (Knesset 23→24) affected all major Haredi cities, followed by complete recovery. Model diagnostics confirming adequate fit are provided in Appendix B.

Table 1 presents the election dates and inter-election intervals for all Knesset elections covered in this study.

Table 1. Election Dates and Intervals.

Knesset	Election Date	Months Since Previous
18	February 10, 2009	— (baseline)
19	January 22, 2013	47.4
20	March 17, 2015	25.8
21	April 9, 2019	48.8
22	September 17, 2019	5.3
23	March 2, 2020	5.5
24	March 23, 2021	12.7
25	November 1, 2022	19.3

Note: Knesset 18 serves as the baseline for computing the first transition (18→19) but is not itself analyzed independently. The transition analysis covers Knessets 19–25 as described in Methods.

 doi:10.20944/preprints202511.2148.v2

Country-Level Transitions

At the national level, the transition matrices reveal strong voter loyalty within both Haredi parties. Shas retained on average more than 90% of its voters across elections, while UTJ consistently preserved above 95%. However, the magnitude of "within-bloc permeability"—estimated transitions between Shas and UTJ—varied notably over time.

The March 2020–March 2021 transition (Knesset 23→24) showed an unusual and dramatic decline in intra-Haredi loyalty, particularly among Shas voters. At the country level (Figure 2), Shas-to-Shas loyalty plummeted from 98.9% (in the September 2019–March 2020 transition, Knesset 22→23) to just 73.5% in the 23→24 transition. Simultaneously, the estimated Shas-to-UTJ transition probability jumped from near zero to 12.3%. UTJ retention also declined, dropping from 96.6% to 87.9%, with an estimated 4.6% UTJ-to-Shas transition. The estimated decline in party retention corresponds to roughly one parliamentary seat per party, illustrating the political significance of even modest swings in Haredi voting patterns. The temporary drop thus had a tangible potential to alter coalition outcomes, yet it was masked in aggregate results by offsetting trends among non-Haredi voters. This cross-flow pattern represents an unprecedented disruption in the typically stable Haredi voting bloc. Critically, this disruption was observed across multiple geographic scales: both at the national level (Figure 2) and across individual cities (Figure 4), indicating a system-wide rather than localized phenomenon—consistent with the geographic synchronization expected under elite-coordinated disruption. The dominant direction of switching was within the Haredi bloc (Shas ↔ UTJ) rather than toward external parties, consistent with within-bloc direction as expected when elite directives operate through intra-communal networks. Notably, this disruption was temporary: in the subsequent March 2021–November 2022 transition (Knesset 24→25), loyalty rates recovered substantially (Shas: 96.9%, UTJ: 95.5%), consistent with the rapid, synchronized recovery expected once a coordinating directive is withdrawn.

Important clarification: This recovery in retention probabilities does not indicate that individual voters who had "strayed" from Shas returned to the party. Rather, it reflects that the "leak" from Shas to other parties stopped (see Methods section on the distinction between voting behavior probabilities and individual voter movements). The actual voter movements are captured by the off-diagonal elements of the transition matrix, the flows between parties, not by the diagonal retention rates themselves. Critically, none of the off-diagonal transitions into Shas (UTJ→Shas, Other→Shas, Abstain→Shas) showed unusual spikes in the March 2021–November 2022 transition (Knesset 24→25), confirming that the Shas-to-UTJ transitions from the 23→24 disruption were not reversed.

Paradoxically, despite losing core Haredi vote share in 23→24 without recovering it in 24→25, Shas's national vote share increased from 7.17% to 8.25% (9 to 11 seats in the Knesset). Since Haredi population hubs show no corresponding Shas influx, this growth originated from voters outside major Haredi centers. This illustrates how aggregate vote-share growth can mask internal dynamics: Shas simultaneously lost votes in its core ultra-Orthodox base (to UTJ in 23→24) while gaining peripheral traditional Sephardic supporters.

As Table 1 shows, the short inter-election intervals during this period make it very unlikely that the observed transitions reflect demographic change through migration or generational replacement rather than genuine voter switching.

Figure 2 shows that Haredi abstention rates remained low and stable, contrasting with fluctuating non-Haredi participation. Cross-over voting between Haredi and non-Haredi parties remained marginal, indicating persistent political segmentation despite broader electoral turbulence.

 doi:10.20944/preprints202511.2148.v2

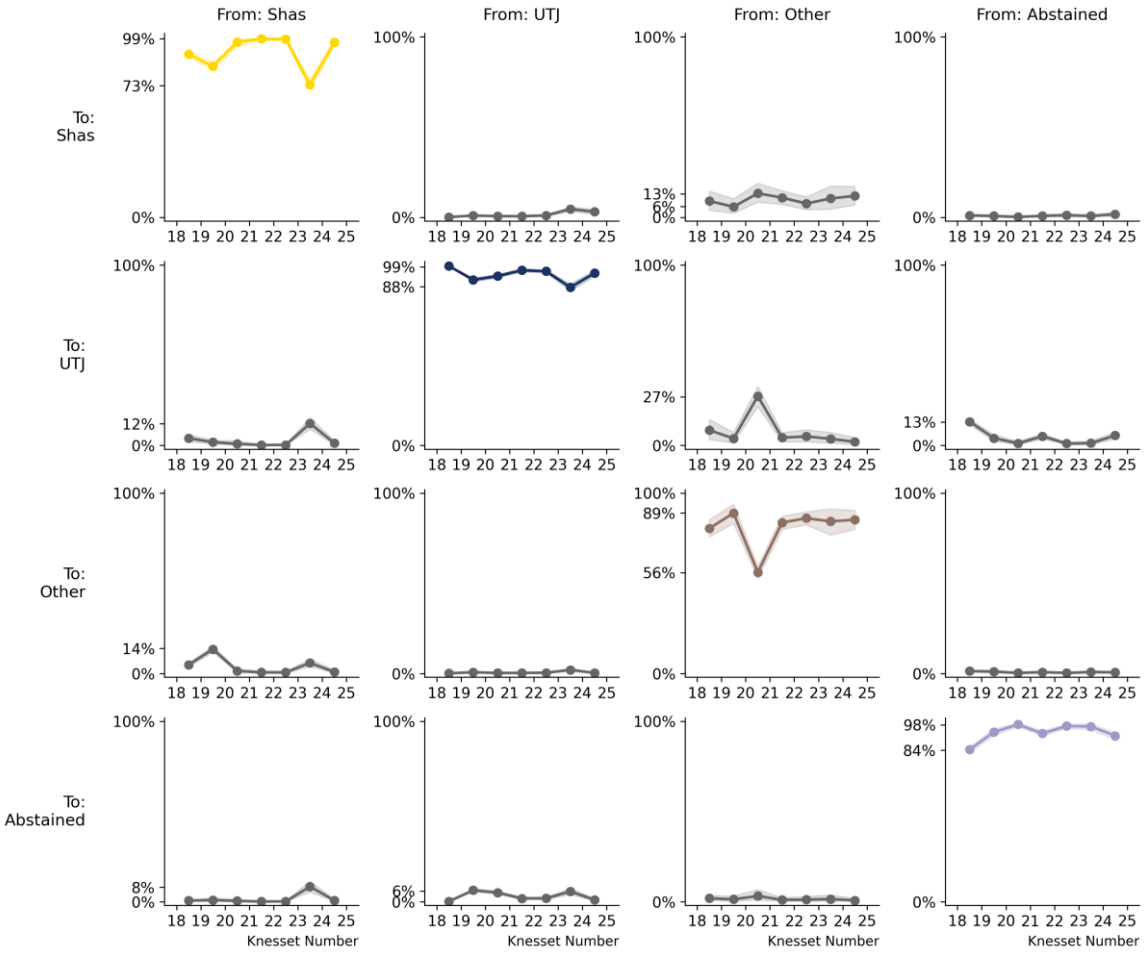

Figure 2. Shas loyalty collapsed and recovered; UTJ held. Country-level transition matrices across all election pairs. The March 2020–March 2021 disruption was exceptional in both magnitude and brevity. Haredi abstention rates stayed low throughout, and cross-over voting between Haredi and non-Haredi parties remained marginal.

Raw Vote Shares Across Haredi Hubs

Before examining city-level model estimates, it is useful to inspect the raw data. Figure 3 plots the election-over-election change in Shas vote share (in percentage points) for each city, computed directly from Haredi-filtered polling stations (those where Shas + UTJ exceed 75% of legal votes) and requiring no modeling assumptions. The 23→24 transition is the only one in which Shas's vote share declined in every city simultaneously; in all other transitions, some cities showed gains while others showed losses. The magnitude of the dip is modest (1–3 percentage points) because raw vote shares compress the underlying signal—a voter switching from Shas to UTJ depresses Shas's share while boosting UTJ's, partially canceling in the aggregate. The ecological inference model (below) disentangles these cross-flows and reveals a much larger disruption in loyalty probabilities. But the raw data already establish, without any model, that something happened simultaneously across geographically dispersed cities during the March 2020–March 2021 period. For evidence that these patterns are robust to model specification, see Appendix C.

 doi:10.20944/preprints202511.2148.v2

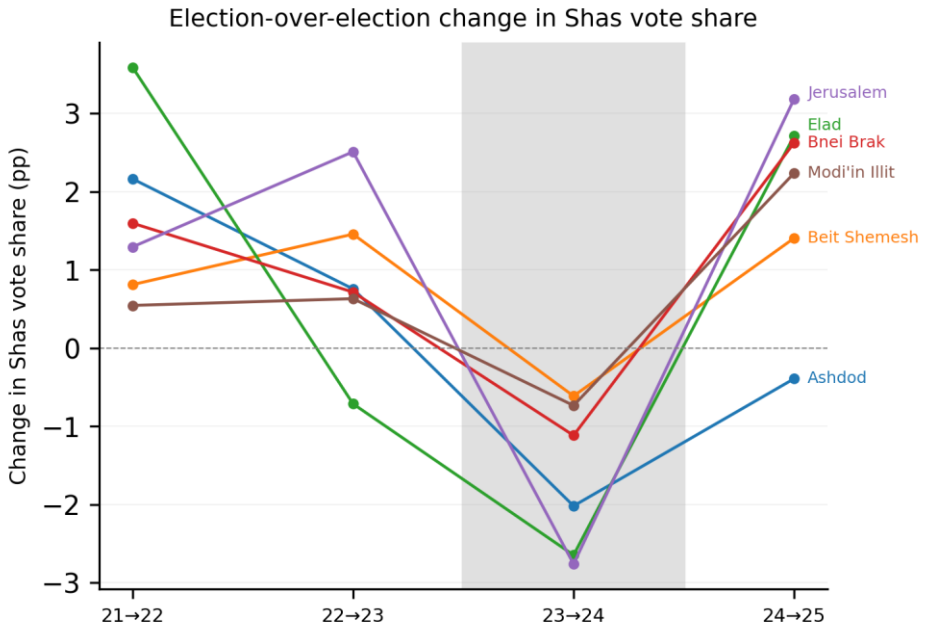

Figure 3. All cities declined simultaneously only in 23→24. Election-over-election change in Shas vote share (percentage points) in Haredi-filtered polling stations. Each line represents one city; the gray band marks the 23→24 transition (March 2020–March 2021), the only period in which every city moved in the same direction.

City-Level Variation

City-level analysis reveals substantial variation in the magnitude of the March 2020–March 2021 (23→24) loyalty disruption across Haredi strongholds. Figure 4 shows that the disruption was universal, every city experienced reduced Shas loyalty in the 23→24 transition, but the magnitude varied substantially. Ashdod (blue line) shows the steepest drop to approximately 65%, while other cities cluster between 70–78%. The recovery in the March 2021–November 2022 transition (24→25) was equally universal and nearly complete, with all cities returning to loyalty rates above 95%.

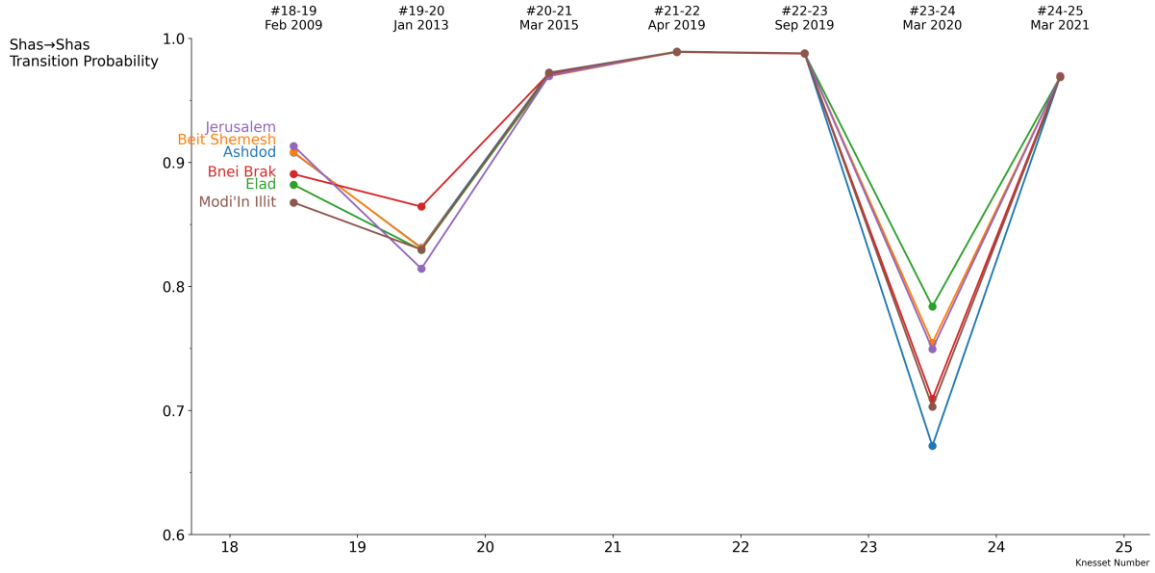

Figure 4. Synchronized collapse and recovery across all cities. Shas-to-Shas transition probabilities across cities and election pairs. All cities show the 23→24 loyalty drop followed by 24→25 recovery. Ashdod (blue) exhibited the most extreme deviation, dropping to approximately 65%.

 doi:10.20944/preprints202511.2148.v2

Ashdod exhibited the most dramatic deviation from national patterns. Notably, Ashdod already showed early signs of instability before the system-wide shock: Shas loyalty declined from 91.8% in the 21→22 transition to 83.6% in 22→23. In the 23→24 disruption itself, as shown in Figure 5, Shas-to-Shas loyalty dropped to just 67.1% (compared to 73.5% nationally), while the Shas-to-UTJ switching rate surged to 19.3%, more than 50% higher than the national rate of 12.3%. In the subsequent 24→25 transition, Ashdod's Shas loyalty recovered to 96.9%, closely tracking the national pattern.

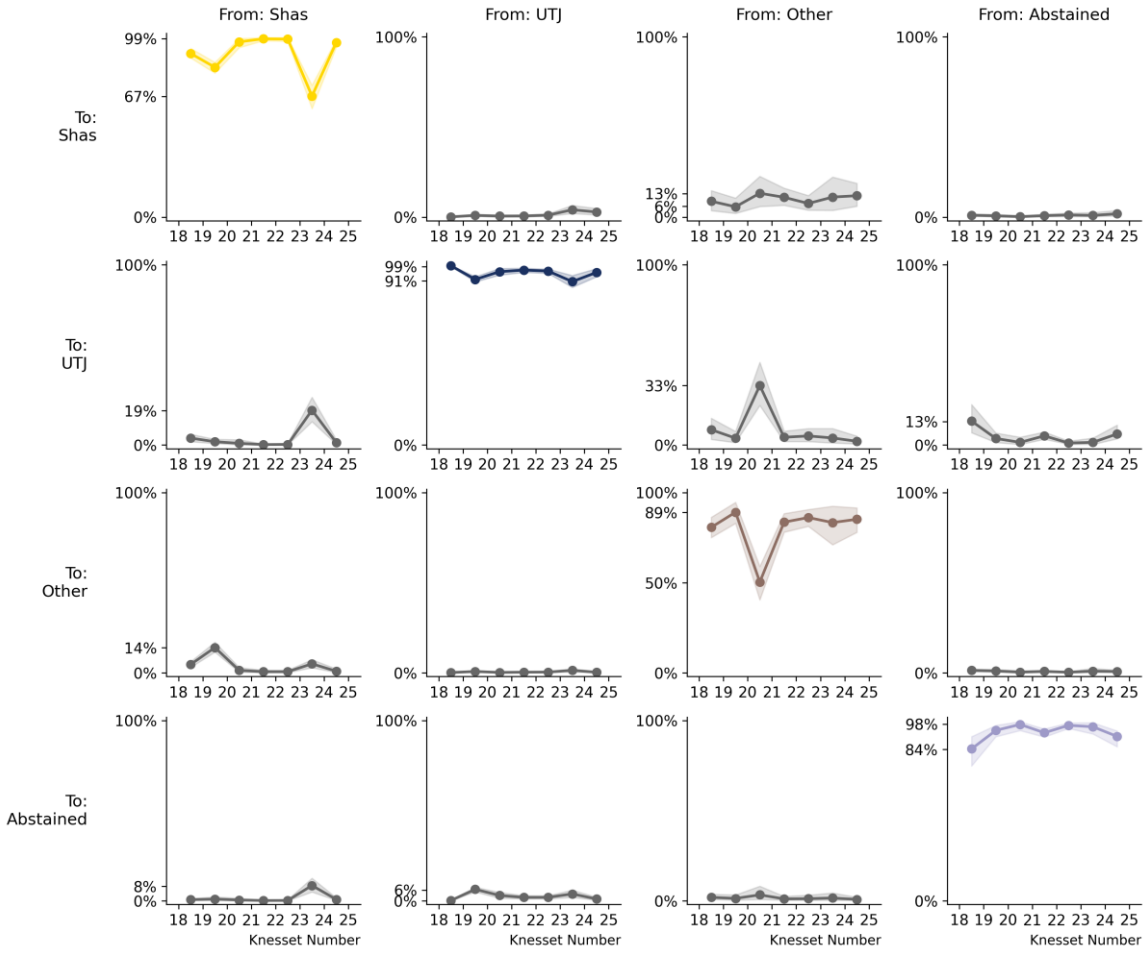

Figure 5. Ashdod showed the largest cross-ethnic switching. Full transition matrices for Ashdod across all election pairs. Shas-to-UTJ switching reached 19.3% in the 23→24 transition, more than 50% above the national rate.

Other major Haredi cities showed similar but more moderate disruptions. Beit Shemesh experienced a Shas loyalty drop to 75.4% (with 9.5% switching to UTJ), while Bnei Brak, despite being a predominantly Ashkenazi stronghold, saw Shas loyalty fall to 70.9% (with 15.8% switching to UTJ). Both cities fully recovered by the 24→25 transition. Detailed transition matrices for these cities are provided in Appendix A.

The consistency of this temporal pattern across cities, sharp disruption followed by full recovery, is consistent with a system-wide rather than city-specific phenomenon. The variation in magnitude of the estimated transitions correlates with demographic composition: cities with more integrated Sephardic-Ashkenazi populations (Ashdod, Bnei Brak) showed larger cross-flows than more homogeneous communities, suggesting that disruption magnitude may reflect local demographic conditions. Both parties maintained their coalition positions despite the internal reshuffling.

 doi:10.20944/preprints202511.2148.v2

A key concern is whether the hierarchical model structure mechanically produces the observed inter-city synchronization. To test this, I fitted two alternative specifications: a relaxed hierarchical model allowing substantially greater city-specific deviations (+166% variability) and a fully independent model with no hierarchical pooling (+248% variability). Despite cities diverging more from national patterns under both alternatives, the synchronized drops and recoveries in Shas loyalty persisted across all specifications (see Appendix C for full details and figures).

Taken together, the transition matrices reveal a pattern of high baseline stability punctuated by a sharp disruption in the 23→24 transition that fully reversed by 24→25 (smaller episodic deviations occurred in other transitions—notably a Shas loyalty drop to 85.0% with an estimated 13.8% transitioning to non-Haredi parties in the 19→20 transition—but none approached the 23→24 magnitude or geographic synchronization). The disruption affected all major Haredi population centers but varied in magnitude, with Ashdod showing the most extreme deviations. The short intervals between elections make demographic explanations implausible, pointing instead to genuine behavioral switching. The Conclusions section interprets these patterns within the "rigidity with stress fractures" framework and explores their theoretical and methodological implications.

Conclusions

This study demonstrates that ethnic voting blocs governed by centralized religious authority can experience sudden, large-scale electoral realignments that reverse just as rapidly—a pattern of "rigidity with stress fractures" that challenges both assumptions of permanent stability and theories of gradual dealignment. The Israeli Haredi case reveals how a small number of religious leaders can temporarily redirect the voting behavior of hundreds of thousands of geographically dispersed voters through institutional channels, producing synchronized disruptions that leave underlying political boundaries intact. These findings carry implications beyond Israel: any identity-based community with centralized institutional authority, dense organizational networks, and sensitivity to exogenous crises may exhibit similar dynamics, suggesting that electoral volatility in cohesive populations is better understood as a sign of institutional discipline than of emerging independence.

Geographic Variation in Disruption Magnitude

While the 23→24 disruption affected all major Haredi population hubs, its magnitude varied across localities, with Ashdod exhibiting the most extreme deviation (Shas loyalty dropping to 67.1%, Shas-to-UTJ switching reaching 19.3%). However, the consistency of the temporal pattern—sharp disruption followed by recovery across all cities—confirms that the disruption was system-wide rather than driven by city-specific factors alone.

Despite its magnitude, this disruption appears to have received limited attention in contemporary political discourse. Both parties maintained their coalition positions despite the internal reshuffling, illustrating how substantial voter transitions within the Haredi sector can remain largely hidden beneath the surface of stable electoral outcomes.

Temporal Dynamics and Demographic Interpretations

With five national elections held over just three and a half years (April 2019 to November 2022), the compressed electoral calendar makes it very unlikely that observed transitions reflect demographic change through migration or generational replacement rather than genuine voter switching, especially given the near-instantaneous nature of both the disruption and the recovery.

This finding has important implications for interpreting the timing and reversibility of electoral volatility in highly cohesive communities. Even populations characterized by strong institutional discipline can experience rapid political realignments, yet the equally rapid recovery indicates that such disruptions need not signal permanent weakening of ethnic-political boundaries. As the corpus analysis suggests (see "What Caused the 23-24 Disruption?" below), the evidence is most consistent with coordinated rabbinic guidance as the primary mechanism: community leaders explicitly

 doi:10.20944/preprints202511.2148.v2

directed cross-ethnic voting as a strategic response, which would account for the synchronized shifts observed across geographically dispersed cities within a single election cycle. The partial withdrawal of these directives by the subsequent election would similarly account for the rapid return to baseline loyalty.

Connecting Spatial and Political Boundaries

This study was motivated by an anomaly in residential segregation indices documented in Gorelik (2025): Ashdod's dissimilarity index dropped sharply between the March 2020 and March 2021 elections despite no demographic change. The transition matrices confirm that this anomaly was electoral rather than demographic—Ashdod's exceptionally high Shas-to-UTJ switching rate (19.3%) altered the party-vote proxies used to measure ethnic segregation, creating the appearance of residential integration without actual movement.

What Caused the 23–24 Disruption?

The timing coincided with extraordinary political and social stress: unprecedented political instability (three elections in 12 months), the COVID-19 pandemic, intense coalition negotiations, and Haredi leadership tensions (Gidron, Sheffer, and Mor 2022). Several mechanisms could plausibly account for the observed disruption:

- **Strategic voting:** Voters may have anticipated coalition outcomes and switched parties to maximize Haredi representation. However, strategic voting would not explain why switching was specifically cross-ethnic (Sephardic→Ashkenazi party) rather than toward non-Haredi alternatives, nor why it occurred simultaneously across geographically dispersed cities.
- **Protest voting:** Dissatisfaction with Shas leadership or performance could have driven defection. Yet protest voting typically produces gradual, staggered changes across localities with idiosyncratic timing, not the synchronized drops observed simultaneously in Jerusalem, Bnei Brak, Ashdod, Beit Shemesh, and Modi'in Illit. Moreover, the fact that Shas did not lose its overall representation in the Knesset further supports this interpretation, as the switching was within the Haredi bloc rather than away from it.
- **COVID-19 backlash:** The pandemic created widespread political grievances within Haredi communities. However, a survey of Haredi respondents conducted by Kikar HaShabbat, a leading ultra-Orthodox news website (see Appendix D, items D14–D15), found that Shas's handling of the pandemic was rated significantly better than UTJ's across all sub-sectors, so COVID-related dissatisfaction alone would predict movement *away from* UTJ, not toward it.
- **Coalition expectation effects:** Voters may have shifted based on predictions about government formation. This does not account for the specifically cross-ethnic direction of switching (Sephardic voters moving to an Ashkenazi party) or the rapid, synchronized recovery.

A systematic search of approximately 58,000 Haredi news articles and forum posts from this period (from Behadrey Haredim, a major ultra-Orthodox online forum, and Kikar HaShabbat, March 2020 to March 2021) points to coordinated rabbinic guidance as the most plausible mechanism—that is, the explanation most consistent with the observed patterns—specifically, a top-down campaign by UTJ's political and rabbinic leadership to recruit Sephardi voters.

The corpus documents a coordinated campaign in which UTJ's political and rabbinic leadership recruited Sephardi voters through a chain of authority—from senior MKs to prominent Sephardi rabbis (notably Rabbi Meir Mazuz and Rabbi Yitzhak Barda) to yeshiva heads to individual voters—spanning the entire national geography of Haredi settlement (see Appendix D for verbatim sources and full citations).

By the Knesset 24 elections (March 2021), this campaign had fragmented. Mazuz retreated from his exclusive UTJ endorsement, directing some supporters toward Religious Zionism Party instead. The splintering of rabbinic instruction is consistent with the partial recovery observed in the transition matrices.

 doi:10.20944/preprints202511.2148.v2

COVID-19 likely served as an enabling factor rather than a direct cause: while 31.7% of surveyed Haredi respondents said the pandemic would affect their vote (Appendix D, item D14), the direction of COVID-related dissatisfaction (favoring Shas over UTJ) is inconsistent with the observed Shas-to-UTJ switching. COVID may have created political mobility that could be channeled by rabbinic directives, but does not appear to have independently driven voters from Shas.

The elite coordination interpretation is consistent with the synchronized disruption and recovery documented in the quantitative analysis: a campaign operating through centralized rabbinic networks with country-wide reach would account for simultaneous loyalty drops across dispersed cities, and the withdrawal of those directives would account for the equally synchronized recovery. This interpretation best explains the specific features of the disruption that alternative mechanisms cannot individually account for: the synchronized timing across cities, the specifically cross-ethnic direction of switching, and the rapid recovery coinciding with the withdrawal of rabbinic directives. However, the corpus documents public statements and reported events, not the private deliberations of individual voters; the link between rabbinic instruction and voter behavior, while supported by the convergence of qualitative and quantitative evidence, cannot be established with certainty from these sources alone. While the qualitative evidence strongly suggests that rabbinic directives played an important role, the analysis cannot establish a direct causal link between elite statements and individual voting behavior.

Theoretical Implications

Mapping Findings onto the Framework

The empirical findings are consistent with the "rigidity with stress fractures" framework: all three observable expectations—geographic synchronization, within-bloc direction, and rapid recovery—are confirmed by the transition matrices. Existing theories each capture part of the observed pattern: cleavage theory (Lipset and Rokkan 1967; Bartolini and Mair 1990) accounts for the high baseline stability, as the Haredi sector's dense organizational encapsulation—yeshivot, synagogues, rabbinic courts—creates the socio-organizational bonds that freeze electoral alignments. Electoral volatility typologies (Pedersen 1979; Mainwaring and Zoco 2007; Mainwaring, Gervasoni, and España-Najera 2017) distinguish within-system from extra-system volatility and recognize elite-driven disruptions, but do not specify when such disruptions reverse. Punctuated equilibrium (Baumgartner and Jones 1993) captures the rhythm of stasis-punctuation-recovery but does not address the specific institutional mechanisms that trigger and terminate disruptions in identity-based blocs.

The Haredi case illustrates how all three scope conditions can operate simultaneously.

First, **centralized institutional authority**: rabbinic leaders appear to have exercised direct electoral influence through binding rulings, consistent with Horowitz's (1985) framework of elite-orchestrated ethnic voting and Freedman's (2020) documentation of rabbinic electoral effects in Israel.

Second, **organizational network reach**: the yeshiva system provided channels for transmitting directives across geographically dispersed cities within a single electoral cycle, producing the synchronized disruption observed in Jerusalem, Bnei Brak, Ashdod, Beit Shemesh, and Modi'in Illit.

Third, **an exogenous, temporary crisis**: the triggering conditions—COVID-19, coalition instability, inter-party recruitment campaigns—were circumstantial rather than structural. Once rabbinic guidance redirected voters to UTJ, those who heeded the directive switched; when no further directive was issued in the opposite direction, no countervailing movement occurred, and baseline loyalty reasserted itself.

This framework helps distinguish the Haredi pattern from two alternatives. Unlike the gradual realignment described by Goodwin, Kaufmann, and Larsen (2022), where identity-driven issues produce lasting partisan shifts in open electorates, the Haredi disruption started and stopped suddenly and synchronously across cities—the leak halted, even though individual voters did not switch back. And unlike the cleavage erosion implied by dealignment theories, the underlying ethnic-

 doi:10.20944/preprints202511.2148.v2

political boundaries remained intact: the same institutional structures that enabled the disruption also terminated it. The closest documented parallel comes from Italy's Christian Democracy, where Lanzara et al. (2024) show that the identity of the bishop in office explained significant variation in party vote share across Italian dioceses from 1948 to 1992, with bishop replacements producing swings of up to three percentage points. As in the Haredi case, a centralized religious authority operating through organizational networks shaped electoral outcomes across dispersed communities. However, the Italian evidence captures cross-sectional variation in bishop influence, not a sudden synchronized temporal disruption followed by recovery—underscoring why the Haredi case, with its repeated elections over a compressed period, provides a uniquely informative window into the dynamics of elite-coordinated volatility.

Synchronized Disruption as an Indicator of Discipline

The sharp drops in intra-party loyalty during the 23→24 transition might initially suggest emerging voter independence. However, the **sudden and synchronized nature** of both disruption and recovery across geographically dispersed cities is more consistent with **institutional discipline** than with independent thinking. If voters were exercising individual judgment, one would expect gradual, staggered changes across localities with uncoordinated recovery. Instead, simultaneous loyalty drops followed by synchronized recovery are consistent with a single rabbinic directive transmitted through institutional channels spanning the national geography of Haredi settlement, producing a 12-15% swing across all major population centers within the same electoral cycle. That baseline loyalty reasserted itself when those directives were withdrawn supports the interpretation that both the disruption and the recovery were expressions of the same underlying discipline. This parallels theoretical expectations for other identity-based voting blocs with centralized institutional authority, though comparable micro-level evidence of synchronized reversible switching has not been documented elsewhere. A useful contrast comes from Ward and Denney's (2025) study of South Korea, where party organizations are constantly renamed and reformed yet voters maintain stable partisan preferences anchored in programmatic divides. The Haredi case inverts this pattern: organizational structures are exceptionally stable, but voter loyalty can be temporarily channeled through those very structures. Together, the two cases suggest that voter loyalty and organizational stability are partially independent—and that the institutional channels connecting leaders to voters, not organizational permanence per se, determine the capacity for coordinated electoral shifts.

The persistence of ethnic-political boundaries despite temporary volatility parallels Andersen and Yaish's (2003) finding that Israeli electoral reform failed to shift cleavage voting patterns, and is consistent with Bartolini and Mair's (1990) argument that high organizational encapsulation enables temporary disruption without permanent realignment—an argument that Martin, de Lange, and van der Brug (2022) have shown still holds in contemporary Western Europe, where parties with stronger civil-society linkages maintain more stable electorates even in an era of rising volatility. The Haredi case suggests that in communities meeting the three scope conditions, change may occur through punctuated disruption and recovery rather than slow drift.

The Limited Predictive Power of Structural Integration

The geographic variation in disruption magnitude might suggest that residential proximity facilitates political permeability. However, universal recovery to baseline loyalty argues against a simple integration-drives-switching model. If residential proximity weakened political boundaries, I would expect gradual convergence, not a temporary surge of switching followed by an abrupt halt. Instead, demographic integration creates latent potential for cross-party switching that activates during crises. Mixed cities like Ashdod serve as bellwethers precisely because their populations make boundaries more permeable when system-wide shocks occur, not because integration itself drives ongoing fluidity.

 doi:10.20944/preprints202511.2148.v2

Methodological Contributions

This natural experiment, Israel's 2019–2022 political deadlock, provided a rare opportunity to observe voter transitions over a short period and demonstrates the utility of hierarchical Bayesian ecological inference for analyzing such transitions in contexts where individual-level data are unavailable due to ballot secrecy. The framework successfully captures both national baseline patterns and city-specific deviations while accounting for the inherent uncertainty in ecological inference. The temporal extension—using posterior distributions from one election pair to inform priors for the next—allows the model to trace gradual evolution while remaining flexible enough to detect sharp disruptions.

The approach is replicable across other contexts where: -Aggregate data (ballot-box or precinct-level) are available -Individual voting patterns are unobserved -Both national trends and local variation are of interest -Multiple time periods permit temporal modeling

More broadly, the Ashdod case demonstrates that studies using party votes as proxies for ethnic composition should account for the fragility of this assumption in disciplined populations, where elite directives may rapidly redirect voting behavior and produce apparent demographic shifts that are purely electoral artifacts (see Gorelik 2025 for the segregation analysis that first relied on this proxy).

Final Reflections

This study reveals that ethnic-political boundaries within Israel's ultra-Orthodox society are simultaneously rigid and vulnerable—and the evidence presented here is consistent with both their rigidity and their vulnerability deriving from the same source: the considerable authority of spiritual leaders over the electoral behavior of hundreds of thousands of voters dispersed across the country. The "rigidity with stress fractures" framework proposed here captures this duality: high baseline stability punctuated by sudden, elite-coordinated disruptions that cease as abruptly as they begin, leaving the underlying institutional structure intact.

What the data clearly show is that assumptions of permanent stability, or conversely, ongoing fluidity, both mischaracterize ethnic voting dynamics in cohesive religious communities. The convergence of quantitative and qualitative evidence suggests that centralized religious authority may play a powerful coordinating role in Haredi electoral behavior. The synchronized switching patterns observed here are consistent with the possibility that directives from a small number of influential rabbis can temporarily redirect voting behavior across geographically dispersed communities. While the available data cannot directly link individual voters to specific instructions, the temporal alignment between elite appeals and the aggregate transition patterns points to elite coordination as a plausible—and, given the evidence, the most parsimonious—mechanism.

Funding: This study was partially funded by the Azrieli College of Engineering—Jerusalem Research Fund.

Data Availability Statement: The electoral data analyzed in this study are publicly available from the Israeli Central Elections Committee (https://bechirot.gov.il/). The datasets and analysis code used in this study are available on GitHub at https://github.com/bgbg/segregation.

Acknowledgments: The author thanks Menachem Keren-Kratz, an independent researcher, for his very important discussions and suggestions, which significantly improved the depth of this work, as well as for generously providing ready-to-analyze files of earlier election results.

Conflicts of Interest: The author declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Ethics Statement: This study uses ecological inference applied to aggregated, publicly available electoral data published by the Israeli Central Elections Committee. No individual-level data were collected or analyzed, and no human subjects were involved. Accordingly, institutional review board (IRB) or ethics committee approval was not required.

 doi:10.20944/preprints202511.2148.v2

Declaration of Use of AI in the Research and Writing Process: The following AI tools were used during the preparation of this manuscript:

- ChatGPT (OpenAI) was used to assist with literature search and review.
- Claude (Anthropic) was used to assist with structuring the theoretical framework and generating initial drafts.
- Claude Code (Anthropic) was used to assist with formatting references and citations, generating tables and illustrations, and ensuring consistency in formatting.

The author affirms that all AI-generated content was reviewed and edited by the author to ensure accuracy and relevance, and that the final manuscript reflects the author's original ideas and interpretations.

Appendix

Appendix A City-Level Transition Matrices

The following figures provide detailed transition matrices for all analyzed cities across consecutive elections. These supplement the summary analysis presented in the main text, particularly the condensed discussion in the City-Level Variation section.

Ashdod

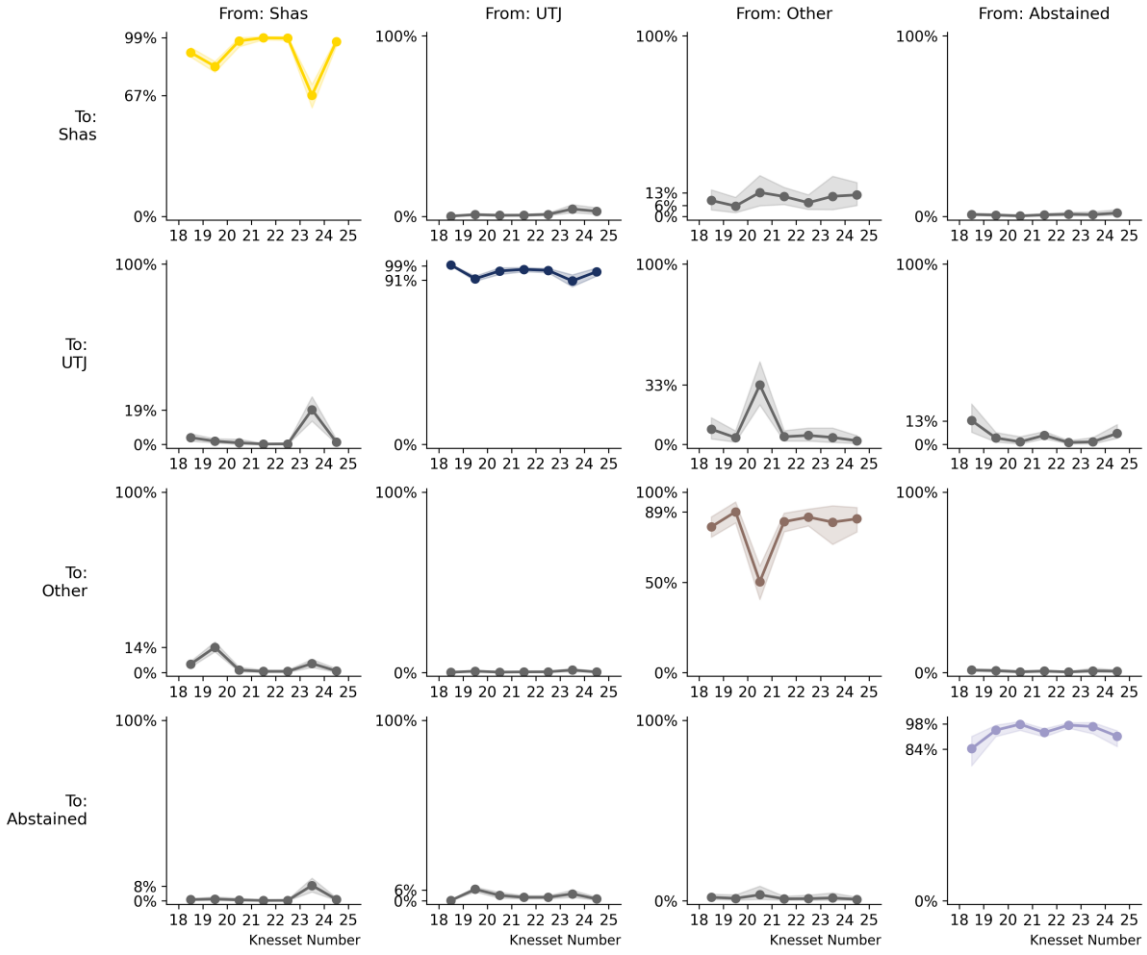

Figure A1. Ashdod exhibited the most extreme disruption during the March 2020–March 2021 transition (2324), with Shas loyalty dropping to 67.1% and 19.3% switching to UTJ.

 doi:10.20944/preprints202511.2148.v2

Beit Shemesh

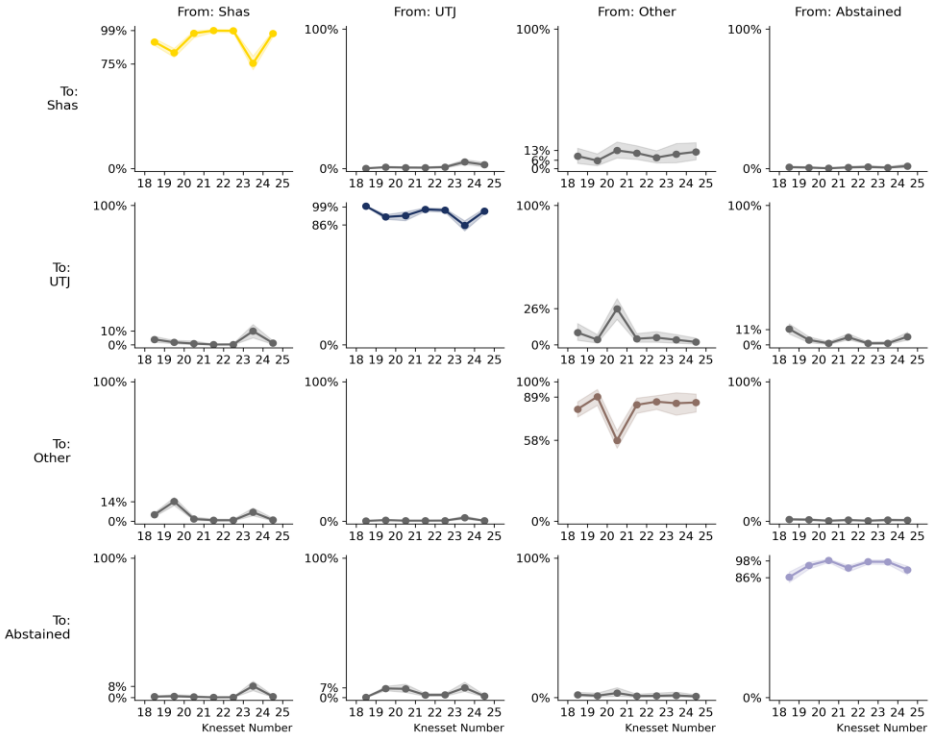

Figure A2. Beit Shemesh showed moderate disruption with Shas loyalty falling to 75.4% during the 2324 transition.

Elad

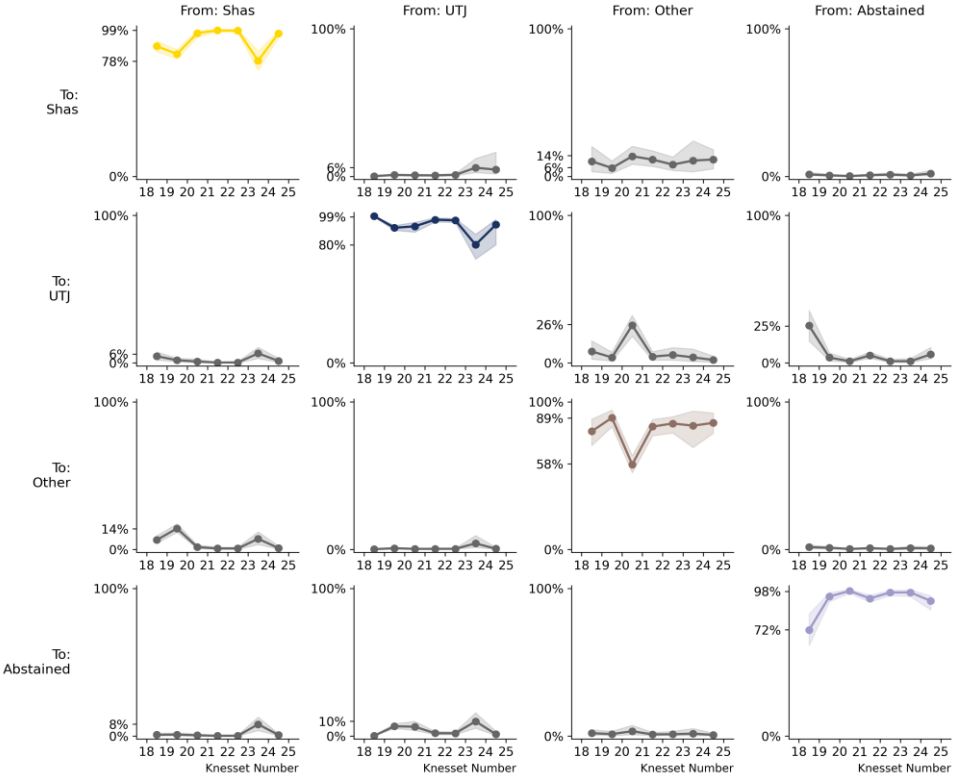

Figure A3. Elad transition patterns across all election pairs.

 doi:10.20944/preprints202511.2148.v2

Bnei Brak

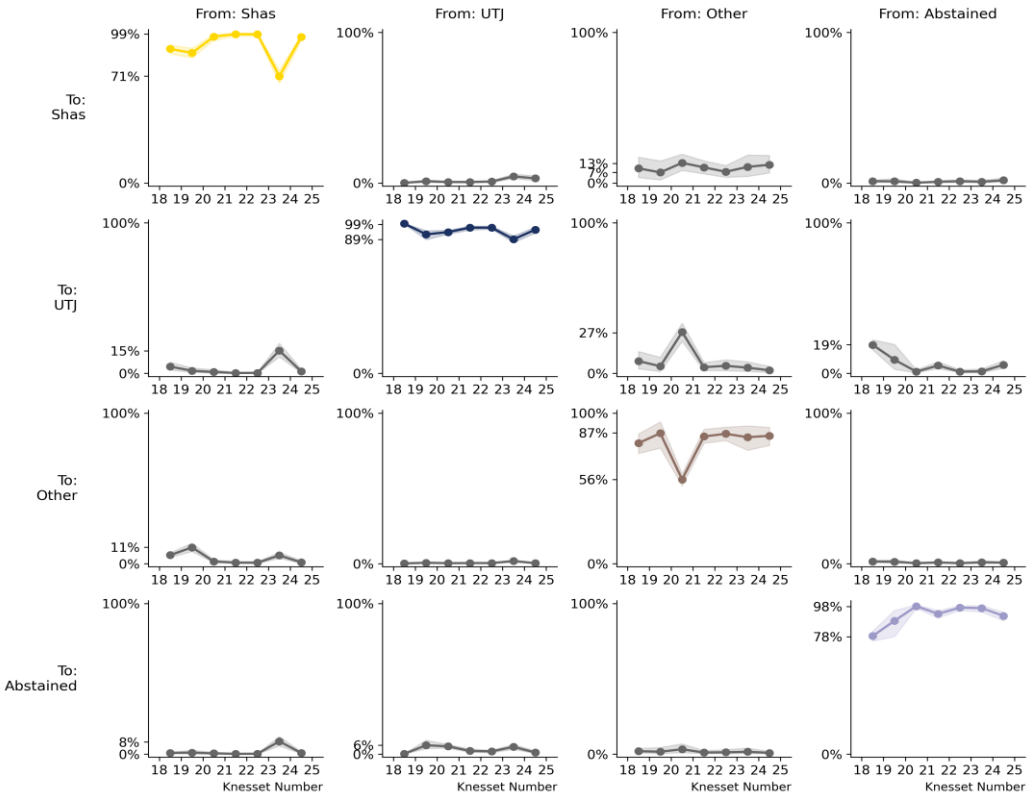

Figure A4. Bnei Brak, despite being predominantly Ashkenazi, experienced a sharp Shas loyalty drop to 70.9% during the 2324 transition.

Jerusalem

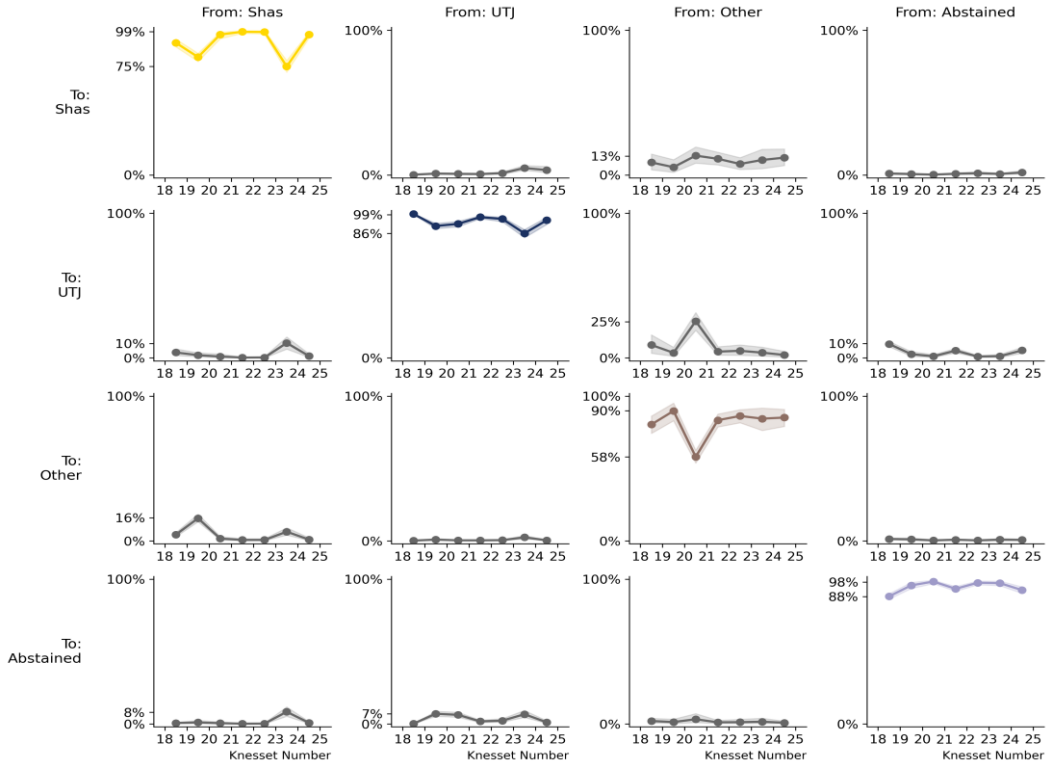

Figure A5. Jerusalem transition patterns across all election pairs.

 doi:10.20944/preprints202511.2148.v2

Modi'in Illit

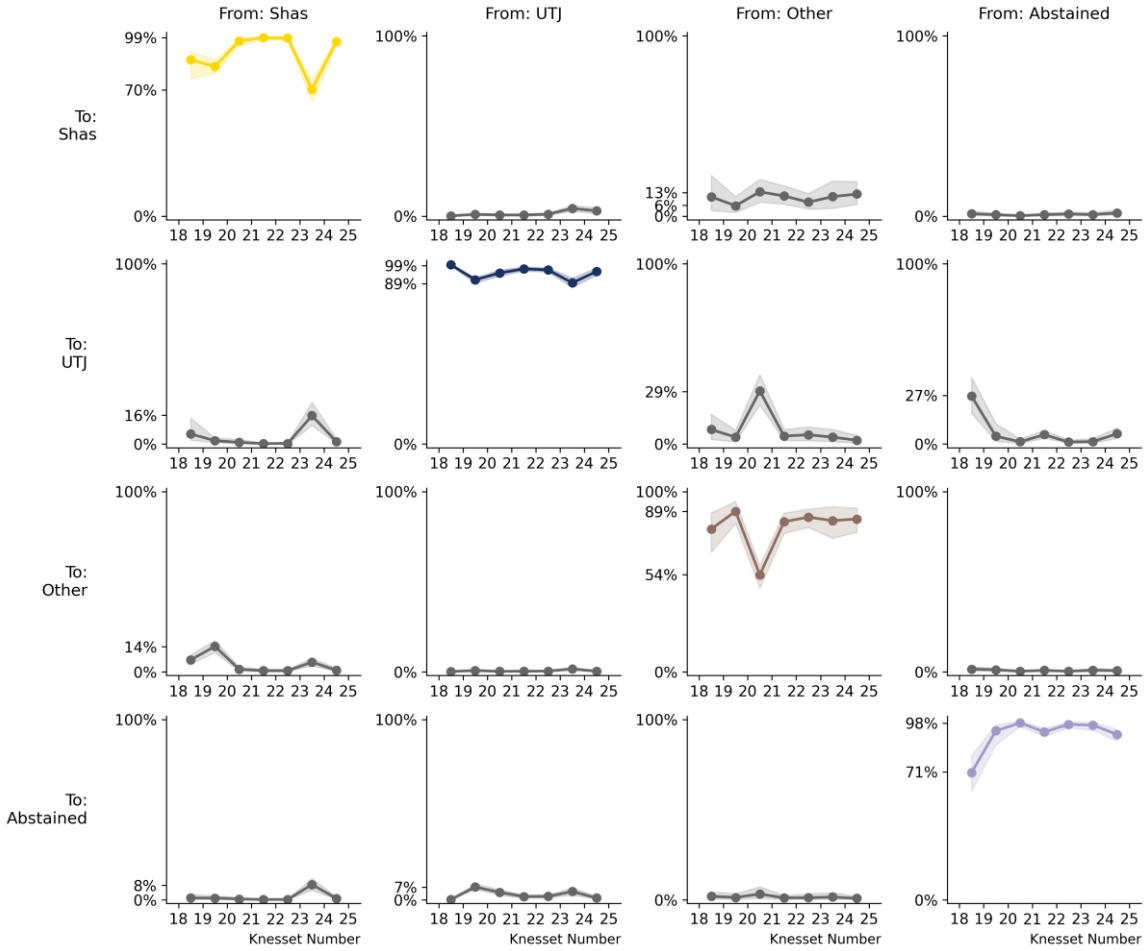

Figure A6. Modi'in Illit transition patterns across all election pairs.

Appendix B Model Validation and Diagnostics

This section documents model validation through posterior predictive checks and convergence diagnostics for transparency and reproducibility.

Model Validation

Posterior predictive checks confirm that the model adequately reproduces empirical vote counts across all elections. All models were sampled with 4 chains, 3,000 draws and 5,000 tuning steps per chain, target acceptance rate 0.99, and random seed 42 for reproducibility.

Convergence Diagnostics

Country-level transition matrix parameters (Z_country, diag_bias) converge well across all election pairs, with R-hat < 1.01 and ESS > 6,500 for all parameters. City-deviation parameters (D, delta) show non-convergence in some pairs due to the multiplicative structure of the city-deviation model (delta * D), which creates a scaling non-identifiability between the deviation pattern and its city-specific scaling factors. This does not affect country-level transition estimates, which are the primary quantities of interest.

 doi:10.20944/preprints202511.2148.v2

Transition	Country R-hat max	Country ESS min	Divergences
Kn 18-19 (Feb 2009–Jan 2013)	1.001	9,816	0
Kn 19-20 (Jan 2013–Mar 2015)	1.001	6,584	0
Kn 20-21 (Mar 2015–Apr 2019)	1.001	11,096	0
Kn 21-22 (Apr 2019–Sep 2019)	1.001	8,017	0
Kn 22-23 (Sep 2019–Mar 2020)	1.001	10,430	0
Kn 23-24 (Mar 2020–Mar 2021)	1.001	8,154	0
Kn 24-25 (Mar 2021–Nov 2022)	1.001	12,487	0

No divergent transitions were observed in any model. City-deviation parameters (D, delta_city) exhibit poor mixing in transitions 18-19, 19-20, 20-21, and 23-24 due to the multiplicative non-identifiability described above; city-level transition estimates for these pairs should be interpreted with caution.

Representative Diagnostics are shown below.

Figure A7. Rank plots for the Kn 2122 (April 2019–September 2019) transition showing uniform distributions across chains, indicating good mixing.

 doi:10.20944/preprints202511.2148.v2

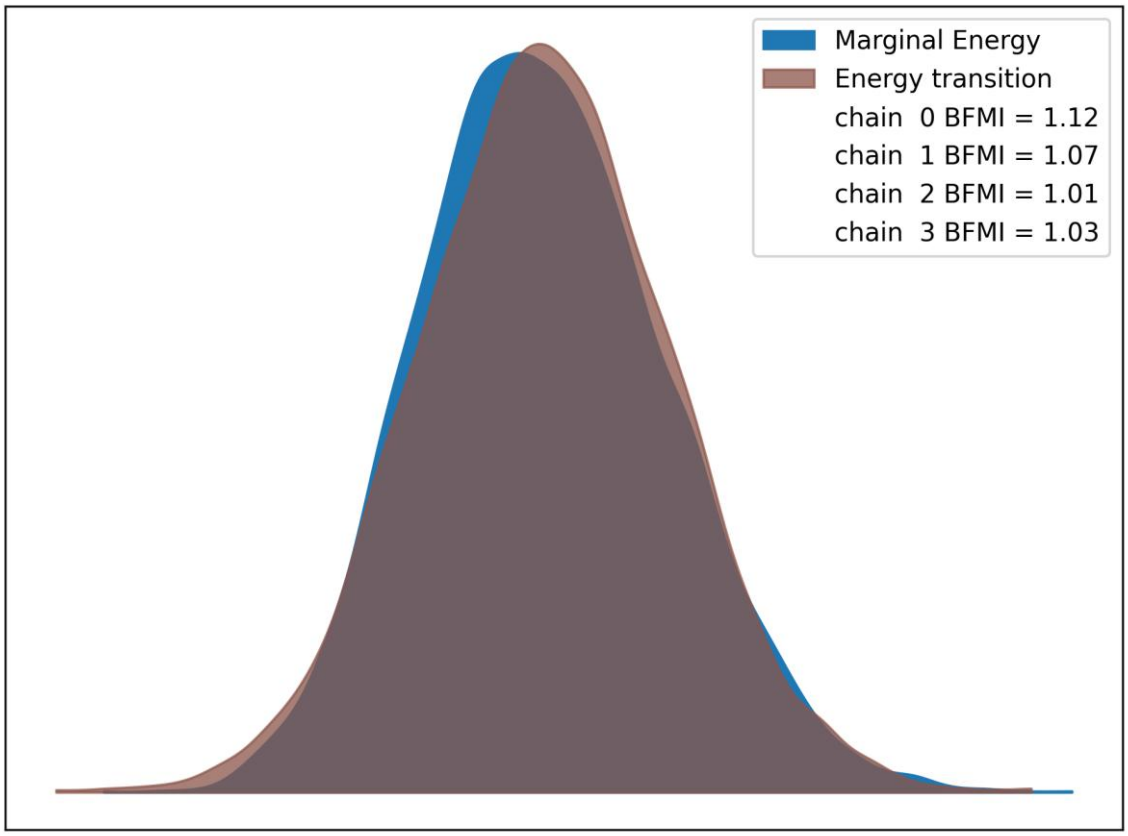

Figure A8. Energy plots showing no evidence of divergent transitions or geometric pathologies in the posterior.

 doi:10.20944/preprints202511.2148.v2

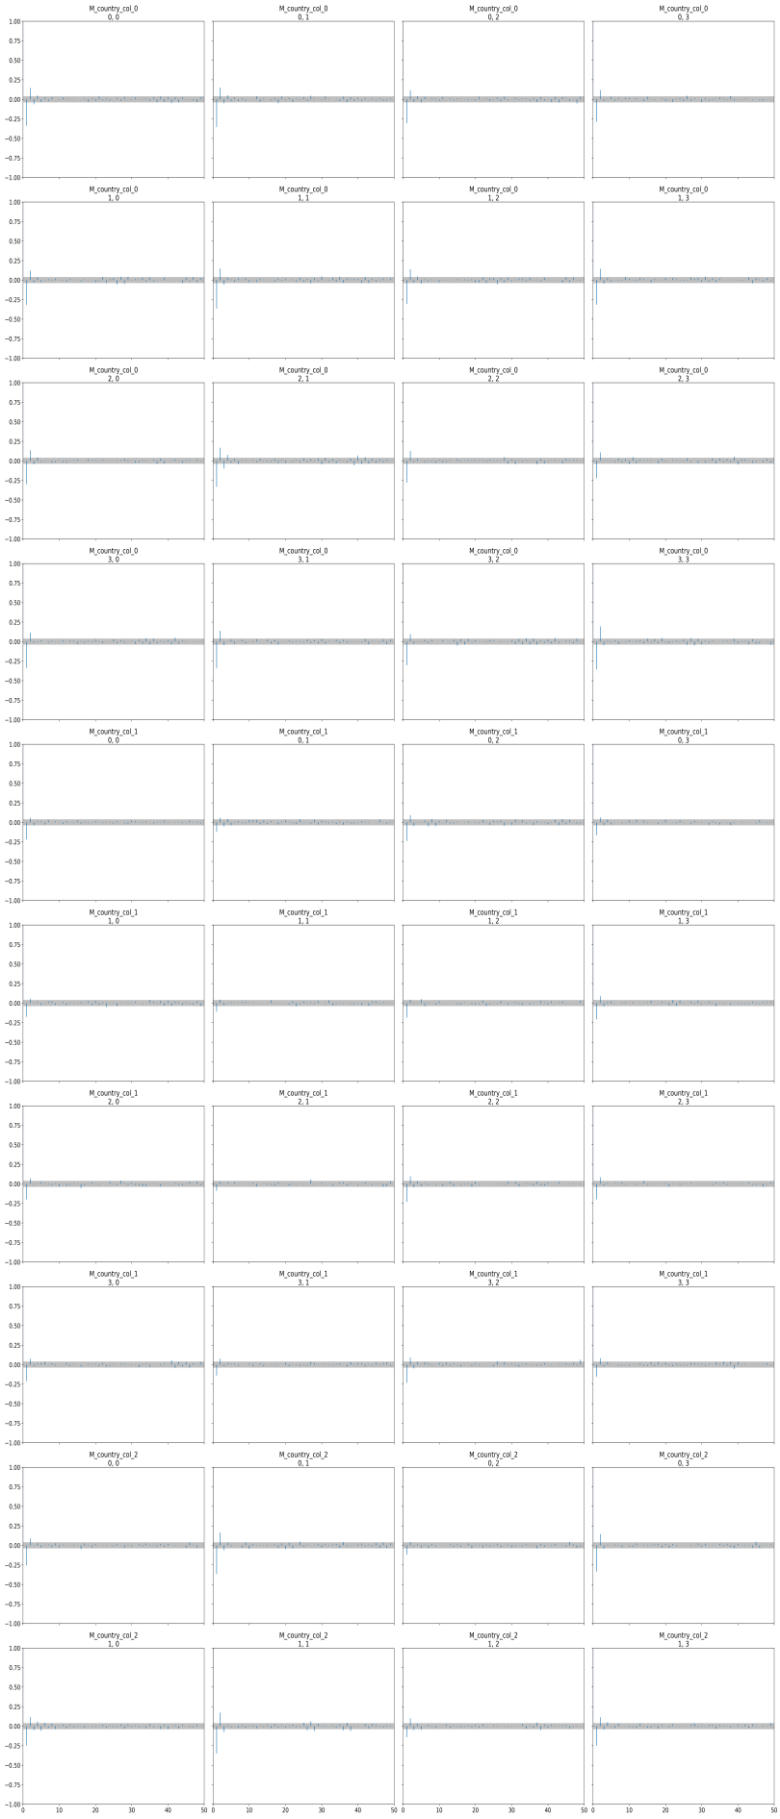

Figure A9. Autocorrelation plots demonstrating rapid decorrelation of MCMC samples for key model parameters.

 doi:10.20944/preprints202511.2148.v2

Appendix C Model Robustness Testing

To verify that the observed synchronization of voter transitions across cities represents genuine coordinated behavior rather than an artifact of the hierarchical Bayesian model structure, I tested three alternative model specifications with varying levels of flexibility for city-specific patterns.

Model Specifications

Original Hierarchical Model (baseline): The model described in the Methods section uses hierarchical pooling with moderately tight priors to stabilize estimates while allowing cities to deviate from national patterns. Key parameters: sigma_D = 0.3, delta_scale = 0.5, D_sigma = 0.3.

Relaxed Hierarchical Model: Same hierarchical structure but with substantially increased variability parameters to allow greater city-specific deviations: sigma_D = 0.8 (+167%), delta_scale = 1.5 (+200%), D_sigma = 0.8 (+167%). This specification tests whether tighter priors artificially constrain city-level patterns.

Independent Model: Each city is fitted separately without hierarchical pooling. The country-level transition matrix is estimated first, then used as a prior mean for independent city-level fits with sigma_city = 0.8. This specification completely removes structural constraints toward similarity between cities.

Results

Despite the increased flexibility allowing cities to diverge substantially from national patterns (248% increase in inter-city variability for the independent model), the synchronized drops and recoveries in Shas loyalty remained evident across all model specifications.

Figure C1 presents one panel per model, showing all cities together within each specification. The inter-city standard deviation (shown in gray shading) increases substantially in the relaxed and independent models, confirming that these specifications successfully allow greater divergence. Yet the temporal correlation across cities remains evident in all three panels.

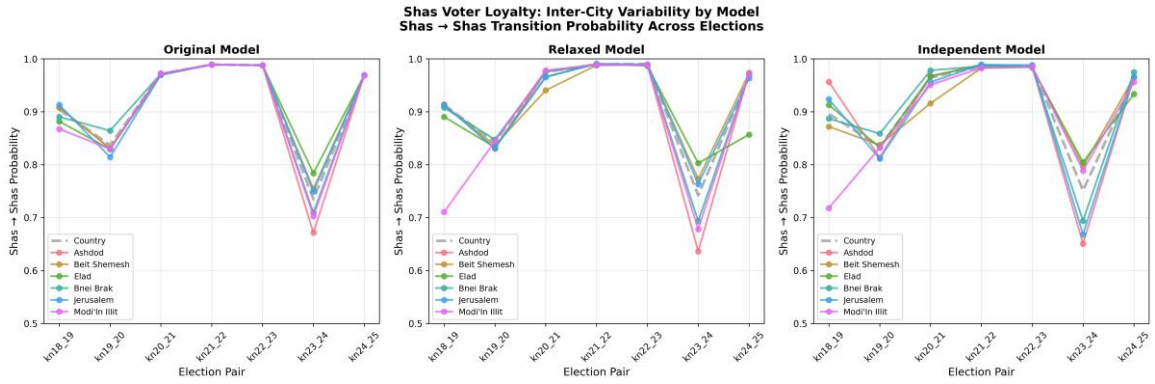

Figure A10. Shas→Shas retention rates by model specification. Each panel shows all cities within one model. Country estimates shown as thick gray dashed line; individual cities as colored solid lines. Gray shading indicates inter-city standard deviation. Despite increased flexibility in relaxed and independent models, synchronized patterns persist.

Quantitative Comparison

Mean Absolute Deviation (MAD) from country estimates across all transition pairs and electoral categories:

- Original hierarchical: MAD = 0.0044
- Relaxed hierarchical: MAD = 0.0116 (+166%)
- Independent: MAD = 0.0152 (+248%)

 doi:10.20944/preprints202511.2148.v2

The substantial increases in inter-city variability confirm that the alternative specifications successfully relax constraints. The persistence of synchronized transitions across all specifications demonstrates that the observed coordination reflects genuine features of the electoral data rather than modeling artifacts.

Full implementation details, including model code and configuration files, are available in the GitHub repository.

Appendix D—Corpus Evidence for the 23-24 Disruption

This appendix provides the full citations supporting the claims made in the "What Caused the 23-24 Disruption?" section of the Conclusions. Each citation includes the source outlet, article identifier, publication date, original URL, a verbatim Hebrew excerpt, and an English translation. The corpus comprises approximately 58,000 items scraped from Behadrey Haredim (bhol.co.il) and Kikar HaShabbat (kikar.co.il), covering March 2020 to March 2021. All URLs were accessed and archived between February 18 and March 1, 2026.

Cross-Ethnic Voter Recruitment Campaign

D1. Rabbi Mazuz instructed followers to vote UTJ on election morning, Knesset 23 (March 2020). > Source: Kikar HaShabbat | ID: kikar-349463 | Date: 2020-03-01 | URL: https://www.kikar.co.il/haredim-news/349463 > Quote: "רבות שמועות אחרי, של קולותיהם על ממושך וקרב הבוקר הורה מאזוז מאיר רבי הגאון (ראשון) להצביע יהדות התורה'ל'." > [After many rumors and a prolonged battle over the votes of the Kisse Rahamim community, Rabbi Meir Mazuz instructed this morning (Sunday) to vote for United Torah Judaism.]

D2. UTJ operatives engineered the Mazuz endorsement; Shas accused UTJ of "Ashkenazi interference." > Source: Kikar HaShabbat | ID: kikar-349484 | Date: 2020-03-01 | URL: https://www.kikar.co.il/haredim-news/349484 > Quote: "האשכנזית ההתערבות על זועמים ס"ש בתנועת הספרדי הציבור בתוך המתמשכת... הרב מאזוז הורה כי לכתחילה יש להצביע יהדות התורה'ל'." > [In the Shas movement they are furious about the ongoing Ashkenazi interference in the Sephardi public… Rabbi Mazuz ruled that ideally one should vote for United Torah Judaism.]

D3. Rabbi Yitzhak Barda also instructed to vote UTJ in Knesset 23. > Source: Kikar HaShabbat | ID: kikar-349484 | Date: 2020-03-01 | URL: https://www.kikar.co.il/haredim-news/349484 > Quote: "אשר וקרליץ אף רקמו את תמיכת הגאון רבי יצחק ברדא מאשקלון התורה ביהדות." > [Asher and Karlitz also arranged the support of Rabbi Yitzhak Barda of Ashkelon for United Torah Judaism.]

D4. Gafni publicly thanked Mazuz after the Knesset 23 results. > Source: Kikar HaShabbat | ID: kikar-349726 | Date: 2020-03-02 | URL: https://www.kikar.co.il/haredim-news/349726 > Quote: "בנאומו, הכהן והרב מאזוז לרב גפני הודה." > [In his speech, Gafni thanked Rabbi Mazuz and Rabbi HaCohen.]

D5. For Knesset 24, Mazuz split his endorsement—Haredi parties for Torah students, Smotrich for others. > Source: Behadrey Haredim | ID: bhol-article-1198514 | Date: 2021-03-20 | URL: https://www.bhol.co.il/news/1198514 > Quote: "תמיכה מאזוז מ"הגר העניק בהם הקודמים לסבבים בניגוד, התורה יהדות' למפלגת בלעדית בשיעורו, הערב השבועי מ"הגר הכריע לפצל את חלוקת הקולות ל-2 מפלגות." > [Unlike previous rounds in which Rabbi Mazuz gave exclusive support to United Torah Judaism, in his weekly lecture this evening he decided to split the vote between two parties.]

Source: Kikar HaShabbat | ID: kikar-387953 | Date: 2021-03-20 | URL: https://www.kikar.co.il/haredim-news/387953 Quote: "חרדית למפלגה ורק אך להצביע צריכים התורה בני, אחרות למפלגות להצביע שחושבים כאלו, תורה בני לא שהם אלו אבל סמוטריץ' ובן גביר." [Torah students must vote exclusively for a Haredi party, but those who are not Torah students, those thinking of voting for other parties, should vote for Smotrich and Ben Gvir.]

D6. Rabbi Barda instructed to vote Smotrich in Knesset 24. > Source: Behadrey Haredim | ID: bhol-article-1199123 | Date: 2021-03-22 | URL: https://www.bhol.co.il/news/1199123 > Quote: "אחרי, בצלאל של ט' להצביע היא שהמצווה להודיע לנכון ראיתי, מעמיקה בדיקה סמוטריץ'." > [After thorough

 doi:10.20944/preprints202511.2148.v2

examination, I saw fit to announce that the religious obligation is to vote Tet (ballot letter) for Bezalel Smotrich.]

D7. Ashkenazi yeshiva head instructed Sephardi students to vote UTJ, Knesset 24. > Source: Kikar HaShabbat | ID: kikar-388061 | Date: 2021-03-22 | URL: https://www.kikar.co.il/yeshiva-world/388061 > Quote: "בן' שבשם מי כל של חובה זו …התורה ליהדות להצביע חובה—עדה מאיזה משנה לא תורה' ייקרא." > [It doesn't matter which ethnic community you belong to—it is an obligation to vote for United Torah Judaism… This is the duty of everyone who bears the name "Torah scholar."]

Inter-Party Friction

D8. Gafni launched a campaign to recruit Sephardi voters from Shas. > Source: Behadrey Haredim | ID: bhol-article-1194886 | Date: 2021-03-10 | URL: https://www.bhol.co.il/news/1194886 > Quote: "אלו ,הספרדים התורה בני של תמיכתם את לגייס החליטו התורה ובדגל הבחירות לפני משבועיים פחות שותקים אך זועמים ס"בש. אשכנזים במוסדות לומדים שילדיהם." > [Less than two weeks before the elections, Degel HaTorah decided to recruit the support of Sephardi Torah scholars, those whose children study in Ashkenazi institutions. In Shas they are furious but silent.]

D9. Deri attacked Gafni for telling Sephardim to disregard Hakham Shalom Cohen. > Source: Kikar HaShabbat | ID: kikar-387209 | Date: 2021-03-10 | URL: https://www.kikar.co.il/haredim-news/387209 > Quote: "כהן שחכם מה ,כהן שלום לחכם תשמעו אל ויגיד ,ידידי והוא ,הכבוד כל עם גפני משה יבוא לכם אומר ,תורה דעת ברורה צרופה… אל תשמעו. לו למה?" > [Moshe Gafni will come, with all due respect—and he is my friend—and say: don't listen to Hakham Shalom Cohen, what Hakham Cohen tells you, a clear and pure Torah ruling… don't listen to him. Why?]

D10. Gafni called Shas a "racist party." > Source: Kikar HaShabbat | ID: kikar-387269 | Date: 2021-03-10 | URL: https://www.kikar.co.il/bchirot-2021/387269 > Quote: "אשכנזי שהוא אחד אף שם אין, שכמוהו חצוף. 'גזענים יותר החרדים' ואומר שלנו השונאים עם הולך הוא. ספרדים שם הכל." > [There isn't a single Ashkenazi there, they're all Sephardim. He goes to our enemies and says "the Haredim are more racist." What chutzpah.]

D11. Shas MKs counterattacked, noting UTJ's own voter losses to Smotrich. > Source: Kikar HaShabbat | ID: kikar-387269 | Date: 2021-03-10 | URL: https://www.kikar.co.il/bchirot-2021/387269 > Quote: "'התורה יהדות' מצביעי של קולות להחזיר ולנסות החרדיות המפלגות של הכח בהגדלת להתרכז שבמקום חבל ס"ש'ב. רפש לשלוח גפני כ"ח בוחר… 'סמוטריץ לטובת שנטשו." > [It's a shame that instead of focusing on strengthening the Haredi parties and trying to recover United Torah Judaism voters who defected to Smotrich… MK Gafni chooses to sling mud at Shas.]

D12. Hakham Shalom Cohen broke down crying, pleading Sephardim to remain loyal to Shas. > Source: Behadrey Haredim | ID: bhol-article-1197676 | Date: 2021-03-17 | URL: https://www.bhol.co.il/news/1197676 > Quote: "אנחנו הספרדים יש לנו הכרת טובה רק לתנועת ש"ס. לא לומדים אנחנו איפה משנה ס"לש רק להצביע ברורה בהוראה חייב ספרדי כל—." > [We Sephardim owe gratitude only to the Shas movement. It doesn't matter where we study—every Sephardi is obligated by a clear directive to vote only for Shas.]

COVID-19 as Enabling Factor

D13. Litzman resigned from the government over the holiday lockdown. > Source: Behadrey Haredim | ID: bhol-article-1137568 | Date: 2020-09-13 | URL: https://www.bhol.co.il/news/1137568 > Quote: "למעשים מאיומים: מהממשלה מתפטר הוא כי הודיע ליצמן יעקב השר בעקבות ה"בר סגר להטיל ההחלטה." > [From threats to action: Minister Yaakov Litzman announced he is resigning from the government following the decision to impose a lockdown on Rosh Hashanah.]

D14. Kikar HaShabbat survey: 31.7% of Haredi voters said COVID would affect their vote. > Source: Kikar HaShabbat | ID: kikar-386120 | Date: 2021-02-23 | URL: https://www.kikar.co.il/haredim-news/386120 > Quote: "60.02% מהנשאלים כי השיבו משבר הקורונה לא ישפיע על ההצבעה… 31.7% 60.02%] < ." מהנשאלים השיבו כי הקורונה כן תשפיע of respondents said the COVID crisis would not affect their vote… 31.7% of respondents said COVID would affect it.]

 doi:10.20944/preprints202511.2148.v2

D15. Shas rated significantly better than UTJ on COVID handling across all sub-sectors. > Source: Kikar HaShabbat | ID: kikar-386120 | Date: 2021-02-23 | URL: https://www.kikar.co.il/haredim-news/386120 > Quote: "גם החסידים והליטאים סבורים כי חברי הכנסת של ש"ס התנהלו טוב יותר מאשר חברי הכנסת של יהדות התורה במשבר הקורונה." > [Even the Hasidim and Lithuanians believe that Shas MKs handled the COVID crisis better than United Torah Judaism MKs.]

Forum Evidence

D16. Forum thread: thousands of Haredim debating whether to vote Smotrich. > Source: Behadrey Haredim Forums | ID: bhol-forum-3172808 | Date: 2021-01-18 | URL: https://forums.bhol.co.il/forums/topic.asp?cat_id=4&topic_id=3172808&forum_id=771 > Quote: "בליבם מערערים חרדיות למפלגות בעבר שהצביעו חרדים אלפי אלה בימים: את לתת הבחינות מכל ונכון ראוי לא האם סמוטריץ לבצלאל קולם?" > [These days thousands of Haredim who previously voted for Haredi parties are questioning in their hearts: is it not right and proper in every respect to give their vote to Bezalel Smotrich?]

D17. Forum users credited Smotrich with advocating for Haredi interests during COVID. > Source: Behadrey Haredim Forums | ID: bhol-forum-3172808 | Date: 2021-01-18 | URL: https://forums.bhol.co.il/forums/topic.asp?cat_id=4&topic_id=3172808&forum_id=771 > Quote: "ביטא ואף ,שמע גם שמע הוא קולנו את אבל .22.00 בשעה הבחירות יום במוצאי קולנו את ספר לא אמנם בצלאל זעקתנו את ברמה." > [Bezalel may not have counted our votes on election night at 10 PM. But our voice he heard loud and clear, and he expressed our outcry at the highest level.]

References

1. Andersen JG and Yaish M (2003) Social cleavages, electoral reform and party choice: Israel's 'natural' experiment. *Electoral Studies* 22(3): 399–423. https://doi.org/10.1016/S0261-3794(02)00006-9.
2. Bartolini S and Mair P (1990) Identity, Competition, and Electoral Availability: The Stabilisation of European Electorates 1885–1985. Cambridge: Cambridge University Press.
3. Baumgartner FR and Jones BD (1993) *Agendas and Instability in American Politics.* Chicago: University of Chicago Press.
4. Blum A (2007) Rubinstein's mistake [הטעות של רובינשטיין]. *Behadrey Haredim*, 9 November. Available at: https://www.bhol.co.il/news/66534 (accessed 5 March 2026). [Hebrew]
5. Brown PJ and Payne CD (1986) Aggregate data, ecological regression and voting transitions. *Journal of the American Statistical Association* 81(394): 453–460. https://doi.org/10.1080/01621459.1986.10478293.
6. Campbell DE, Green JC and Layman GC (2011) The party faithful: Partisan images, candidate religion, and the electoral impact of party identification. *American Journal of Political Science* 55(1): 42–58. https://doi.org/10.1111/j.1540-5907.2010.00474.x.
7. Cincotta R (2013) Government without the ultra-Orthodox? *Demography and the Future of Israeli Politics.* FPRI eNotes, December. Philadelphia: Foreign Policy Research Institute.
8. Clarke HD, Sanders D, Stewart MC et al. (2004) *Political Choice in Britain.* Oxford: Oxford University Press.
9. Curiel CP and Zeedan R (2024) Social identity and voting behavior in a deeply divided society: The case of Israel. *Societies* 14(9): 177. https://doi.org/10.3390/soc14090177.
10. Dödtmann A (2022) Haredi fundamentalism in the State of Israel. *Religions* 13(12): 1149. https://doi.org/10.3390/rel13121149.
11. Forcina A and Pellegrino D (2019) Estimation of voter transitions and the ecological fallacy. *Quality & Quantity* 53: 1859–1874. https://doi.org/10.1007/s11135-019-00845-1.
12. Freedman M (2020) Vote with your rabbi: The electoral effects of religious institutions in Israel. *Electoral Studies* 67: 102241. https://doi.org/10.1016/j.electstud.2020.102241.
13. Gelman A and Hill J (2007) *Data Analysis Using Regression and Multilevel/Hierarchical Models.* Cambridge: Cambridge University Press.
14. Gidron N, Sheffer L and Mor G (2022) Validating the feeling thermometer as a measure of partisan affect in multi-party systems. *Electoral Studies* 80: 102542. https://doi.org/10.1016/j.electstud.2022.102542.

 doi:10.20944/preprints202511.2148.v2

15. Glynn AN and Wakefield J (2010) Ecological inference in the social sciences. *Statistical Methodology* 7(3): 307–322. https://doi.org/10.1016/j.stamet.2009.09.003.
16. Goodman LA (1953) Ecological regressions and the behavior of individuals. *American Sociological Review* 18(6): 663–664. https://doi.org/10.2307/2088111.
17. Goodwin MJ, Kaufmann E and Larsen EG (2022) Asymmetric realignment: Immigration and right party voting. *Electoral Studies* 80: 102551. https://doi.org/10.1016/j.electstud.2022.102551.
18. Gorelik B (2025) Ethnic divisions within unity: Insights into intra-group segregation from Israel's ultra-Orthodox society. *Social Sciences* 14(3): 169. https://doi.org/10.3390/socsci14030169.
19. Greiner J and Quinn K (2010) R×C ecological inference: Bounds, correlations, flexibility and transparency of assumptions. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 173(1): 55–81. https://doi.org/10.1111/j.1467-985X.2009.00606.x.
20. Hamza A and Maeda K (2025) Beyond binary: A new measure of party ethnicization. *Party Politics*. Epub ahead of print. https://doi.org/10.1177/13540688251376485.
21. Horowitz DL (1985) *Ethnic Groups in Conflict.* Berkeley: University of California Press.
22. Keren-Kratz M (2025) The turnaround in Israel's Haredi society in the late 20th century: A data-based analysis. *Religions* 16(2): 145. https://doi.org/10.3390/rel16020145.
23. King G (1997) A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data. Princeton, NJ: Princeton University Press.
24. Lanzara G, Lazzaroni S, Masella P et al. (2024) Do bishops matter for politics? Evidence from Italy. *Journal of Public Economics* 238: 105177. https://doi.org/10.1016/j.jpubeco.2024.105177.
25. Leon N (2014) Mizrachi ultra-Orthodoxy: Strict ideology, liquid identity. *Journal for the Study of Haredi Society* 1: 1–20. [Hebrew]
26. Lipset SM and Rokkan S (1967) Party Systems and Voter Alignments: Cross-National Perspectives. New York: Free Press.
27. Mainwaring S, Gervasoni C and España-Najera A (2017) Extra- and within-system electoral volatility. *Party Politics* 23(6): 623–635. https://doi.org/10.1177/1354068815625229.
28. Mainwaring S and Zoco E (2007) Political sequences and the stabilization of interparty competition: Electoral volatility in old and new democracies. *Party Politics* 13(2): 155–178. https://doi.org/10.1177/1354068807073852.
29. Martin N, de Lange SL and van der Brug W (2022) Holding on to voters in volatile times: Bonding voters through party links with civil society. *Party Politics* 28(2): 354–364. https://doi.org/10.1177/1354068820980304.
30. Malach G and Cahaner L (2025) *Statistical Report on Ultra-Orthodox Society in Israel 2024.* Jerusalem: Israel Democracy Institute.
31. Pedersen MN (1979) The dynamics of European party systems: Changing patterns of electoral volatility. *European Journal of Political Research* 7(1): 1–26. https://doi.org/10.1111/j.1475-6765.1979.tb01267.x.
32. Puig X and Ginebra J (2015) Ecological inference and spatial variation of individual behavior: National divide and elections in Catalonia. *Geographical Analysis* 47(3): 262–283. https://doi.org/10.1111/gean.12056.
33. Robinson WS (1950) Ecological correlations and the behavior of individuals. *American Sociological Review* 15(3): 351–357. https://doi.org/10.2307/2087176.
34. Romero R, Pavía JM, Martín J et al. (2020) Assessing uncertainty of voter transitions estimated from aggregated data. Application to the 2017 French presidential election. *Journal of Applied Statistics* 47(13-15): 2711–2736. https://doi.org/10.1080/02664763.2020.1804842.
35. Rosen O, Jiang W, King G et al. (2001) Bayesian and frequentist inference for ecological inference: The R×C case. *Statistica Neerlandica* 55(2): 134–156. https://doi.org/10.1111/1467-9574.00162.
36. Seawright J and Gerring J (2008) Case selection techniques in case study research: A menu of qualitative and quantitative options. *Political Research Quarterly* 61(2): 294–308. https://doi.org/10.1177/1065912907313077.
37. Simas EN and Lothamer L (2025) Not just who, but how: Further probing the connection between primary election dissatisfaction and general election voting behavior. *Electoral Studies* 88: 102969. https://doi.org/10.1016/j.electstud.2025.102969.

 doi:10.20944/preprints202511.2148.v2

38. Wakefield J (2004) Ecological inference for 2×2 tables (with discussion). *Journal of the Royal Statistical Society: Series A* 167(3): 385–445. https://doi.org/10.1111/j.1467-985x.2004.02046.x.
39. Ward P and Denney S (2025) Partisan voters in party systems with ephemeral parties: Evidence from South Korea. *Party Politics*. Epub ahead of print. https://doi.org/10.1177/13540688251339976.
40. Zur R and Bakker R (2025) The Israeli parties' positions in comparative perspective. *Party Politics* 31(2): 323–334. https://doi.org/10.1177/13540688231218917.