

Governing Through Narrative: Text-as-Data Evidence of Spin Dictatorship in Bukele's El Salvador

Course Project — Text as Data, Spring 2026

Abstract

This paper examines whether Nayib Bukele's El Salvador exhibits communication patterns consistent with “spin dictatorship” theory: the hypothesis that contemporary autocrats govern primarily through narrative control rather than overt repression. Drawing on a novel corpus of 139,637 news articles from more than 20 Salvadoran and international outlets (2015–2025) and 161,788 tweets from five official government accounts, I apply log-odds ratio analysis, cosine similarity measures, and event-framing comparisons to trace how the Bukele administration has constructed and coordinated its public narrative. I find three patterns consistent with spin dictatorship: (1) government accounts underwent significant lexical convergence following the institutional capture of May 2021, with the Asamblea Legislativa’s cosine similarity to @nayibbukele increasing by 52 percent; (2) government messaging during the Estado de Excepción (2022–present) systematically avoids the vocabulary of rights and accountability that dominates independent media coverage of the same period; and (3) the administration’s distinctive bigrams reveal a stylized “security success” frame that crowds out alternative interpretations of the mass incarceration policy. These findings suggest that institutional capture in El Salvador has extended beyond formal power to encompass the communicative function of nominally independent state bodies.

Keywords: spin dictatorship, El Salvador, text-as-data, authoritarianism, narrative control, computational social science

1. Introduction

On May 1, 2021, the first day that Bukele’s Nuevas Ideas party controlled a legislative supermajority, the newly seated Asamblea Legislativa voted to dismiss El Salvador’s Constitutional Court and Attorney General within hours of convening. The move was condemned by the United States, the European Union, and international human rights organizations as a constitutional coup. It was livestreamed on Twitter and accompanied by a government messaging campaign framing the dismissals as anti-corruption housecleaning (Human Rights Watch, 2022a). Within weeks, the replaced officials had begun posting in a register strikingly similar to the president’s own account.

This paper asks whether that observation (institutions speaking with one voice after capture) is a systematic feature of Bukele’s communications strategy or an artifact of normal political coordination. The question matters because scholars of democratic backsliding have increasingly emphasized that the most durable contemporary autocrats consolidate power not through violent repression but through information control (Guriev and Treisman, 2022). On this “spin dictatorship” account, the real substrate of authoritarian power is not the truncheon but the narrative: the ability to shape what information citizens receive and how they interpret it.

El Salvador is a theoretically important case. Bukele was elected in 2019 with genuine popular support, promising to break the ARENA–FMLN duopoly that had governed the country since the end of its civil war. By conventional indicators (Freedom House scores, V-Dem indices, press freedom rankings) El Salvador has experienced sharp democratic backsliding under his leadership (Freedom House, 2024; V-Dem Institute, 2024). Yet his approval ratings have consistently exceeded 80 percent, and he won re-election in February 2024 with 85 percent of the vote. This combination (institutional erosion alongside popular legitimacy) is precisely the pattern Guriev and Treisman (2022) identify as characteristic of spin autocrats.

Despite its theoretical salience, El Salvador under Bukele has received limited systematic empirical treatment in the political science literature. The primary language of political communication is Spanish, and the relevant texts (official communications, news archives, social media) have not previously been assembled into an analyzable corpus. This paper makes two contributions. First, it introduces a large-scale, multi-source corpus of Salvadoran political text covering 2015–2025, combining government Twitter data, domestic news archives, and international press coverage filtered for political relevance. Second, it uses computational text analysis to trace the communicative signatures of institutional capture and narrative management that theory predicts but that have not been empirically documented for this case.

The remainder of the paper proceeds as follows. Section 2 reviews the theoretical literature on spin dictatorship, competitive authoritarianism, and the El Salvador context. Section 3 describes the data collection and corpus construction. Section 4 presents the analytical methods. Section 5 reports the empirical results. Section 6 discusses implications and limitations, and Section 7 concludes.

2. Theoretical Background

2.1 Spin Dictatorship and Informational Authoritarianism

Guriev and Treisman (2019, 2022) distinguish two ideal types of autocratic rule. “Fear dictators” maintain power through violence and surveillance; the population obeys because the costs of resistance are lethal. “Spin dictators” maintain power through information manipulation; the population acquiesces because it is persuaded (or at least left uncertain) that the leader is competent, popular, and indispensable. Modern spin autocrats are associated with competitive elections they reliably win, high approval ratings, media environments that are technically pluralistic but substantively captured, and a preference for legal mechanisms of control over extra-legal ones (Guriev and Treisman, 2022, pp. 14–22).

Several predictions follow for communications behavior. First, spin regimes should exhibit coordinated messaging: as institutions are captured, their communications should converge toward the president's preferred framing. Second, spin regimes should invest heavily in positive agenda-setting during crises, foregrounding security gains, economic metrics, or national pride while avoiding the vocabulary associated with accountability and rights. Third, spin regimes should target international critics rhetorically rather than physically, using social media to delegitimize rather than to threaten (Gunitsky, 2015).

Related work on “informational autocracy” (Guriev and Treisman, 2019) emphasizes that spin regimes are particularly prevalent in middle-income countries with educated populations and access to foreign information. In such settings, overt censorship is counterproductive; subtler narrative management is the equilibrium strategy. El Salvador, with a large diaspora, significant U.S. media exposure, and a digitally active urban population, fits this profile.

2.2 Competitive Authoritarianism and Institutional Capture

Levitsky and Way (2010), building on the electoral authoritarianism framework of Schedler (2006), define competitive authoritarian regimes as polities in which formal democratic institutions exist and are regularly contested, but where incumbents systematically violate the rules such that the competition is not meaningfully fair. The key mechanism is the abuse of state resources (legal, financial, and communicative) to disadvantage opponents and extend incumbency.

Bermeo (2016) identifies “executive aggrandizement” (the incremental accumulation of power through formally legal acts) as the dominant pathway of democratic backsliding in the contemporary period. The May 2021 dismissal of El Salvador's Constitutional Court judges exemplifies this pattern: the act was formally legislative, carried out by an elected majority, but violated the constitutional requirement that the court serve fixed terms (IACHR, 2021).

The communications dimension of competitive authoritarianism has received less systematic attention. Gunitsky (2015) argues that social media platforms, once expected to democratize information flows, have in practice become tools for authoritarian counter-mobilization and narrative management. Autocrats use official accounts to flood the information environment, undermine trust in independent media, and build direct relationships with supporters that bypass journalistic intermediaries.

2.3 El Salvador: Context, Genuine Appeal, and Performative Governance

Understanding Bukele requires understanding what came before him. For three decades after the 1992 peace accords, El Salvador alternated between the right-wing ARENA and the left-wing FMLN. By the mid-2010s both parties were widely associated with corruption, clientelism, and failure to address the country's most urgent problem: gang violence. In 2015, El Salvador recorded 103 homicides per 100,000 inhabitants, making it the most violent country in the Western Hemisphere (Statista, 2024). Communities in gang-controlled neighborhoods lived under systematic extortion and the constant threat of violence. The political establishment had no credible answer.

Bukele ran in 2019 as a genuine outsider. A former publicist and then mayor of San Salvador, he was expelled from the FMLN and formed his own movement, Nuevas Ideas. He won with 53 percent of the first-round vote, the first candidate in El Salvador's democratic history to win without a runoff. His appeal was not primarily ideological; it was anti-institutional. He promised to dismantle a corrupt system that had failed ordinary Salvadorans (Wolf, 2017). That promise resonated.

The security improvements under his administration are real. Following the Estado de Excepción declared in March 2022, the government reported a homicide rate of approximately 2 per 100,000 in 2024, down from 53 per 100,000 in 2018 (Statista, 2024). Many Salvadorans who could not previously walk in their own neighborhoods describe a transformation in daily life.

Whatever the methodological caveats around official homicide statistics (El Faro, 2023), the perceived improvement in public safety is politically consequential and genuinely experienced.

Bukele has weaponized this reality rhetorically. When human rights organizations, foreign governments, or international journalists criticize conditions in the prisons or the suspension of due process, he frames them as wealthy elites, removed from the reality of gang violence, who prioritize the rights of criminals over the safety of working-class communities. This framing has been effective: his approval ratings have not meaningfully declined despite sustained international criticism. The argument lands because it contains a grain of truth. The populations most affected by gang extortion were poor and largely ignored by the previous political establishment.

What distinguishes Bukele from a traditional security-focused leader is his mastery of spectacle. A former publicist, he self-labeled “the world’s coolest dictator” on social media and deployed slick propaganda videos to promote CECOT, the megaprison holding 40,000 detainees, as a symbol of state power. Bitcoin City was announced in a theatrical seaside event with pyrotechnics. His White House visits were choreographed as cinematic productions. Governance, for Bukele, is inseparable from performance: the message is as carefully produced as the policy. This is precisely the communicative mode spin dictatorship theory predicts.

3. Data

3.1 *Government Twitter Corpus*

The primary data on government communications consist of 161,788 tweets collected from five official Salvadoran government accounts via the twitterapi.io full-archive search API: @AsambleaSV (54,466 tweets; the Asamblea Legislativa), @FGR_SV (36,519; the Fiscalía General de la República), @Gobierno_SV (33,370; Gobierno de El Salvador), @PresidenciaSV (30,438; Casa Presidencial), and @nayibbukele (6,995; Nayib Bukele’s personal account). Collection was conducted in February–March 2026, paginating through the full available timeline for each account. The corpus spans 2015–2026, though coverage before 2018 is thinner for accounts created later.

The five accounts represent distinct institutional roles. @AsambleaSV covers the legislative branch; @FGR_SV the public prosecutor; @Gobierno_SV and @PresidenciaSV are executive communications accounts; @nayibbukele is the president’s personal account, used as his primary public communications channel. The institutional capture of May 2021 replaced the leadership of the Asamblea and FGR, making these accounts a natural quasi-experimental site: messaging before and after the leadership change should differ if capture induces communicative convergence.

3.2 *News Article Corpus*

The news article corpus comprises 139,637 articles drawn from six source streams: (1) HuggingFace Salvadoran news datasets (Martinez, 2023), contributing 85,912 articles primarily from El Mundo (diario.elmundo.sv) and El Diario de Hoy (elsalvador.com); (2) direct scrapes of additional outlets including Revista Factum, Foco STV, and the GDELT-linked Salvadoran press (17,629 articles); (3) the original article collection via sitemaps and GDELT BigQuery (8,815 articles); (4) repaired articles recovered from broken URLs via HTTP redirect-following and Wayback Machine CDX lookup (8,157 articles); (5) new outlet scrapes (640 articles); and (6)

internationally-sourced articles filtered for political relevance to El Salvador (30,849 articles from Reuters, Deutsche Welle, The Guardian, Proceso.hn, Univision, and others).

Domestic outlets cover the full political spectrum. El Faro and Revista Factum are critical investigative outlets; La Página and El Mundo range from center to pro-government; La Prensa Gráfica is establishment conservative. International sources were filtered using a content relevance test requiring at minimum three mentions of “El Salvador” or two mentions of “Bukele” combined with key political terms, to exclude articles about Salvadorans abroad that are unrelated to domestic politics.

The corpus covers 2016–2025. The median article length is 386 words; the mean is 508 words. Article volume is highest in 2018–2022, with strong 2025 coverage owing to international attention on Bukele’s second term and the ongoing Estado de Excepción.

3.3 Data Quality and Limitations

Several limitations warrant acknowledgment. Twitter data collection was budget-constrained (\$30 API credit), limiting retrievable tweets per account; @nayibbukele coverage is particularly truncated relative to his full tweet history. La Prensa Gráfica is partially paywalled, reducing article counts. Gatoencerrado.news, an independent investigative outlet, could not be scraped due to bot protection. Our term frequency analyses treat government accounts as a unified corpus in some analyses, potentially masking within-government variation. Finally, the corpus captures text production but not consumption: we cannot observe which audiences receive which messages or how they interpret them.

4. Methods

The analyses below apply standard text-as-data methods (Grimmer and Stewart, 2013) to the government tweet and news article corpora described in Section 3.

4.1 Term Frequency Analysis

To trace the temporal evolution of government messaging priorities, we compute the rate per 1,000 tweets at which key political terms appear in government tweets by year. Terms are selected based on their theoretical salience to the spin dictatorship hypothesis and verified against the corpus vocabulary. We examine security-related terms (pandillas, excepción, seguridad, terroris-), Bitcoin-related terms, and accountability-related terms (derechos, corrupción, democracia). Rate normalization accounts for variation in annual tweet volume.

4.2 Log-Odds Ratio for Distinctive Vocabulary

To identify vocabulary that distinguishes individual accounts and outlets from one another, we apply the log-odds ratio with Dirichlet prior smoothing (Monroe, Colaresi, and Quinn, 2008). For each term w and source s , the log-odds is computed as the difference between the logit of term frequency in source s and the logit of term frequency in all other sources combined. A smoothing constant of 0.5 per term per source prevents extreme values for rare terms. We report the top 15 most distinctive terms per source.

4.3 Cosine Similarity for Institutional Convergence

To measure linguistic similarity between institutional accounts over time, we construct term frequency vectors over the unigram vocabulary (after removing stop words, URLs, and account mentions) and compute pairwise cosine similarity: $\cos(A, B) = (A \cdot B) / (\|A\| \times \|B\|)$. We compare each institutional account to @nayibbukele before and after May 1, 2021. The before period includes all tweets through April 30, 2021; the after period runs from May 1, 2021 through the collection date. An increase in cosine similarity is interpreted as evidence of communicative convergence following institutional capture.

4.4 Event Framing Windows

To compare government and media framing of key events, we construct text windows of plus or minus 30 days around six politically significant dates. Within each window, we build separate term-frequency vectors for government tweets and media articles, weight by TF-IDF, and visualize the most distinctive terms using word clouds. Events examined include Bukele's inauguration (June 2019), the COVID Emergency declaration (March 2020), the May 2021 institutional takeover, the Bitcoin legal tender law (September 2021), the declaration of the Estado de Excepción (March 2022), and the opening of CECOT (2023).

Figure 1. Word Clouds: COVID Emergency Decree (March 2020, ± 30 days)






Figure 2. Word Clouds: CECOT Mega-Prison Opens (November 2023, ±30 days)

Government tweets (red)



Media articles (blue)

5. Results

5.1 Narrative Evolution in Government Communication

Table 1 presents the rate per 1,000 tweets at which key terms appear in government communications by year (2019–2025). Three patterns stand out.

First, the Estado de Excepción (declared March 2022) produced the sharpest single-period shift in government messaging in the corpus. The term *pandill-* (gang-related) rises from 12.5 per 1,000 tweets in 2021 to 44.2 in 2022 and 76.3 in 2023. The term *excepci-* rises from 1.8 per 1,000 in 2021 to 24.7 in 2022. These trends are consistent with the administration prioritizing a security-success frame following the emergency declaration.

Second, Bitcoin-related vocabulary exhibits a sharp pulse. The term *bitcoin* registers 24.7 per 1,000 tweets in 2021 (the year of its adoption as legal tender) before declining precipitously to 4.7 in 2022 and 0.5 in 2023 and 2024. This pattern suggests strategic rather than sustained emphasis: the administration heavily publicized the policy at launch, then allowed coverage to fade as international criticism mounted and the IMF expressed concern (International Monetary Fund, 2021).

Third, accountability language remains flat or declines across the period. *Democracia* registers between 1.2 and 3.3 per 1,000 tweets throughout, with no clear trend. *Corrupci-* spikes in 2021 (16.9 per 1,000), consistent with framing institutional capture as anti-corruption reform, then falls in subsequent years (10.6 in 2024).

Table 1. Key Term Frequency per 1,000 Government Tweets by Year

Term	2019	2020	2021	2022	2023	2024	2025
<i>pandill-</i>	9.6	14.0	12.5	44.2	76.3	51.9	86.5
<i>excepci-</i>	0.7	3.7	1.8	24.7	29.8	27.5	16.2
<i>seguridad</i>	41.9	38.1	44.9	57.0	66.9	75.8	48.5
<i>bitcoin</i>	0.0	0.0	24.7	4.7	0.5	0.5	2.4
<i>derechos</i>	16.6	14.5	11.4	16.8	14.1	10.0	2.8
<i>corrupci-</i>	11.9	3.9	16.9	14.5	24.7	10.6	14.6
<i>democracia</i>	2.0	1.4	2.9	1.4	2.0	3.3	1.2

N tweets	13,707	17,278	13,767	16,250	10,262	7,572	2,473
----------	--------	--------	--------	--------	--------	-------	-------

Note. All five government accounts combined. Terms matched as substrings (e.g., pandill- matches pandilla, pandillas, pandillero).

5.2 Institutional Messaging Convergence After May 2021

Table 2 reports cosine similarity between each institutional Twitter account and @nayibbukele before and after the May 2021 institutional capture, computed over unigram frequency vectors after stop word removal.

Table 2. Cosine Similarity to @nayibbukele Before and After May 2021

Account	Before May 2021	After May 2021	Change (%)
@AsambleaSV	0.293	0.445	+51.9%
@FGR_SV	0.253	0.319	+26.1%
@PresidenciaSV	0.404	0.613	+51.7%
@Gobierno_SV	0.540	0.698	+29.3%

All four accounts show increased linguistic similarity to @nayibbukele after May 2021. The two accounts whose leadership was most directly replaced by Bukele loyalists (@AsambleaSV and @PresidenciaSV) show the largest increases (both approximately +52%). @FGR_SV and @Gobierno_SV show smaller but still substantial increases (+26% and +29%, respectively).

The pattern is consistent with spin dictatorship theory: institutional capture produces not merely administrative alignment but communicative alignment. The captured Asamblea and PresidenciaSV accounts adopt the president’s vocabulary, tone, and framing. This convergence suggests either active messaging coordination across accounts or self-selection of personnel who internalize the administration’s communicative norms.

The FGR’s smaller convergence is also informative. The Fiscalía is a prosecutorial institution whose public communications necessarily include legal terminology, case descriptions, and procedural language that differs structurally from a political account. That it

nonetheless converged by 26 percent suggests the effect extends even to institutionally constrained accounts.

5.3 Framing Divergence: *The Estado de Excepción*

Table 3 compares the rate per 1,000 units at which key terms appear in government tweets versus domestic media articles for the 2022–2024 period, the *Estado de Excepción* window.

Table 3. Term Frequency per 1,000 Units, Government vs. Domestic Media (2022–2024)

Term	Gov. Tweets / 1k	Media Articles / 1k	Ratio (Gov:Med)
pandill-	92.0	119.4	0.77
seguridad	64.2	195.2	0.33
excepci-	26.8	118.6	0.23
derechos	14.5	158.9	0.09
capturado	20.7	60.7	0.34
terroris-	23.4	37.2	0.63
arbitrar-	0.9	41.2	0.02
inocente	0.9	23.9	0.04
tortur-	0.2	19.1	0.01
hacinam-	0.0	4.0	0.00

Note. Government tweets $n = 34,084$; domestic media articles $n = 14,101$. International sources excluded.

Two features of this table are theoretically significant. First, government tweets are not simply silent on the *Estado de Excepción*: they actively use the terminology of gang violence and security (pandill-, terroris-, excepci-). The security frame is actively constructed, not merely absent. Second, the vocabulary of accountability and due process (derechos, arbitrar-, inocente, tortur-, hacinam-) is systematically underrepresented in government tweets relative to media coverage. The ratio for arbitrar- is 0.02: government tweets use this term at 2 percent the rate of independent media coverage of the same period. For tortur- the ratio falls to 0.01.

This pattern cannot be explained by the structural difference between tweets and long-form articles alone. Both rates are normalized per 1,000 tokens, which controls for format-driven length differences: government tweets average roughly 25 tokens each, while media articles

average 386 words. Under this normalization, a term that appears once in an average tweet would register approximately 40 per 1,000 -- well above what we observe for accountability terms in government tweets. That *pandillas* and *seguridad* appear at comparable or higher per-1,000 rates in government tweets than in media articles confirms that short format is not the binding constraint. The administration actively engages the security topic at high frequency; the avoidance is selective, targeting precisely the sub-vocabulary that would invite accountability.

5.4 Distinctive Vocabulary by Source

Log-odds ratio analysis reveals systematic differences in the vocabulary most distinctive to each government account and each media outlet. Among government accounts, @nayibbukele is most distinctively associated with terms of international positioning: *mundo* (world), *global*, and country names appear alongside economic framing terms. @AsambleaSV (post-2021) is most distinctive for legislative procedure framing (*sesión*, *diputados*, *aprovó*), language that normalizes institutional capture as routine lawmaking. @FGR_SV's most distinctive terms cluster around criminal prosecution (*detenidos*, *imputados*, *fiscales*), a frame that presents mass arrests as orderly due-process events.

Among media outlets, El Faro's distinctive vocabulary includes *constitución*, *derechos humanos*, *arbitraria*, and *víctimas*: the accountability lexicon largely absent from government communications. La Página's distinctive terms skew toward crime and security reporting. Diario El Mundo's vocabulary overlaps substantially with government framing terms, consistent with its generally sympathetic editorial orientation toward the administration.

6. Discussion

The three empirical patterns documented here (narrative evolution, institutional convergence, and framing divergence) collectively suggest that El Salvador under Bukele exhibits the communicative hallmarks of spin dictatorship. The administration has not merely changed policy; it has systematically reorganized the information environment in ways consistent with the theoretical predictions.

The institutional convergence finding (Table 2) is the most novel contribution. While scholars have documented formal institutional capture in El Salvador, the communicative dimension of that capture has not been previously quantified. Our analysis shows that institutional capture produces a measurable shift in the language that captured institutions use. Communicative convergence may be a more sensitive early-warning indicator of authoritarian consolidation than formal institutional change: messaging can converge before formal structures are fully aligned.

The framing divergence findings (Table 3) are consistent with a large body of qualitative reporting from Salvadoran journalists and human rights organizations. What the computational analysis adds is scale and systematicity: the gap holds across 34,000 government tweets and 14,000 media articles and is asymmetric in theoretically meaningful ways. The government does engage the *Estado de Excepción* topic but does so in a vocabulary that forecloses accountability frames.

These findings should be interpreted alongside the genuine complexity of the Salvadoran case. Bukele's narrative management has been effective in part because it responds to real grievances. The dramatic reduction in homicide rates (from 53 per 100,000 in 2018 to approximately 2 per 100,000 in 2024) represents a transformation in daily life for millions of Salvadorans. His framing of international critics as disconnected elites who prioritize criminal rights over citizen safety resonates because the prior system did, in fact, fail working-class

communities. This is precisely what makes spin dictatorship analytically distinct from pure propaganda: the narrative works because it contains enough truth to be persuasive. The question this paper raises is not whether Bukele has delivered security improvements, but whether those improvements are being used to foreclose legitimate scrutiny of how they were achieved.

Several limitations deserve emphasis. First, cosine similarity measures aggregate lexical overlap rather than semantic content. Two accounts might use identical words with different tone or implication. Embedding-based similarity measures would strengthen this analysis. Second, the before/after comparison for institutional convergence is not a clean natural experiment. Three changes occurred simultaneously on May 1, 2021: (a) new personnel loyal to Bukele replaced the leadership of @AsambleaSV and @FGR_SV; (b) a new Nuevas Ideas supermajority took over the legislature, shifting the institution's political agenda regardless of any coordination directive; and (c) the post-pandemic political environment and the impending Bitcoin rollout altered the broader communications landscape for all accounts. The design cannot fully disentangle personnel replacement from these concurrent shifts. One partial check is available: @Gobierno_SV and @PresidenciaSV were already under executive control before May 2021, yet they also converged toward @nayibbukele post-capture (+29% and +52%, respectively), suggesting the shift was not purely an artifact of leadership replacement in the directly captured institutions. Nevertheless, the finding should be interpreted as consistent with the capture hypothesis rather than a causal identification of it. Third, we do not model audience reception. Government messaging and independent media may reach different audiences, in which case the framing divergence we document may not translate into measurable attitudinal differences among citizens.

Future work should extend this analysis in three directions. First, a structural topic model (Roberts et al., 2014) applied to the full corpus would allow finer-grained mapping of topic prevalence across sources and time. Second, embedding-based institutional similarity measures would allow detection of semantic convergence that surface-level lexical overlap misses. Third,

linking the text data to public opinion survey data (e.g., LAPOP, Latinobarómetro) would allow assessment of whether narrative coordination corresponds to measurable attitude formation.

7. Conclusion

This paper has assembled the first large-scale, multi-source corpus of Salvadoran political text and applied computational methods to test whether El Salvador under Bukele exhibits the communicative hallmarks of spin dictatorship. The answer, across three distinct analytical approaches, is yes. Government messaging has become more internally coordinated since the May 2021 institutional capture; the administration systematically avoids accountability language while foregrounding security frames; and the captured institutions have adopted a vocabulary markedly closer to the president's personal account than they used before the takeover.

These findings have implications beyond the Salvadoran case. If communicative convergence is a general signature of institutional capture, text-as-data methods offer a new toolkit for democratic backsliding detection that complements existing indicators based on formal institutional structure and election integrity. The Bukele case also illustrates a core tension in contemporary authoritarian politics: the most durable narrative management is not pure fabrication but selective emphasis, foregrounding real improvements while systematically suppressing the vocabulary of cost and accountability. The underlying corpus, analysis code, and interactive visualizations are publicly available to support replication and extension at <https://github.com/elizabethavargas/elsalvador> and <https://el-salvador-news.netlify.app>.

References

- Amnesty International. (2023, April). One year into state of emergency, authorities are systematically committing human rights violations. <https://www.amnesty.org/en/latest/news/2023/04/el-salvador-state-emergency-systematic-human-rights-violations/>
- Bermeo, N. (2016). On democratic backsliding. *Journal of Democracy*, 27(1), 5–19.
- El Faro. (2023, March). Radiografía de un año de régimen de excepción. https://elfaro.net/es/202303/el_salvador/26786/radiografia-de-un-ano-de-regimen-de-excepcion
- Freedom House. (2024). Freedom in the world 2024: El Salvador. Freedom House. <https://freedomhouse.org/country/el-salvador/freedom-world/2024>
- Grimmer, J., and Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3), 267–297.
- Gunitsky, S. (2015). Corrupting the cyber-commons: Social media as a tool of autocratic stability. *Perspectives on Politics*, 13(1), 42–54.
- Guriev, S., and Treisman, D. (2019). Informational autocrats. *Journal of Economic Perspectives*, 33(4), 100–127.
- Guriev, S., and Treisman, D. (2022). *Spin dictators: The changing face of tyranny in the 21st century*. Princeton University Press.
- Human Rights Watch. (2022a). World report 2022: El Salvador. <https://www.hrw.org/world-report/2022/country-chapters/el-salvador>
- Human Rights Watch. (2022b). “We can arrest anyone we want”: Widespread human rights violations under El Salvador’s state of emergency. <https://www.hrw.org/report/2022/12/07/we-can-arrest-anyone-we-want/widespread-human-rights-violations-under-el>
- Inter-American Commission on Human Rights (IACHR). (2021). IACHR condemns the removal of the judges of the Constitutional Chamber of the Supreme Court of Justice in El Salvador (Press Release 110/21). https://www.oas.org/en/iachr/media_center/preleases/2021/110.asp
- International Monetary Fund. (2021). El Salvador: 2021 Article IV consultation. IMF Country Report No. 21/237.
- Martinez, J. (2023). salvadoran-news [Data set]. HuggingFace. <https://huggingface.co/datasets/justinian336/salvadoran-news>

- Levitsky, S., and Way, L. A. (2010). *Competitive authoritarianism: Hybrid regimes after the cold war*. Cambridge University Press.
- Monroe, B. L., Colaresi, M. P., and Quinn, K. M. (2008). Fightin' words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4), 372–403.
- Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., and Rand, D. G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58(4), 1064–1082.
- Schedler, A. (2006). The logic of electoral authoritarianism. In A. Schedler (Ed.), *Electoral authoritarianism: The dynamics of unfree competition* (pp. 1–23). Lynne Rienner.
- Statista. (2024). Homicide rate in El Salvador 2025. <https://www.statista.com/statistics/696152/homicide-rate-in-el-salvador/>
- V-Dem Institute. (2024). *Democracy report 2024: Democracy winning and losing at the ballot*. University of Gothenburg.
- Wolf, S. (2017). *Mano dura: The politics of gang control in El Salvador*. University of Texas Press.