Markets for Extremity: Digital Attention, Fringe Supply, and Radical-Right Agenda Setting

Arne Holverscheid[1]

[1]Northwestern University

Abstract

How do radical right parties decide which narratives to prioritize once they have entered the electoral mainstream, and why do their formulations so often track the most extreme currents of online discourse? I argue that digital embeddedness is not merely a communications strategy: in low-cost, high-velocity attention markets, fringe actors supply boundary-pushing frames that diffuse through forwarding networks and can be taken up by party-linked communicators. To test this bottom-up mechanism against a top-down elite-curation account, I assemble a new corpus of public German-language Telegram messages (2015–2025) from 16,790 channels and groups connected by 188,711 observed forwarding ties. I link weekly, directed diffusion networks to automated content measures produced by a locally run large language model that assigns binary labels for electoral talk and delegitimizing-election framing, alongside related indicators of extremist rhetoric. Exploiting the annulment of Romania's 2024 presidential first-round vote as a common narrative shock to election-integrity discourse, I implement difference-in-differences tests that compare AfD-affiliated chats with higher versus lower *pre-shock* exposure to fringe sources. The design pairs (i) a general measure of extremist rhetoric with (ii) a narrative-specific measure of *conditional framing* (delegitimizing frames among election-related messages), and evaluates whether pre-existing susceptibility to diffusion predicts sharper post-shock shifts.

Introduction

In recent decades, the far-right has become a major political force across Europe and in other democratic systems. Once relegated to the margins of political life, radical right parties (RRPs) have achieved significant electoral breakthroughs, forming part of governing coalitions in countries such as Italy, the Netherlands, and Austria, and emerging as major opposition forces in Germany and France. As they have entered the electoral mainstream, many of the newer RRPs have undergone a process of institutionalization (Janda 1970; Panebianco and Panebianco 1988). Often emerging from reactionary movements, ethno-nationalist networks, and fringe ideological milieus, these parties have consolidated into enduring political organizations with formal leadership structures, paid staff, and organizational routines that mirror those of established parties (De Lange and Art 2011; Heinisch and Mazzoleni 2016; Art 2018).

Mainstreaming and institutionalization, however, raise a central question about radical right politics once it moves from outsider agitation to sustained electoral relevance: how do RRPs decide which issues to champion, which narratives to prioritize, and how far to push their rhetoric when new events create opportunities for agenda-setting? Classic accounts emphasize party competition and relatedly the party family's capacity for strategic differentiation and issue entrepreneurship. While these mechanisms clearly matter for explaining the direction of RRP positioning, they leave an empirical puzzle that has become harder to ignore as RRPs have grown: even when the underlying direction of their appeals is clear, RRPs frequently adopt formulations that sit at the extreme end of that direction, and they often do so in a way that appears to follow the rhythms of fast-moving online discourse rather than only electoral calendars or legislative bargaining. Understanding where these extreme formulations come from, and why party-linked actors adopt them, is essential because RRPs can reshape the broader issue agenda and prompt downstream adaptation by other parties (Spoon and Klüver 2020; Abou-Chadi and Krause 2020).

This paper advances a theory of radical right ideological priorities that treats digital embeddedness as more than a communications strategy. RRPs are often described as centralized and leader-driven organizations, with message control routed through a narrow leadership circle. At the same time, their historical outsider status and the constraints of social norms and exclusion have encouraged many RRPs to cultivate alternative infrastructures for mobilization and attention, including highly permissive online spaces (Muis, Klein, and Dijkstra 2020). In such environments, political content circulates through networked dynamics of selection and amplification: intermediaries forward and reframe messages, fringe actors supply sharper or more conspiratorial formulations, and engagement rewards transgression. A central theoretical possibility developed in this paper is that these digital dynamics can feed back into party-linked discourse in a bottom-up manner. Rather than merely projecting a pre-formulated agenda

outward, party-affiliated actors embedded in a networked media ecosystem may draw on frames and cues that circulate in more extreme online milieus and, as they encounter and reuse them, incorporate these formulations into party-linked discourse. This mechanism is especially likely to matter for narratives that are both highly salient and normatively charged, because they create incentives for boundary-pushing claims that are attention-grabbing and easily repackaged.

The empirical puzzle is therefore not whether RRPs use the internet, which is well documented, but whether the internet can shape what RRPs say and prioritize once they have entered the mainstream. This motivates the paper's guiding question: does deep embeddedness in online far-right ecosystems pull RRPs toward more extreme issue positions than they would otherwise take? A competing perspective is explicitly top-down: party elites strategically curate their agenda and deploy extreme frames instrumentally, while online discourse reflects amplification rather than influence. The design is structured to distinguish between these mechanisms by combining observed information flow with variation in pre-existing exposure to extreme sources.

To evaluate these claims, I leverage a new corpus of public messages from German-language Telegram groups and channels between 2015 and 2025, assembled via Telegram's API. Telegram is analytically useful because it functions as a central infrastructural space for far-right communication while also providing observable traces of content diffusion through its forwarding function. I use forwarding metadata to construct a directed, time-varying network of information flow and link this network to automated measures of topic salience and extremist content in messages. The empirical focus is the Alternative für Deutschland (AfD) and its affiliated channels, embedded within a much larger constellation of non-party actors that spans populist influencers, movement entrepreneurs, and openly extremist communities.

The research design exploits a common narrative shock that plausibly raises the salience of a specific theme across the ecosystem, and then tests whether the post-shock evolution of AfD-linked discourse varies systematically with pre-shock exposure to extreme sources. In substantive terms, the expectation implied by the bottom-up mechanism is straightforward: after the shock, AfD-affiliated channels that were more exposed beforehand to fringe sources should show a larger increase in extremist rhetoric and, conditional on discussing the election, should be more likely to adopt delegitimizing or conspiratorial framings. The design does not require that the shock differentially "treat" units directly; instead, it leverages heterogeneity in susceptibility to diffusion, operationalized as pre-period exposure measured from observed forwarding ties. The empirical analysis combines descriptive panel patterns with difference-in-differences style tests using this exposure contrast.

This paper contributes to research on radical right parties and digital politics in three ways. First, it advances our understanding of the mechanisms through which the most extreme currents of far-right

discourse enter mainstream politics. Rather than treating radicalization as a purely strategic choice made within party competition, the argument highlights a pathway by which fringe formulations can be taken up by party-linked actors and incorporated into party-relevant communication. Second, it speaks to broader debates about how the internet and social media shape politics. Although the influence of online communication on political behavior is a well-established research agenda, both digital infrastructures and party politics have changed at extraordinary speed. Testing whether online embeddedness helps determine the ideological content of party discourse updates this agenda by focusing on position formation. Third, the paper contributes novel data and an empirical strategy for causal inference with networked social media text, pairing large-scale platform data with network structure and large-language-model (LLM)-based measurement in a way that supports theory-driven inference about diffusion and ideological uptake.

The remainder of the paper proceeds as follows. The next section situates the argument in the literatures on issue competition, party organization, and the radical right's development as an outsider party family. I then develop the theoretical account of how networked online environments can convert fringe supply into party-linked uptake and derive empirical expectations about extremism activation and framing. The subsequent sections introduce the Telegram data, explain the network construction and content-coding strategy, and lay out the identification approach that leverages pre-shock exposure heterogeneity. Please note: Data collection and annotation are still in progress, so I present the theoretical argument and research design here, with empirical results to follow.

Position Formation Inside the Radical Right

As RRPs begin to dominate public discourse, they drive the salience of issues they assert ownership over and shift other parties' positions as a result (Spoon and Klüver 2020; Abou-Chadi and Krause 2020; Völker and Saldivia Gonzatti 2024). This makes it a crucial empirical puzzle to understand how, which, and why they choose the issues and topics they do, beyond what their general ideological identity as a party family defined by nativism, authoritarianism, and populism would predict (Mudde 2007). Kitschelt and McGann (2005) offer a foundational account of RRP positioning, theorizing that issue adoption is a function of strategic differentiation in a multidimensional policy space in which more moderate parties converge to a median voter. Others have similarly highlighted the role of party competition as a driving factor for RRP positioning (Meguid 2005; De Lange 2007) and, relatedly, issue entrepreneurship (Hobolt and Tilley 2016; De Vries and Hobolt 2020). For instance, some RRPs have changed their stance on biological race to more ideational criteria of nativism in the successful attempt to appeal to a broader group of supporters (Koopmans and Statham 1999; Halikiopoulou, Mock, and Vasilopoulou 2013). Similarly, economic policy has long been particularly fluid among RRPs as they oscillate between

economic liberalism and welfare chauvinism (Enggist and Pinggera 2022), which exemplifies their ability to shift ideologically in response to electoral demand (De Lange 2007; Ivaldi 2015; Enggist and Pinggera 2022). But examining the determinants of party positions primarily through the lens of party competition necessarily limits our theoretical focus, leaving unexamined the internal organizational processes through which positions are debated, framed, and institutionalized.

The literature typically describes RRPs as highly centralized and "top-heavy", with a small leadership circle and, more often than not, a single, charismatic figure like Marine Le Pen. Decision-making is routed through vertical chains of command, with few institutionalized channels for internal debate or participatory input. This leader-centric model has even been linked to the organizational stability of RRPs, because the absence of strong leadership under such a model can prove destabilizing, as was the case with List Pim Fortuyn or the German Republikaner (Ellinas 2009; Heinisch and Mazzoleni 2016; Ellinas and Lamprianou 2017). At the same time, however, sustained persistence depends on more than a compelling programme or charismatic leadership: RRPs also need to institutionalize a well-structured party organization by recruiting, training, and socializing competent personnel who reliably support the party's goals (De Lange and Art 2011). Art (2012) brings into focus that the trajectory of RRPs is shaped by the composition of their activist base, especially in the party's formative years. While they eventually develop into hierarchical organizations (Michels 1915), this suggests a potentially more fluid top-bottom structure at least internally.

Unlike other challenger parties, the outsider status of RRPs has not only entailed the absence of established organizational resources or institutional legacies; it has also been shaped by a longstanding *cordon sanitaire* and social norm against the far-right (Mudde 2010). Historically excluded from mainstream institutions as a socially unacceptable party family, RRPs have faced formal barriers such as intentional self-censoring by the media (De Jonge 2021) as well as informal political barriers such as the exclusion from coalition government or the social ostracization of (potential) party members (Art 2012). The erosion of such norms in turn has been identified as a key condition of the far right's recent gains (Halikiopoulou and Vasilopoulou 2018; Wondreys and Mudde 2022; Valentim 2024).

As a response to both social norms and institutional constraints, many RRPs have turned to online spaces and social media, where low entry costs and the relative absence of social norms have allowed them to scale up their organizational structure and amplify their message with minimal accountability. Gerbaudo (2019) characterizes these formations as "digital parties", arguing that they adopt digital tools not only to bypass institutional deficits but also to project a narrative of participatory authenticity (Bennett, Segerberg, and Knüpfer 2018) and to communicate with the public and their supporters. RRPs have portrayed themselves as the sole representatives of a supposedly silenced "true" people, a rhetorical

strategy characteristic of such movements (Mudde 2007, 2019), and the appearance of direct engagement with "the people" online has become central to how these parties justify their representative claims, even if control over the party may remain centralized.

Position Taking and the Internet

Considering their unique political development as both: (1) historical outsiders with less institutionalized organizational structures; and (2) digital parties, existing research has overlooked the potential role online spaces may play in RRP position taking. While the use of the internet as a tool of the far-right is well documented (Wodak 2015; Muis, Klein, and Dijkstra 2020), including the role of echo-chambers (Barberá 2020; Boulianne, Koc-Michalska, and Bimber 2020; Ludovic Terren and Rosa Borge-Bravo 2021; Koiranen et al. 2022), we lack empirical evidence of the effect the internet conversely has on far-right parties.

I argue that online environments do not merely extend the communicative reach of RRPs but also reshape the internal calculus that determines how far along a given ideological dimension they travel. Parties represent coalitions whose core grievances anchor the direction of positioning (Aldrich 1995), yet this alone does not explain why RRPs often adopt positions at the extreme of that direction that outpace the median preferences of their own supporters. As noted above, strategic differentiation in party competition can sharpen positions to distinguish RRPs from centrist rivals and it may not substantially erode support if underlying demand for anti-immigration or anti-establishment appeals is inelastic. This explanation, however, implicitly relies on two demanding conditions. First, it presumes a degree of voter loyalty that immunizes RRPs against backlash from boundary-pushing rhetoric. Second, it presumes a level of message discipline and coordinated framing among party figures and affiliates that RRPs have historically struggled to maintain.

The internet offers an alternative explanatory mechanism that relaxes both requirements. In highly networked, low-cost spaces, attention is allocated through continuous contests for visibility, and content that is sharper, more transgressive, or more conspiratorial is disproportionately supplied by fringe actors, forwarded by intermediaries, and rewarded with engagement. These feedbacks lower the perceived reputational cost of extremity for party actors embedded in such ecosystems and increase the availability of ready-made frames when related topics rise in salience. I theorize that the consequence is not merely passive exposure; these dynamics can increase the likelihood that party figures adopt and integrate more extreme articulations into their own communications and programmatic outputs. In this sense, networked online environments can convert fringe supply into party-level uptake, shifting expressed positions toward the tail of the party's ideological direction even when the median supporter

would not privately endorse the most extreme formulation. To evaluate this mechanism empirically, I ask the following research question:

Research question: Given a newly emerging narrative, do actors within a radical right party who are more exposed to extreme actors in their media ecosystem adopt more extreme positions on that narrative?

The two perspectives we have outlined so far are not mutually exclusive. Electoral strategy and online embeddedness can operate in tandem: party figures may recognize that extreme framings circulate and gain traction in their digital surroundings, and they may allow this process to proceed with limited internal correction, anticipating benefits from boundary-pushing content without always issuing it under the most official imprimatur. In this interpretation, a finding that greater exposure to extreme actors predicts sharper positions on a newly salient narrative would not imply an absence of agency. It would indicate that networked environments provide a channel through which extreme formulations become available and low cost to adopt, while elite strategy may consist in tolerating, amplifying, or selectively internalizing those formulations.

By contrast, if such exposure does not cause position taking, the observation that parties nonetheless articulate positions more extreme than the median preferences of their supporters, and that these formulations usually originate online, would support a more deliberate mechanism: elites would be actively curating and deploying specific extreme frames for strategic ends rather than drifting toward them through ambient digital influence. The distinction concerns mechanism rather than direction; whether extreme output reflects permissive uptake from the digital milieu or purposeful orchestration by party elites, the internal calculus behind either route lies beyond the scope of the present study.

I assemble a large corpus of public German-language Telegram groups and channels via the platform's API, collecting full message histories with timestamps, text, and forwarding metadata. Messages are aggregated to a weekly panel at the chat level, and AfD-affiliated chats are identified ex ante. As an exogenous narrative shock around "election integrity," I leverage the annulment of Romania's 2024 presidential first-round vote, announced just before the planned 8 December runoff, which catalyzed a wave of delegitimizing-election discourse across European right-wing ecosystems (CeMAS – Center für Monitoring, Analyse & Strategie 2025). I then proceed in two steps. First, I use descriptive two–way fixed–effects panel regressions to summarize pre– and post–shock changes in two outcomes: (i) a measure of *extremist rhetoric* generally and (ii) *delegitimizing (conspiratorial) election framing* among election-related messages. Second, I implement a difference-in-differences design that compares higher versus lower pre-shock exposure to fringe sources, estimating whether post-shock changes in each outcome are larger among AfD-affiliated chats that were more embedded in fringe information flows prior to the

shock.

H1 (Extremism activation). After the narrative shock, AfD-affiliated chats with higher pre-shock exposure to fringe sources will, on average, exhibit a larger post-shock increase in *extremist rhetoric* than lower-exposure chats.

H2 (Frame uptake). Conditional on discussing the election, AfD-affiliated chats with higher pre-shock exposure will, on average, post a larger share *among election-related messages* that employ delegitimizing-election frames than lower-exposure chats.

Data and Research Design

I draw on a novel dataset collected from the direct messaging application Telegram, which has become a central space for far-right digital infrastructure in recent years (Holnburger 2023; Knüpfer and Hoffmann 2025). This is in large part due to the platform's refusal to enforce any noteworthy content moderation, meaning that it is effectively possible to post or share anything without fear of getting banned or the content being deleted. Unlike more closed platforms such as WhatsApp, Telegram also supports more public communication formats. Two features are especially relevant for this study: public groups, in which all members can post and comment, and public channels, where only the administrator can post content, but any user can subscribe and read posts in a broadcast-style feed. Importantly, Telegram's API is exceptionally open, as anyone can access and download all public messages and metadata without significant restrictions, aside from basic rate limits.

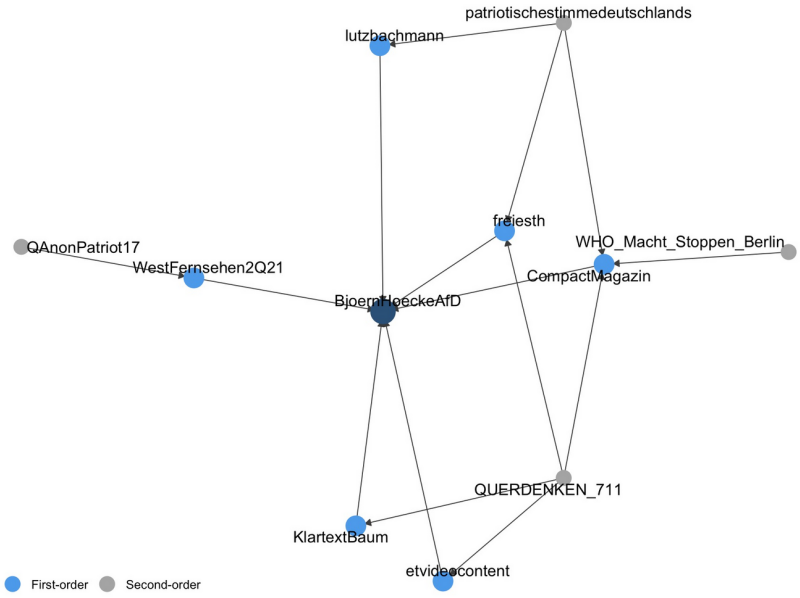

Figure 1: Partial snapshot of a two-step forwarding sub-network. Edges run from the original source toward the channel that forwarded the message, so arrow direction reflects downstream content flow. First-order nodes are channels that @BjoernHoeckeAfD reposted directly, while second-order nodes are the sources those channels, in turn, have forwarded from.

Beyond being remarkably accessible for data collection, Telegram's architecture can be leveraged to observe how its various communities are interlinked. A built-in *forwarding* function enables users, whether individual members in public groups or administrators of public channels, to repost content from one public chat to another (Jost et al. 2023). These forwarded messages constitute observable ties between different chats, allowing us to treat the Telegram ecosystem as a directed social network: public chats (channels or groups) function as nodes, and each forwarded message creates a directed edge from the source to the destination chat. Unlike other platforms where relationships must be inferred (e.g. based on retweets or follower overlap), Telegram's forward function provides a visible trace of information flow and channel connectivity. This network structure reflects both content diffusion and likely patterns of attention and affiliation, since forwarding implies that actors behind a chat not only consume content from other chats but likely follow and identify with them as well.

Data and Case Selection

This study focuses on the German case, where the radical right has both become more institutionalized politically with electoral success, as well as become deeply embedded in digital spaces and ecosystems. Germany's far-right milieu is especially active on Telegram, and, crucially for this study, this includes not

only formal party structures such as the Alternative für Deutschland (AfD), but also a broad constellation of actors spanning the ideological spectrum. These range from extreme-right and openly neo-Nazi groups, to populist influencers aligned with AfD messaging, to local groups that mobilize around COVID-19 restrictions or immigration, or at times post entirely apolitical content about lost pets or local flea markets. This heterogeneity makes the German Telegram ecosystem analytically valuable: it enables the study of a radical right party embedded in a wider, active, and ideologically diverse digital field. It also allows us to observe whether and how formal party actors, who are typically understood as hierarchical and controlled, interact with, and potentially absorb, content from a decentralized, more extreme online environment.

To construct the dataset, I use a snowball sampling approach based on Telegram's forward function. I begin with a curated set of seed accounts consisting of (1) AfD-affiliated accounts, including major national and subnational branches and party figures, and (2) prominent far-right non-party actors. The latter group was identified based on (a) research conducted by independent monitoring organizations such as CEMAS (Holnburger 2023; Dittrich, Düker, and Müller 2023), and (b) manual confirmatory inspection of Telegram activity and follower engagement. From this initial set, I collect all public messages posted since January 2015. I then extract any channels or groups that these seed accounts forwarded messages from. This process is repeated two additional times (i.e., two "rounds"), expanding the sample to include chats connected through forwarded content. The result is a directed communication network consisting of 16,790 nodes and 188,711 edges, representing a large and internally connected subnetwork of German-language far-right Telegram activity.

Each node in the network represents a public Telegram channel or group, and each edge encodes the forwarding of messages from one chat to another. The underlying content of the network consists of the full message histories of these chats, with each message associated with a timestamp, text body, and the chat in which it was posted. To facilitate longitudinal analysis, each message is assigned to a calendar week based on its timestamp. This allows all subsequent modeling to be conducted at the weekly level. Weekly binning offers a balance between granularity and stability: daily-level data can be too sparse, especially for less active chats, while monthly aggregation risks obscuring short-lived or fast-moving topic dynamics. I create a binary variable in the panel identifying whether each chat is affiliated with the AfD. All other chats are coded as unaffiliated. A subset of AfD-affiliated chats exhibit zero observed incoming forwards in the pre-shock window; I report their prevalence and treat them separately in exposure-based analyses where $X_{i,\text{pre}}^{\text{fringe}}$ would otherwise be undefined or uninformative.

The analysis relies on an automated content–coding pipeline that assigns binary labels to every Telegram message along six conceptually distinct dimensions: conspiracism, electoral talk, extra-legal coercion or

harassment, out-group demonization, supernatural claims, and endorsements of violence. The classification is achieved using a locally run large language model (Gemma 3 12B) accessed via an API. For every dimension, the model is prompted with a short definition and two to three illustrative examples and is instructed to return a single token in {`Yes`, `No`}. Inference uses temperature = 0 and otherwise default decoding, yielding deterministic outputs given a fixed prompt. Let $I_m^{(c)} \in \{0, 1\}$ denote the normalized label for message m and category c, where $I_m^{(c)} = 1$ iff the model returns `Yes`.

Outputs are normalized and stored with message identifiers and timestamps. From these base labels, I construct two outcome measures at the chat-week level that separate *extremism activation* from *conditional framing*. First, *extremism activation* captures the weekly prevalence of extremist rhetoric, $S_{i,t}^{\text{ext}}$. Second, *delegitimizing-election framing* is defined conditional on election talk: among election-related messages in a given chat–week, it is the share that employ delegitimizing frames (operationalized as conspiracism within election talk), $C_{i,t}^{\text{delig}|\text{elec}}$. Weekly quantities are computed by aggregating labeled messages within each chat and week; conditional framing is defined only for chat–weeks with at least one election-related message. These unit-time series are then used in descriptive panels and the difference-in-differences design.

I use a prompted, decoder-only large language model to code message-level categories because it offers a practical way to combine conceptual nuance with scale. Recent work argues that generative LLMs can relax the classic trade-off between small-n human coding (high validity, low coverage) and fully automated methods that scale only by sacrificing interpretability, since LLMs can apply detailed codebooks in zero- or few-shot form without retraining (Linegar, Kocielnik, and Alvarez 2023; Bail 2024). Empirically, evaluations in political science show that prompted LLMs often match or exceed both human coders and supervised classifiers on classification tasks, with high intercoder-style agreement and reliable performance in complex multi-class settings (Gilardi, Alizadeh, and Kubli 2023; Mellon et al. 2024). Related evidence suggests that LLMs also handle context-sensitive inference tasks that are difficult to operationalize with keyword rules or fixed embeddings, including cross-national and multilingual settings (Törnberg 2024). In this project, these properties make LLM coding a transparent and reproducible way to score the full corpus at message scale under a fixed prompt and deterministic decoding.

To characterize the pre-shock information environment, I summarize non-electoral labels into a message-level extremism score and then average these scores within channels over the pre-shock period. Channels above a high quantile of this distribution are designated as *fringe*. Pre-shock exposure is then defined using observed forwarding ties as the share of a chat's forwards that originate from channels designated fringe. This exposure measure serves as the treatment contrast in the empirical tests that follow.

I construct a sequence of time-varying adjacency matrices W_t, where each matrix records observed

forwarding behavior between chats in week t. A directed edge from chat j to chat i indicates that i forwarded at least one message originating from j during that period. Edges are interpreted as realized diffusion events: forwarding is taken as direct evidence that content from j was transmitted into i's feed. I measure two weekly outcomes at the chat–week level, separating *extremism activation* from *conditional framing*. First, I construct an extremism prevalence outcome that captures how much extremist rhetoric appears in a chat in a given week. Let $\mathcal{C}$ denote the five extremist dimensions (conspiracism, out-group demonization, extra-legal coercion or harassment, supernatural claims, and endorsements of violence). For each message m, define an extremism score

$$E_m = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} I_m^{(c)}.$$

Or, simply put, this number represents how extreme each message is based on how many of the scored categories it fits. I then aggregate to the chat–week level:

$$S_{i,t}^{\text{ext}} \;=\; \frac{1}{|\mathcal{M}_{i,t}|} \sum_{m \in \mathcal{M}_{i,t}} E_m.$$

This yields a weekly extremism score for each Telegram channel. Second, I measure *delegitimizing-election framing conditional on election talk*. Let $I_m^{\text{elec}} \in \{0,1\}$ flag whether message m is about electoral politics and let $I_m^{\text{consp}} \in \{0,1\}$ be the conspiracism label. Define a delegitimizing-election indicator as the intersection

$$I_m^{\text{delig}} \;=\; I_m^{\text{consp}} \cdot I_m^{\text{elec}}.$$

The conditional framing outcome is then

$$C_{i,t}^{\text{delig|elec}} \;=\; \frac{\sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{delig}}}{\sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{elec}}},$$

which is defined only for chat–weeks with $\sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{elec}} > 0$. Intuitively, $C_{i,t}^{\text{delig|elec}}$ isolates *how elections are framed* from *how frequently elections are discussed*.

For completeness, I also report the unconditional prevalence of delegitimizing-election messages,

$$S_{i,t}^{\text{delig}} \;=\; \frac{1}{|\mathcal{M}_{i,t}|} \sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{delig}},$$

as a supplementary outcome that mechanically combines election talk and framing.

To classify upstream sources as *fringe* for the exposure measure, I aggregate E_m to the channel level

over the pre-shock period $\mathcal{T}_{\text{pre}}$:

$$\bar{E}_{j,\text{pre}} = \frac{1}{|\mathcal{T}_{\text{pre}}|} \sum_{s \in \mathcal{T}_{\text{pre}}} \left(\frac{1}{|\mathcal{M}_{j,s}|} \sum_{m \in \mathcal{M}_{j,s}} E_m \right).$$

Channels with $\bar{E}_{j,\text{pre}}$ at or above the empirical 0.75 quantile of the $\{\bar{E}_{j,\text{pre}}\}_j$ distribution are labeled *fringe*; I denote this set by $\mathcal{F}$. As a robustness check, I also use the 0.90 quantile cutoff. In other words, pre-shock exposure to the fringe for an AfD-affiliated channel i is defined as the share of its *incoming forwards* that originate from fringe sources (channels whose average extremism message-score is at a high percentile).

Let $W_{ji,s}$ be the count of messages in week $s \in \mathcal{T}_{\text{pre}}$ that i forwarded from j. The exposure is

$$X_{i,\text{pre}}^{\text{fringe}} = \frac{\sum_{s \in \mathcal{T}_{\text{pre}}} \sum_{j \in \mathcal{F}} W_{ji,s}}{\sum_{s \in \mathcal{T}_{\text{pre}}} \sum_{j} W_{ji,s}},$$

the fraction of all pre-shock forwards into i that came from channels classified as fringe.

Analysis

Preliminary panel OLS. As a descriptive first pass, I estimate a two–way fixed–effects panel regression that captures the average change in extremist rhetoric at the shock:

$$S_{i,t}^{\text{ext}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \delta\, \mathbf{1}\{t \geq t_0\} \;+\; \varepsilon_{i,t}, \tag{1}$$

where $S_{i,t}^{\text{ext}}$ is the extremism prevalence in chat i during week t, α_i and γ_t are chat and week fixed effects, and t_0 indexes the week of the annulment shock. The coefficient δ summarizes the average post–shock shift among AfD–affiliated chats. Standard errors are clustered at the chat level.

Difference–in–differences with continuous exposure. To test H1, I leverage pre–shock exposure heterogeneity measured by $X_{i,\text{pre}}^{\text{fringe}}$. The continuous DiD specification interacts this time–invariant measure with a post–shock indicator:

$$S_{i,t}^{\text{ext}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \beta \left(X_{i,\text{pre}}^{\text{fringe}} \times \mathbf{1}\{t \geq t_0\} \right) \;+\; \varepsilon_{i,t}. \tag{2}$$

Here, β captures whether chats with greater pre–shock exposure to fringe sources exhibit a larger post–shock increase in extremist rhetoric. Identification follows the standard parallel–trends assumption: absent the shock, the evolution of $S_{i,t}^{\text{elec}}$ would have been similar across exposure levels. I report estimates using the raw $X_{i,\text{pre}}^{\text{fringe}}$ and, for ease of interpretation, a standardized version.

Difference–in–differences with binary exposure. For a more interpretable contrast, I also estimate a binary DiD that compares high–exposure units to others. Let $T_i^{\text{hi}} = \mathbf{1}\{X_{i,\text{pre}}^{\text{fringe}} \geq \tau\}$, with τ set to the empirical 0.75 quantile (and 0.90 in robustness checks). The model is:

$$S_{i,t}^{\text{ext}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \beta_D \left(T_i^{\text{hi}} \times \mathbf{1}\{t \geq t_0\}\right) \;+\; \varepsilon_{i,t}. \tag{3}$$

The coefficient β_D is the average treatment effect on the treated under parallel trends between high– and lower–exposure chats. I report models with and without restricting the comparison group to a symmetric band around τ to avoid leverage from very low–exposure units.

H2: Conditional delegitimizing framing. To test H2, I estimate analogous specifications using the conditional outcome $C_{i,t}^{\text{delig|elec}}$, restricting the estimating sample to chat–weeks with $\sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{elec}} > 0$. A descriptive two–way fixed–effects specification is:

$$C_{i,t}^{\text{delig|elec}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \delta^{(2)}\, \mathbf{1}\{t \geq t_0\} \;+\; \varepsilon_{i,t}, \tag{4}$$

and the continuous exposure DiD is:

$$C_{i,t}^{\text{delig|elec}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \beta^{(2)} \left(X_{i,\text{pre}}^{\text{fringe}} \times \mathbf{1}\{t \geq t_0\}\right) \;+\; \varepsilon_{i,t}. \tag{5}$$

For the binary exposure contrast T_i^{hi}, I estimate:

$$C_{i,t}^{\text{delig|elec}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \beta_D^{(2)} \left(T_i^{\text{hi}} \times \mathbf{1}\{t \geq t_0\}\right) \;+\; \varepsilon_{i,t}. \tag{6}$$

These specifications isolate changes in *framing* from changes in the overall volume of election talk and therefore map directly onto the "uptake" interpretation in H2.

Causal identification and interference. The annulment is a common narrative shock and does not differentially treat units. The design therefore does not interpret $\mathbf{1}\{t \geq t_0\}$ as treatment assignment. Instead, it asks whether AfD-affiliated chats that were *more embedded pre-shock* in fringe information flows exhibit larger post-shock changes in extremism activation and framing. I operationalize this pre-existing embeddedness as $X_{i,\text{pre}}^{\text{fringe}}$, a time-invariant summary of incoming forwarding ties to channels classified as fringe in the pre-period, and estimate a reduced-form (intention-to-treat) effect: whether the post-annulment shift in outcomes scales with $X_{i,\text{pre}}^{\text{fringe}}$. In this formulation, interference is not assumed away; it is part of the mechanism that generates the narrative and allows it to reach some chats more readily than others, and $X_{i,\text{pre}}^{\text{fringe}}$ serves as a low-dimensional proxy for susceptibility to such diffusion.

Formally, I replace strict SUTVA with a stable *exposure-mapping* assumption: conditional on chat fixed effects and week fixed effects, a unit's post-shock potential outcomes may depend on the broader network only through its own pre-shock exposure score $X_{i,\text{pre}}^{\text{fringe}}$ and the shock indicator $\mathbf{1}\{t \geq t_0\}$. Because $X_{i,\text{pre}}^{\text{fringe}}$ is computed exclusively from pre-shock forwarding behavior and pre-shock content, it is fixed prior to the shock and avoids conditioning on post-treatment diffusion. Week fixed effects absorb shocks common to all AfD chats; identification then comes from differential post-shock changes by pre-shock exposure. This does not rule out *type confounding*: high-exposure AfD chats may differ systematically (e.g., baseline activity, connectivity, or ideological profile) in ways that also generate differential post-shock responses.

To address this type confounding concern, I (i) report event-study diagnostics that test for differential pre-trends by exposure; (ii) estimate specifications that allow pre-period chat characteristics (e.g., baseline activity, degree/forwarding volume, and baseline outcome levels) to have their own post-shock shifts via interactions with $\mathbf{1}\{t \geq t_0\}$; and (iii) show robustness to alternative exposure parameterizations (standardized continuous; binary high-exposure with $\tau \in \{0.75, 0.90\}$), trimming/down-weighting high-degree hubs, and restricting comparisons to a symmetric band around τ. Together, these checks assess whether post-shock divergence by $X_{i,\text{pre}}^{\text{fringe}}$ is consistent with diffusion-mediated uptake rather than exposure-correlated differential shocks.

Dynamic diagnostics. To probe the plausibility of parallel trends and to visualize dynamics, I estimate an event–study that saturates relative–time indicators interacted with exposure:

$$S_{i,t}^{\text{ext}} \;=\; \alpha_i \;+\; \gamma_t \;+\; \sum_{\ell \neq -1} \beta_\ell \left(D_\ell(t) \times T_i^{\text{hi}} \right) \;+\; \varepsilon_{i,t}, \tag{7}$$

where $D_\ell(t) = \mathbf{1}\{t - t_0 = \ell\}$ are relative–week dummies with $\ell = -1$ omitted. Pre–shock coefficients β_ℓ for $\ell < 0$ provide a test for differential pre–trends; post–shock coefficients trace the evolution of treatment effects.

Estimation details and reporting. All models are estimated on the AfD–affiliated sample at the chat–week level. I cluster standard errors at the chat level and verify robustness to two–way clustering by chat and week. Because outcomes are shares, estimates are reported on the percentage–point scale. Because outcomes are estimated with varying precision, I additionally report weighted least squares (WLS). For extremism outcomes ($S_{i,t}^{\text{ext}}$ and the supplementary $S_{i,t}^{\text{delig}}$), weights are $|\mathcal{M}_{i,t}|$. For conditional framing ($C_{i,t}^{\text{delig|elec}}$), weights are the conditional denominator $\sum_{m \in \mathcal{M}_{i,t}} I_m^{\text{elec}}$. I also verify that winsorizing outcomes at the upper tail leaves inferences unchanged. Specifications include the baseline (1), the continuous DiD (2), the binary DiD (3), and the event–study (7). Figures display coefficient paths with

95 percent confidence intervals and binned residual plots for model fit diagnostics.

Results

As of this writing, I do not present empirical results, as data collection is still in progress. This is primarily due to the scale and complexity of the Telegram data. The dataset is being assembled through a computationally intensive workflow that runs on a high-performance computing cluster, with careful storage and backup to an external drive. Both the Telegram API scraping and the LLM-based scoring require long runtimes due to rate limits and the processing demands of the models.

Discussion

References

Abou-Chadi, Tarik, and Werner Krause. 2020. "The Causal Effect of Radical Right Success on Mainstream Parties' Policy Positions: A Regression Discontinuity Approach." *British Journal of Political Science* 50, no. 3 (July): 829–847. ISSN: 0007-1234, 1469-2112, accessed June 28, 2023. https://doi.org/10.1017/S0007123418000029.

Aldrich, John Herbert. 1995. *Why Parties? The Origin and Transformation of Political Parties in America.* American Politics and Political Economy Series. Chicago: The University of Chicago Press. ISBN: 978-0-226-01271-1 978-0-226-01272-8.

Art, David. 2012. *Inside the Radical Right: The Development of Anti-Immigrant Parties in Western Europe.* Repr. New York, NY: Cambridge University Press. ISBN: 978-0-521-72032-8 978-0-521-89624-5.

———. 2018. *Party Organization and the Radical Right.* Edited by Jens Rydgren. Vol. 1. Oxford University Press, February. Accessed May 27, 2025. https://doi.org/10.1093/oxfordhb/9780190274559.013.12.

Bail, Christopher A. 2024. "Can Generative AI Improve Social Science?" *Proceedings of the National Academy of Sciences* 121, no. 21 (May): e2314021121. ISSN: 0027-8424, 1091-6490, accessed August 15, 2025. https://doi.org/10.1073/pnas.2314021121.

Barberá, Pablo. 2020. "Social Media, Echo Chambers, and Political Polarization." In *Social Media and Democracy,* 1st ed., edited by Nathaniel Persily and Joshua A. Tucker, 34–55. Cambridge University Press, August. ISBN: 978-1-108-89096-0 978-1-108-83555-8 978-1-108-81289-4, accessed June 5, 2025. https://doi.org/10.1017/9781108890960.004.

Bennett, W. Lance, Alexandra Segerberg, and Curd B. Knüpfer. 2018. "The Democratic Interface: Technology, Political Organization, and Diverging Patterns of Electoral Representation." *Information, Communication & Society* 21, no. 11 (November): 1655–1680. ISSN: 1369-118X, 1468-4462, accessed May 27, 2025. https://doi.org/10.1080/1369118X.2017.1348533.

Boulianne, Shelley, Karolina Koc-Michalska, and Bruce Bimber. 2020. "Right-Wing Populism, Social Media and Echo Chambers in Western Democracies." *New Media & Society* 22, no. 4 (April): 683–699. ISSN: 1461-4448, 1461-7315, accessed June 5, 2025. https://doi.org/10.1177/1461444819893983.

CeMAS – Center für Monitoring, Analyse & Strategie. 2025. *Bundestagswahl 2025 Monitoring.* https://btw2025.cemas.io/. Accessed January 7, 2026.

De Jonge, Léonie. 2021. "The Curious Case of Belgium: Why Is There No Right-Wing Populism in Wallonia?" *Government and Opposition* 56, no. 4 (October): 598–614. ISSN: 0017-257X, 1477-7053, accessed June 4, 2025. https://doi.org/10.1017/gov.2020.8.

De Lange, Sarah L. 2007. "A New Winning Formula?: The Programmatic Appeal of the Radical Right." *Party Politics* 13, no. 4 (July): 411–435. ISSN: 1354-0688, 1460-3683, accessed June 4, 2025. https://doi.org/10.1177/1354068807075943.

De Lange, Sarah L., and David Art. 2011. "Fortuyn versus Wilders: An Agency-Based Approach to Radical Right Party Building." *West European Politics* 34, no. 6 (November): 1229–1249. ISSN: 0140-2382, 1743-9655, accessed January 6, 2026. https://doi.org/10.1080/01402382.2011.616662.

De Vries, Catherine E., and Sara Binzer Hobolt. 2020. *Political Entrepreneurs: The Rise of Challenger Parties in Europe.* Princeton: Princeton University Press. ISBN: 978-0-691-20007-1.

Dittrich, Miro, Joe Düker, and Martin Müller. 2023. *Where's the Money at? Right-wing Extremist Fundraising over Telegram.* Technical report. Center for Monitoring, Analysis, and Strategy (CeMAS).

Ellinas, Antonis A. 2009. "Chaotic but Popular? Extreme-Right Organisation and Performance in the Age of Media Communication." *Journal of Contemporary European Studies* 17, no. 2 (August): 209–221. ISSN: 1478-2804, 1478-2790, accessed May 27, 2025. https://doi.org/10.1080/14782800903108650.

Ellinas, Antonis A, and Iasonas Lamprianou. 2017. "How Far Right Local Party Organizations Develop: The Organizational Buildup of the Greek Golden Dawn." *Party Politics* 23, no. 6 (November): 804–820. ISSN: 1354-0688, 1460-3683, accessed May 27, 2025. https://doi.org/10.1177/1354068816641337.

Enggist, Matthias, and Michael Pinggera. 2022. "Radical Right Parties and Their Welfare State Stances – Not so Blurry after All?" *West European Politics* 45, no. 1 (January): 102–128. ISSN: 0140-2382, 1743-9655, accessed June 4, 2025. https://doi.org/10.1080/01402382.2021.1902115.

Gerbaudo, Paolo. 2019. *The Digital Party: Political Organisation and Online Democracy.* Digital Barricades. London: Pluto press. ISBN: 978-0-7453-3579-7.

Gilardi, Fabrizio, Meysam Alizadeh, and Maël Kubli. 2023. "ChatGPT Outperforms Crowd Workers for Text-Annotation Tasks." *Proceedings of the National Academy of Sciences* 120, no. 30 (July): e2305016120. ISSN: 0027-8424, 1091-6490, accessed August 14, 2025. https://doi.org/10.1073/pnas.2305016120.

Halikiopoulou, Daphne, Steven Mock, and Sofia Vasilopoulou. 2013. "The Civic Zeitgeist: Nationalism and Liberal Values in the E Uropean Radical Right." *Nations and Nationalism* 19, no. 1 (January): 107–127. ISSN: 1354-5078, 1469-8129, accessed June 4, 2025. https://doi.org/10.1111/j.1469-8129.2012.00550.x.

Halikiopoulou, Daphne, and Sofia Vasilopoulou. 2018. "Breaching the Social Contract: Crises of Democratic Representation and Patterns of Extreme Right Party Support." *Government and Opposition* 53, no. 1 (January): 26–50. ISSN: 0017-257X, 1477-7053, accessed June 4, 2025. https://doi.org/10.1017/gov.2015.43.

Heinisch, Reinhard, and Oscar Mazzoleni. 2016. *Understanding Populist Party Organisation.* Edited by Reinhard Heinisch and Oscar Mazzoleni. London: Palgrave Macmillan UK. ISBN: 978-1-137-58196-9 978-1-137-58197-6, accessed May 27, 2025. https://doi.org/10.1057/978-1-137-58197-6.

Hobolt, Sara B., and James Tilley. 2016. "Fleeing the Centre: The Rise of Challenger Parties in the Aftermath of the Euro Crisis." *West European Politics* 39, no. 5 (September): 971–991. ISSN: 0140-2382, 1743-9655, accessed January 6, 2026. https://doi.org/10.1080/01402382.2016.1181871.

Holnburger, Josef. 2023. *Chronology of a Radicalization.* Technical report. Center for Monitoring, Analysis, and Strategy (CeMAS).

Ivaldi, Gilles. 2015. "Towards the Median Economic Crisis Voter? The New Leftist Economic Agenda of the Front National in France." *French Politics* 13, no. 4 (December): 346–369. ISSN: 1476-3419, 1476-3427, accessed June 4, 2025. https://doi.org/10.1057/fp.2015.17.

Janda, Kenneth. 1970. *A Conceptual Framework for the Comparative Analysis of Political Parties.* Beverly Hills: Sage.

Jost, Pablo, Annett Heft, Kilian Buehling, Maximilian Zehring, Heidi Schulze, Hendrik Bitzmann, and Emese Domahidi. 2023. "Mapping a dark space: Challenges in sampling and classifying non- institutionalized actors on telegram." *Medien & Kommunikationswissenschaft* (Baden-Baden) 71 (3-4): 212–229. ISSN: 1615-634X.

Kitschelt, Herbert, and Anthony J. McGann. 2005. *The Radical Right in Western Europe: A Comparative Analysis.* 1. paperback ed., [Nachdr.] Ann Arbor: Univ. of Michigan Press. ISBN: 978-0-472-08441-8 978-0-472-10663-9.

Knüpfer, Curd, and Matthias Hoffmann. 2025. "Countering the "Climate Cult" – Framing Cascades in Far-Right Digital Networks." *Political Communication* 42, no. 1 (January): 85–107. ISSN: 1058-4609, 1091-7675, accessed June 4, 2025. https://doi.org/10.1080/10584609.2024.2332762.

Koiranen, Ilkka, Aki Koivula, Sanna Malinen, and Teo Keipi. 2022. "Undercurrents of Echo Chambers and Flame Wars: Party Political Correlates of Social Media Behavior." *Journal of Information Technology & Politics* 19, no. 2 (April): 197–213. ISSN: 1933-1681, 1933-169X, accessed June 5, 2025. https://doi.org/10.1080/19331681.2021.1950095.

Koopmans, Ruud, and Paul Statham. 1999. "Challenging the Liberal Nation-State? Postnationalism, Multiculturalism, and the Collective Claims Making of Migrants and Ethnic Minorities in Britain and Germany." *American Journal of Sociology* 105, no. 3 (November): 652–696. ISSN: 0002-9602, 1537-5390, accessed December 29, 2025. https://doi.org/10.1086/210357.

Linegar, Mitchell, Rafal Kocielnik, and R. Michael Alvarez. 2023. "Large Language Models and Political Science." *Frontiers in Political Science* 5 (October): 1257092. ISSN: 2673-3145, accessed August 13, 2025. https://doi.org/10.3389/fpos.2023.1257092.

Ludovic Terren, Ludovic Terren, and Rosa Borge-Bravo Rosa Borge-Bravo. 2021. "Echo Chambers on Social Media: A Systematic Review of the Literature." *Review of Communication Research* 9 (March). ISSN: 2255-4165, accessed June 5, 2025. https://doi.org/10.12840/ISSN.2255-4165.028.

Meguid, Bonnie M. 2005. "Competition Between Unequals: The Role of Mainstream Party Strategy in Niche Party Success." *American Political Science Review* 99, no. 3 (August): 347–359. ISSN: 0003-0554, 1537-5943, accessed November 9, 2023. https://doi.org/10.1017/S0003055405051701.

Mellon, Jonathan, Jack Bailey, Ralph Scott, James Breckwoldt, Marta Miori, and Phillip Schmedeman. 2024. "Do AIs Know What the Most Important Issue Is? Using Language Models to Code Open-Text Social Survey Responses at Scale." *Research & Politics* 11, no. 1 (January): 20531680241231468. ISSN: 2053-1680, 2053-1680, accessed August 15, 2025. https://doi.org/10.1177/20531680241231468.

Michels, Robert. 1915. *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy.*

Mudde, Cas. 2007. *Populist Radical Right Parties in Europe.* 1st ed. Cambridge University Press, August. ISBN: 978-0-521-61632-4 978-0-521-85081-0 978-0-511-49203-7, accessed March 20, 2023. https://doi.org/10.1017/CBO9780511492037.

———. 2019. *The Far Right Today.* Cambridge, UK ; Medford, MA: Polity. ISBN: 978-1-5095-3684-9 978-1-5095-3683-2.

———. 2010. "The Populist Radical Right: A Pathological Normalcy." *West European Politics* 33, no. 6 (November): 1167–1186. ISSN: 0140-2382, 1743-9655, accessed March 22, 2023. https://doi.org/10.1080/01402382.2010.508901.

Muis, Jasper, Ofra Klein, and Guido Dijkstra. 2020. "Challenges and Opportunities of Social Media Research." In *Researching the Far Right,* 1st ed., edited by Stephen D. Ashe, Joel Busher, Graham Macklin, and Aaron Winter, 147–163. Abingdon, Oxon; New York, NY: Routledge, 2020. | Series: Fascism and the far right: Routledge, November. ISBN: 978-1-315-30467-0, accessed January 29, 2024. https://doi.org/10.4324/9781315304670-9.

Panebianco, Angelo, and Angelo Panebianco. 1988. *Political Parties: Organization and Power.* Cambridge Studies in Modern Political Economies. Cambridge: Cambridge University Press. ISBN: 978-0-521-31401-5 978-0-521-30627-0.

Spoon, Jae-Jae, and Heike Klüver. 2020. "Responding to Far Right Challengers: Does Accommodation Pay Off?" *Journal of European Public Policy* 27, no. 2 (February): 273–291. ISSN: 1350-1763, 1466-4429, accessed May 27, 2025. https://doi.org/10.1080/13501763.2019.1701530.

Törnberg, Petter. 2024. "Large Language Models Outperform Expert Coders and Supervised Classifiers at Annotating Political Social Media Messages." *Social Science Computer Review* (September): 08944393241286471. ISSN: 0894-4393, 1552-8286, accessed June 4, 2025. https://doi.org/10.1177/08944393241286471.

Valentim, Vicente. 2024. *The Normalization of the Radical Right: A Norms Theory of Political Supply and Demand.* Oxford Studies in Democratization. Oxford: Oxford University Press. ISBN: 978-0-19-892671-9 978-0-19-892674-0. https://doi.org/10.1093/9780198926740.001.0001.

Völker, Teresa, and Daniel Saldivia Gonzatti. 2024. "Discourse Networks of the Far Right: How Far-Right Actors Become Mainstream in Public Debates." *Political Communication* 41, no. 3 (May): 353–372. ISSN: 1058-4609, 1091-7675, accessed June 11, 2025. https://doi.org/10.1080/10584609.2024.2308601.

Wodak, Ruth. 2015. *The Politics of Fear: What Right-Wing Populist Discourses Mean.* Los Angeles London. ISBN: 978-1-4462-4700-6 978-1-4739-1417-9.

Wondreys, Jakub, and Cas Mudde. 2022. "Victims of the Pandemic? European Far-Right Parties and COVID-19." *Nationalities Papers* 50, no. 1 (January): 86–103. ISSN: 0090-5992, 1465-3923, accessed June 4, 2025. https://doi.org/10.1017/nps.2020.93.