Domesticating Neoreaction: The French Far Right and Illiberal Discourses on YouTube, 2016–2024

Tristan Boursier* Arnaud Miranda† Antoine Lemor‡

2026

Abstract

Is the neoreactionary (NRx) current diffusing into the French far right, and are its themes becoming more prevalent than the classical far-right repertoire? Drawing on the YouPol corpus of French political YouTube, we measure five NRx themes against six classical far-right ones across twenty French far-right channels, 2016 to 2024, combining a distillation classifier (the IDS-NRx) with close reading. The two repertoires diverge rather than converge. NRx content crests in 2019 to 2021 and then recedes, declining within channels over time, while the classical repertoire intensifies around immigration. Most of what reads as neoreactionary is the French far right's own concerns in a borrowed idiom, since a nation-state politics absorbs and neutralises a doctrine built on dissolving it. No NRx theme gains a platform advantage, and its one distinctive trace is in reception, where NRx themes depoliticise comment threads, the fingerprint of an anti-political style.

Introduction

Reactionary and illiberal ideas have long crossed national borders through institutional carriers, intellectual networks, political parties, think tanks and the legacy press (Bures, 2023; Schilk, 2025; Norocel, 2023). Over the past decade a different carrier has moved to the centre of that circulation, the political influencer. Unlike the organic intellectuals of Gramscian theory or the militants documented by scholars of far-right mobilisation, these figures assemble mass audiences, often in the hundreds of thousands, by narrating the day's events and supplying ideological framings under the guise of independent commentary. They seldom present themselves as ideological entrepreneurs, and that apparent neutrality is part of what makes them effective. A creator who reads the news for half a million subscribers can move a vocabulary and a worldview further, and with less resistance, than a manifesto. The content these influencers actually produce, and not only the networks they inhabit, has accordingly become an object of study in its own right.

One body of ideas circulating through these channels is especially puzzling, because it appears to cut against the very far right that carries it. Neoreaction (NRx), reconstructed by Miranda (2026) and tied to the conservative turn of parts of the Silicon Valley tech milieu, is an illiberal grammar

*Université du Québec à Montréal, CEVIPOF, Sciences Po, tristan.boursier@sciencespo.fr
†CEVIPOF, Sciences Po, arnaud.miranda@sciencespo.fr
‡Université de Sherbrooke, antoine.lemor@umontreal.ca

built on natural hierarchy, anthropological pessimism, an explicit rejection of democracy in favour of corporate or monarchical government, a libertarian principle of exit, and a techno-capitalist faith that casts democracy and progressivism as obstacles to acceleration. Illiberal without being classically far-right, it sits at odds with the tradition it circulates within on the very axis that defines that tradition, since where the established far right valorises the people, the nation-state and cultural homogeneity, neoreaction is anti-populist, post-national, and oriented toward dissolving democratic government into an entrepreneurial logic of management (Laruelle, 2024; Carter, 2018). Its vocabulary has nonetheless surfaced in far-right online spaces well beyond the Anglophone world, raising the possibility that an established national far right might take it up and be reshaped by it (Hemmila and Perliger, 2024; Behr and Gianoncelli, 2025).

France offers a good setting in which to examine that possibility. Its far right rests on a dense and old reactionary tradition, from the Nouvelle Droite to the Rassemblement National, with established frameworks, diffusion networks and audiences within which an imported current might find receptive structures (Schir and Laruelle, 2025; Gianoncelli, 2021). At the same time French political culture stands at a particular distance from the American intellectual sphere: close enough, through shared platforms and translated content, to be permeable to neoreactionary imaginaries, yet distinct enough in its republican vocabulary and its own far-right genealogies to translate and reconfigure what it borrows (Boursier, 2022). A French far-right YouTube ecosystem is therefore neither sealed against neoreaction nor a blank surface for it, so whether the current actually takes root there is something only the empirical record can settle.

What place does neoreactionary discourse occupy in what France's far-right YouTube influencers produce, and is that place growing over time? Presence alone of the neoreactionary discourse in the far-right YouTube ecosystem would settle little, since fragments of almost any ideology can be found somewhere online. The question we ask is instead one of *movement* over time, of a trend measured against the classical far-right repertoire that neoreaction would have to push aside, such as nation, immigration, authority, tradition, democracy and progress. Tracked this way, the French far right can do one of two things with the imported neoreactionary discourse. It can take the current up and make itself its standard-bearer, drifting steadily in its direction and adopting the themes most its own, such as libertarianism or techno-capitalism, so that the neoreactionary themes gain over the period until the distance to the classical repertoire closes, the accretion that adjacent work has found in other reactionary settings (Hemmila and Perliger, 2024; Behr and Gianoncelli, 2025). Or it can hold its shape, the classical repertoire keeping its dominance while neoreaction lingers as a marginal and conjunctural presence, the two running in parallel rather than converging.

Telling those trajectories apart calls for an archive large enough to cover the whole ecosystem together with a way of reading what is really *said* inside the videos, sentence by sentence, a combination that research on online reaction has seldom had at its disposal (Ganesh, 2025; Abrahamsen and Williams, 2024). YouPol, a continuously updated archive of French political YouTube (Lemor and Boursier, 2026), supplies both, since every transcript sentence it holds is already annotated, by a family of classifiers distilled from a large language model and validated against human coders, the IDS-NRx (Boursier and Lemor, 2025; Lemor et al., 2026), for eleven ideological themes: the six that define the classical far right, nationalism, immigration, authority, tradition, democracy and progress, and the five through which neoreaction enters, libertarianism, technology, ecology, the (in-)equality theme it shares with the classical far right, and fictional metaphors (i.e., politics cast through fiction and myth such as the *Matrix* "red-pill" image of an awakening from illusion). From this corpus we take the twenty channels YouPol classifies as French and far-right and keep their political sentences over 2016 to 2024, the window across which the ecosystem is active. On that sample we proceed in two steps: we track the yearly prevalence of each theme, to see whether the

neoreactionary themes gain on the classical ones, and we then estimate how the density of each theme in a video relates to the engagement it draws and to the share of its comments that turn political, reading the sentences and comment threads themselves alongside every count so that no figure floats free of the talk behind it.

Our central finding is that the French far-right YouTube ecosystem domesticates neoreaction rather than adopting it. The share of content carrying at least one neoreactionary theme stands near twenty-seven percent in 2016, climbs to about thirty-one percent across the pandemic and "culture-war" window of 2019 to 2021, and falls back to roughly twenty-two percent by 2024, so that the population trend over the decade is essentially flat. What grows instead is the "classical" repertoire of the far-right, organised around immigration, whose prevalence rises by about twenty points per decade over the recent years and pulls nationalism, tradition and authority along with it. Reading the labelled sentences shows why the two never merge, since most of what registers as neoreactionary proves on inspection to be the French far right's own preoccupations in a borrowed, neoreactionary-sounding idiom, the clearest case being libertarianism, which in a political culture as state-centred as France's almost never carries its defining principle of exit and surfaces instead as a defence of national sovereignty and of a protective state.

Audience engagement tells the same story. No neoreactionary theme enjoys a systematic platform advantage, and the theme that most reliably mobilises an audience is again a classical one, immigration, whose denser videos draw about six percent more comments per standard deviation of its prevalence. Where neoreaction does leave a distinctive mark is in how audiences respond rather than in how much is produced, since videos dense in technology and in fictional metaphors are followed, within the same channel, by comment threads that are markedly less political, by about thirteen and eleven percent, whereas immigration, nationalism and democracy raise that rate; neoreactionary content engages audiences in a register that classical political language does not capture, a measurable mark of an "anti-political" style.

The paper proceeds as follows. Section 1 develops the theoretical framework, articulating the three-way relationship between illiberalism, the far right, and neoreaction. Section 2 describes the YouPol corpus and the analytic sample. Section 2.3 presents the two measurement instruments, the IDS-ED and the IDS-NRx, and the distillation pipeline behind the latter. Section 2.4 sets out the empirical strategy. Section 3 maps the presence and longitudinal evolution of NRx themes in the French far-right YouTube ecosystem and the reception of that content. Section 4 discusses the implications for our understanding of ideological mutation and platform-mediated metapolitics.

1 Neoreactionary Discourse, Metapolitics, and the Digital Turn

1.1 The Neoreactionary Assemblage and Its Boundaries

The analysis developed in this paper rests on a precise articulation of three concepts that are frequently conflated in the literature: illiberalism, the far right, and neoreaction. Establishing their relationships clearly is not merely a theoretical exercise, it defines the structure of the corpus, the logic of the IDS instruments, and the empirical question the paper is designed to answer.

The broadest category is illiberalism, understood in the sense developed by Laruelle (2022, 2024) as a new and coherent, if doctrinally fluid, ideological universe constituted by its relational opposition to liberalism. Illiberalism is not a regime type but an ideological grammar (Demias-Morisset, 2024). As Demias-Morisset (2024) has shown, the conceptualisation of illiberalism is structurally complicated by the essentially contested nature of liberalism itself: what counts as

'illiberal' varies according to which grammar of liberalism one employs, and these grammars are not mutually consistent. We adopt Laruelle's framework as a deliberate analytical choice, privileging a relational and ideological definition over procedural or regime-type alternatives. Crucially, Laruelle insists that illiberalism is broader than the far right: it includes mainstream and governmental forces that do not qualify as extreme by conventional criteria, and whose audiences, strategies, and ideological toolkits differ from those of the radical fringe. Not all illiberalism is far-right, and not all far-right formations are straightforwardly illiberal, though the overlap is substantial.

Within this broader universe, the far right constitutes a specific and more radical node. Following Carter (2018), we define far-right ideology by three constitutive properties: opposition to democracy, nativism, and authoritarianism. This triple criterion, rather than any single dimension, is what distinguishes far-right formations from other illiberal currents. It is also the baseline criterion against which we qualify the influencers in our corpus: the channels we analyse are those YouPol classifies as far-right on precisely these grounds, and the IDS-ED measures the thematic expression of that profile..

Neoreaction occupies a third position in this architecture: it is a specific subcategory of the illiberal universe, but one whose relationship to far-right ideology in the classical sense is ambiguous and empirically open. As reconstructed by Miranda (2026), five core components of the NRx ideological grammar appear particularly structuring. First, NRx rests on the conviction that natural hierarchies, social, racial, and sexual, organise human societies in enduring ways, grounding a form of radical elitism justified through "realist" appeals to biological and cognitive differences. Second, it is marked by a pronounced anthropological pessimism: violence is seen as an irreducible feature of human coexistence, and political order is therefore oriented toward security and conflict management in a broadly Hobbesian sense. Third, this leads to an explicit and systematic rejection of democracy, understood as both inefficient and captured by a self-reproducing elite; in its place, NRx theorists often envision monarchical or corporate forms of governance. Fourth, the current incorporates a libertarian-inflected "exit" principle, whereby individuals are reconceived as clients free to leave one jurisdiction for another, thus substituting mobility for political participation. Finally, NRx is characterised by a form of techno-optimism: democracy and progressivism are interpreted as obstacles to technocapitalist development, while post-democratic arrangements are framed as accelerators toward a potentially transhumanist future.

What distinguishes NRx from far-right ideology as defined by Carter is not its illiberal character, NRx is unambiguously illiberal, but the specific form its anti-liberalism takes. The classical far right valorises authoritarianism, (traditional hierarchies) ; national nativism, and reject democracy . NRx is more precise. By contrast, it is explicitly anti-populist, elitist, and techno-capitalist: it does not seek to restore a national community but to dissolve democratic governance into an entrepreneurial logic compatible with the acceleration of capital. Where the far right could fears and mobilises the crowd, NRx distrusts it. These structural differences explain why NRx and far-right ideology require distinct measurement instruments, and why the convergence or divergence of IDS-ED and IDS-NRx scores within the same corpus constitutes a meaningful empirical question: it asks whether these two illiberal formations are hybridising in practice, or circulating along separate tracks (Hemmila and Perliger, 2024). This question of convergence versus fragmentation has been raised at the European level for reactionary currents more broadly (Gianoncelli, 2021); our paper provides an empirical test of it within a specific national and platform context.

Not all five components can serve as discriminating markers. The rejection of democracy is constitutive of both formations — it is one of Carter's three defining properties of the far right — and anthropological pessimism is a background disposition that rarely surfaces as a codable

sentence-level claim. The empirical tracking of neoreaction therefore rests on the components the classical far right does not own — the naturalisation of hierarchy, the exit principle, techno-capitalist acceleration — to which we add two registers through which the current characteristically travels online: the accelerationist, 'eco-realist' reframing of ecological politics and the ironic, mythological 'red-pill' register (Miranda, 2026). Section 2.3.1 operationalises this mapping into the five themes of the NRx bundle, while anti-democracy remains, by construction, part of the classical baseline; any anti-democratic talk in the corpus is thus counted against, not for, the diffusion hypothesis, which makes the test conservative.

1.2 The Political Influencer as Metapolitical Actor

To account for how NRx themes might penetrate the discourse of predominantly far-right influencers, we require a framework for theorising both the actors and the transmission mechanism. We propose to treat French YouTube political influencers as metapolitical actors, in the analytical sense developed in recent scholarship on far-right cultural strategies (Boursier, 2026; Bures, 2023; Ganesh, 2025).

The concept of metapolitics, as Boursier (2026) has shown, oscillates between a category of practice, deployed by far-right actors as a label for their own cultural strategy, and a category of analysis. We adopt it exclusively in the second register. In its analytical form, metapolitics refers to a set of discursive and cultural practices aimed at reshaping collective common sense over the long term, beneath and beside the terrain of formal institutional politics (Schilk, 2025, 2023). Boursier (2026) identifies four dimensions through which this reshaping operates: the construction of affectively compelling narratives, the creation of alternative diffusion channels, the production of cultural material, and the mobilisation of affect. The political influencer condenses all four in a single actor type: they construct narratives through editorial framing, operate an alternative diffusion infrastructure, produce sustained cultural material, and engage audiences through the affective registers of outrage, irony, and performed common sense.

What makes the influencer figure particularly significant for our purposes is the discursive position it occupies. The political commentator who presents themselves as "simply analyzing the news" benefits from an apparent neutrality that makes their ideological content more porous and persuasive than that of the declared militant, and more difficult to detect through conventional methods. This is why, in the context of NRx circulation specifically, intentionality must be treated as an empirical variable rather than a presupposition. Some influencers may be deliberate entrepreneurs of NRx ideas; others may function as structural amplifiers, reproducing and domesticating NRx imaginaries without identifying with the current. Between these poles lies a spectrum of partial, hybrid, and opportunistic engagements that mirrors the logic of reactionary coalition-building identified in adjacent work (Boursier et al., 2026). The IDS-NRx, by measuring the presence of themes in content rather than imputing ideological commitments to producers, is calibrated precisely to this empirical reality.

1.3 Digital Platforms and the French-Speaking Terrain

The circulation of NRx themes among illiberal French influencers does not occur in a technologically neutral environment. YouTube's architecture, its recommendation algorithms, watch-time optimisation, and monetisation incentives, creates structural conditions that favour specific types of content independently of their ideological valence. These dynamics are best understood through the lens of digital illiberalism as theorised by Dall'Agnola (2024): a set of technology-enabled practices that restrict deliberative space and reshape public discourse from within democratic structures, without

the overt censorship associated with digital authoritarianism. The "gamification of the public space" produced by engagement-maximising algorithms, amplifying polarising, emotionally charged, and identity-confirming content, creates conditions structurally favorable to the diffusion of illiberal imaginaries, including NRx themes. We treat this platform dimension as a contextual constraint: it does not determine what influencers say, but it differentially selects and accelerates what their audiences receive and engage with.

France constitutes a particularly revealing terrain for studying NRx penetration within an illiberal far-right ecosystem. Its historically dense far-right tradition, from the Nouvelle Droite to the Rassemblement National, has generated established ideological frameworks, diffusion networks, and audience communities within which NRx themes may find receptive structures (Schir and Laruelle, 2025). At the same time, French political culture maintains a specific relationship to the American intellectual sphere: close enough through shared platform ecosystems and translated content to be permeable to NRx imaginaries, yet sufficiently distinct in its political vocabulary, its republican universalism, and its far-right genealogies to produce systematic processes of translation and reconfiguration (Boursier, 2022). This process of discursive domestication has been documented in adjacent registers of French far-right culture (Schir and Laruelle, 2025), suggesting that the translation of NRx themes follows broader patterns of illiberal cultural adaptation. The question is therefore not simply whether NRx themes are present in the French illiberal YouTube ecosystem, since their presence is our starting assumption, but whether they are growing, which influencers carry them, and how they are adapted as they cross cultural and linguistic boundaries.

1.4 Expectations

This architecture yields one directional expectation that the rest of the paper is built to test, and it is comparative. If the neoreactionary current is diffusing into the French far right and hybridising with it, in the manner that adjacent work reports for other reactionary settings (Hemmila and Perliger, 2024; Behr and Gianoncelli, 2025), then the neoreactionary themes should occupy a growing share of these channels' production over time, and growing specifically against the classical far-right themes that the IDS-ED already captures, so that the distance between the two repertoires narrows as the neoreactionary themes gain ground. On the contrary, the competing expectation is, then, coexistence. The classical repertoire of nation, immigration, authority and tradition retains or extends its dominance while the neoreactionary themes remain a marginal and conjunctural presence, and the two move on separate tracks. The design follows directly from this contrast. We do not study the neoreactionary themes in isolation but measure them against the classical far-right themes in the same corpus, on the same scale and over the same window, so that "growing" is always growing relative to the repertoire the neoreactionary current would have to displace, and so that the test is whether the neoreactionary themes become more prevalent than the classical ones rather than merely whether they are present. A second expectation concerns reception: if the neoreactionary themes carry the anti-political grammar that the theory ascribes to them, their presence should change how audiences engage and not only what producers say, which we test on the comment threads.

2 Data and Methods

This section sets out the data and the methods. We describe the YouPol corpus and the analytic sample (§2.1–§2.2), the two nested instruments that label each sentence (§2.3), and the empirical strategy behind the prevalence, engagement and reception estimates (§2.4). Throughout, the quantitative measurement is paired with a qualitative reading: we read the transcripts and comment threads the classifiers label, interpret them against the political events of the period, and count the occurrences of diagnostic terms over time, so that every statistic is anchored in, and checked against, the text that produced it (§2.5).

2.1 The YouPol corpus

The empirical material comes from YouPol (the YouTube and TikTok Political Observatory and Longitudinal database), a permanently updated research infrastructure that records what political content creators say on video platforms and preserves material the platforms later remove (Lemor and Boursier, 2026). YouPol works as a pipeline, sketched in Figure 1. It identifies political channels, downloads their videos and the attached comment threads on a rolling basis, and archives them so that suppressions and edits do not erase the record. Each video is transcribed and diarised into speaker-attributed segments, those segments are split into sentences, and a battery of fine-tuned classifiers then annotates every sentence (i.e, a supervised language model that reads one sentence and assigns it a label). As of early 2026 the corpus holds more than twenty-five thousand videos from sixty-four channels, close to three million annotated transcript sentences, and more than seven million archived comments. All data access for this paper ismediated by the public YouPol application programming interface and its Python client, so that our analysis is fully reproducible and can be updated as the database grows.

Among YouPol's native classifiers is `pol_detect`, which decides, for every transcript sentence, whether it carries political content. It is trained on a deliberately broad definition of the political that follows political-science usage, covering current affairs, social issues, political actors, power relations and social norms rather than narrow party talk. On held-out data it reaches a macro F_1 of 92.2% and an accuracy of 93.5%, with a human-to-model agreement of Light's $\kappa = 0.787$ on a thousand-sentence overlap sample (Lemor and Boursier, 2026). Every analysis in this paper begins from this gate. We keep only the sentences labeled as *political*. The reason is that the object of the study is the ideological content of political talk, and far-right creators interleave that talk with gaming, chat, product placement and personal aside, so measuring theme prevalence over undifferentiated transcript would dilute the ideological signal with material that is simply off-topic. Restricting to political sentences is what makes a per-video or per-channel composition of themes interpretable, and because the gate is native to YouPol we inherit an audited, reproducible definition of political speech rather than adding one of our own.

2.2 The analytic sample

Because the research question concerns the French far right, the corpus we analyse is the complete set of channels that YouPol classifies as French (`country=FR`) and far-right (`ideas=Far_right`) on YouTube, twenty channels in all. We take every such channel rather than a stability-filtered subset, since the question is about the full ecosystem. We focus on a specific time window of the corpus: 2016–2024. As Figure 2 shows, videos before 2016 the corpus only carry less than two percent of all political sentences, so reading a trend off that period would describe one voice rather

than a field. After 2024 the database holds only a handful of collection stragglers, fewer than sixty videos against more than two thousand in 2024, because harvesting after 2024 is still running at the time of writing. We therefore fix the window at 2016 to 2024, the span over which at least five channels are active each year. This choice also excludes, without any ad hoc rule, the two channels that exist only after 2024.

Within the window the unit of analysis is the video, and we keep videos with at least ten political sentences. The threshold follows from the per-video distribution, whose median is nineteen political sentences: a cut at ten removes the noisy short tail, where a per-video theme share would rest on one or two sentences, while retaining about ninety-six percent of all political sentences and roughly fifty-five hundred videos. Restricting to substantively political videos is what makes a per-video composition meaningful, because the object of the paper is the ideological content of political talk and not the share of a channel's output that happens to be political. The resulting sample is unbalanced, with one channel contributing a large fraction of the sentences. As a consequence, every estimate below is reported both within channels, through channel fixed effects, and across channels given equal weight, and we probe the influence of the dominant channel directly in Appendix B.

2.3 Measuring far-right and neoreactionary discourse

We measure ideological content with two nested instruments built on a common coding tradition. The first is the IDS-ED, the Ideological Diffusion Score for the Extreme Right, the name under which the YouPol infrastructure carries the *Score Idéologique d'Extrême Droite* (SIED) of Boursier and Lemor (2025), and it fixes the baseline far-right profile of the corpus. The second is the IDS-NRx, on which this paper rests, which extends that instrument to the neoreactionary motifs and lets us track the two repertoires on a single scale. A *theme*, in the sense we use throughout, is a recurring and theoretically specified element of a worldview, immigration or natural hierarchy for instance, that a single sentence can carry and that a classifier can detect independently of the others, and these themes are the common unit on which every measurement in this paper rests. The two instruments share a scheme and a measurement logic and differ only in coverage, so the rest of this section builds the IDS-NRx end to end, from the coding scheme to the classifiers we run on the full corpus.

We construct the IDS-NRx by knowledge distillation, following the hybrid annotation pipeline of Lemor et al. (2026). A capable but slow large language model, the *teacher*, annotates a large sample of sentences against an explicit codebook. We then fine-tune classifiers, the *students*, on those labels. The students recover most of the teacher's competence at a small fraction of the inference cost, which matters because the full corpus runs to millions of sentences, and they remain auditable against human judgement. Figure 3 lays out the five steps, and the subsections below take them in order: we define the scheme (§2.3.1), annotate with the teacher (§2.3.2), validate the teacher against human coders (§2.3.3), assemble the distillation corpus (§2.3.4), and finally train

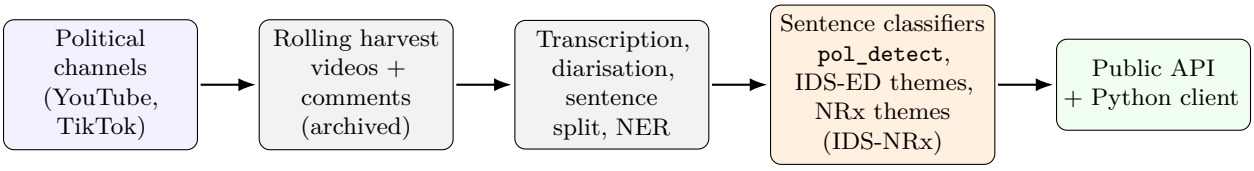

Figure 1: The YouPol infrastructure. Political channels are harvested continuously, their videos and comments archived against later removal, transcribed and diarised, and annotated at the sentence level by the `pol_detect` political-speech classifier and the IDS-ED theme classifiers. Everything is served through a public API (Lemor and Boursier, 2026).

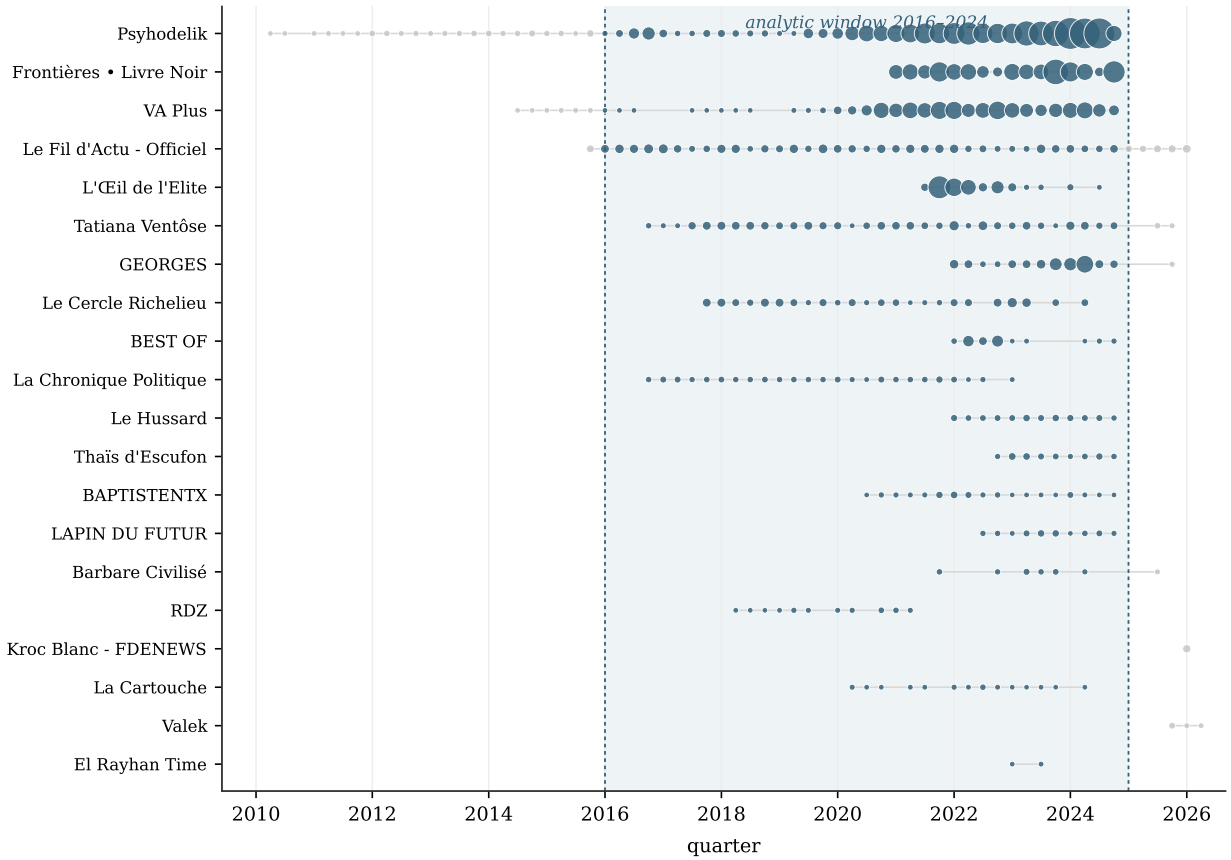

Figure 2: The analytic sample: the twenty French far-right channels (rows, ordered by political-sentence volume) and their activity over time, with marker area proportional to a channel's political-sentence volume in each quarter. The shaded band is the 2016–2024 analytic window. The pre-2016 mono-channel tail and the post-2024 collection stragglers (grey) fall outside it.

and evaluate the students (§2.3.5).[1]

2.3.1 The annotation scheme: extending the SIED

We start from an explicit, theory-driven coding scheme. Our scheme extends the IDS-ED, which is simply the English name for the *Score Idéologique d'Extrême Droite* (SIED) introduced by Boursier and Lemor (2025), the same French-language far-right ideological score under a translated label. Building on Carter's and Mudde's definitions of the far right and on Freeden's morphological theory of ideology, it operationalises far-right ideology as seven *ideological-membership categories* of the far right (in the original French, *catégories d'appartenance idéologique à l'extrême droite*): nationalism, anti-immigration, anti-democracy, anti-progress, authoritarianism, traditionalism, and anti-egalitarianism. Each category is refined by a closed list of contextual sub-dimensions; nationalism, for example, separates ethnic, familial, state-fused, vital, threatened, and colonial-nostalgic framings (Boursier and Lemor, 2025).

[1]Every artifact behind this section, namely the annotation prompt, the 100,000-sentence annotated pool, the rebalanced training set, the full training logs, and the scripts that regenerate every table and figure below, is released in the paper's public repository for end-to-end reproducibility with the LLM Tool package.

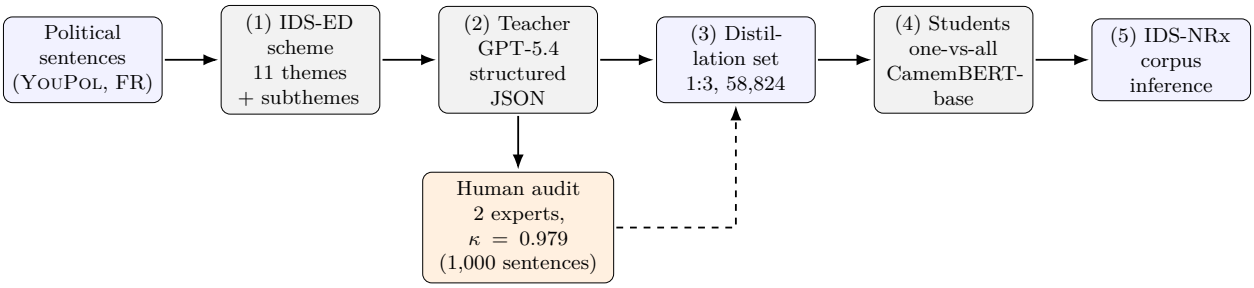

Figure 3: The IDS-NRx construction pipeline. A GPT-5.4 teacher annotates YouPol political sentences against the IDS-ED codebook; its labels are validated on a human-coded audit (dashed gate) and then distilled into eleven one-vs-all CamemBERT-base classifiers used for full-corpus inference. Steps map to the subsections of Section 2.3.

We keep these seven categories and add four themes specific to the Anglophone neoreactionary genealogy: ecology (accelerationist, "eco-realist" framings), libertarianism (the exit principle and jurisdictional arbitrage), technology (acceleration and transhumanism), and fictional metaphors (ironic, mythological, and "red-pill" registers). The result is a scheme of eleven sentence-level themes, each with its own subcategory list, which Table 1 summarises with a short definition and an illustrative example. Five of the eleven make up the NRx bundle that this paper follows: the

Table 1: The eleven IDS-NRx themes: what each captures and an illustrative excerpt drawn from the corpus (NRx themes marked †; excerpts are short translations from the French, with the source channel in parentheses).

Theme	Definition	Illustrative excerpt from the corpus
Nationalism	the nation as an ethnic or organic community, or as under threat	"they want to pull down statues, to revisit our history" (Frontières / Livre Noir)
Immigration	immigration framed as a threat to identity or security, or as needing stricter law	"the great replacement is already done; in twenty years it is over" (Frontières / Livre Noir)
Democracy	democracy as a regime judged corrupt, inefficient, or to be left behind	"today's elected officials are puppets that have been put in place" (VA Plus)
Progress	social or moral progress framed as a threat to identity, or as something to slow down	"they call it safe space, anti-racism, progressivism" (Psyhodelik)
Authority	the strong leader, order and security, the army and the police	"harsher sentences, and more resources for the gendarmerie" (Frontières / Livre Noir)
Tradition	traditional values, the family, heritage, a civilisational project	"I am not Christian; I simply defend Christian civilisation" (Le Fil d'Actu)
Equality†	anti-egalitarianism: natural hierarchies, inequality as an ordering principle	"there is a natural hierarchy, and the left cannot stand it" (Thaïs d'Escufon)
Ecology†	accelerationist or "eco-realist" reframings of ecological politics	"the planet has always had warming periods and cooling periods" (Frontières / Livre Noir)
Libertarianism†	the exit principle, jurisdictional arbitrage, autonomy against the state	"a topic among libertarians: the secession of a deep-state state" (Psyhodelik)
Technology†	acceleration, transhumanism, techno-optimism	"the first step toward a Chinese-style social credit system" (Psyhodelik)
Fict. metaphors†	ironic, mythological, fictional, and "red-pill" registers	"the courage to sacrifice your old life to step out of the matrix" (Thaïs d'Escufon)

four new themes, together with equality, which is the IDS-ED's anti-egalitarianism read here as the polarity-inverted entry point for NRx ideas. The remaining six themes (authority, democracy, immigration, nationalism, progress, and tradition) anchor the instrument in the classical far-right repertoire and serve as contrastive baselines.

One coding rule matters for how the results should be read. Following Boursier and Lemor (2025), a theme is coded present when the idea is actually raised in the passage, that is, when the sentence genuinely engages the theme rather than brushing against a related word in passing; the subcategories then record which specific far-right or neoreactionary framing, if any, the sentence carries. The scheme favours precision and restraint over exhaustiveness, because our aim is to estimate how often each theme occurs across a very large corpus, not to capture every faint allusion.

2.3.2 Teacher annotation with a large language model

With the scheme fixed, we use GPT-5.4 (model `gpt-5.4-2026-03-05`) as the *teacher*, driven by a structured prompt that states the annotator's role and the measurement objective, defines the eleven categories and their subcategories, gives positive and negative examples, and imposes a strict output format: one JSON object per sentence, listing all eleven themes, each with a binary detection flag and a possibly empty list of subcategory codes. Annotation is multi-label, since several themes can co-occur in a single sentence.

We first restrict the YouPol corpus to political speech using the `pol_detect` classifier, then draw a sample of 100,000 political sentences and annotate it with the teacher through the OpenAI batch interface (temperature 1.0). The run proceeded in successive batches between April and May 2026; each model response is checked against the expected schema, with automatic repair and re-submission of the small fraction that fail on the first pass, so that malformed outputs are recovered rather than dropped. Routing annotation through an explicit codebook shifts the human effort from labelling tens of thousands of sentences by hand to designing the scheme and auditing its quality, and it produces labels that stay consistent across the whole sample (Lemor et al., 2026).

2.3.3 Validating the teacher against human experts

Before trusting these labels, we check them against human judgement. Two domain experts (the two first authors of this paper) independently annotated a separate sample of 1,000 sentences with the same scheme. The two coders agreed almost perfectly (Light's $\kappa = 0.979$, with a pairwise F1 of 98.3%), which gives us a reliable human consensus to compare against. On that consensus, GPT-5.4 detects each of the eleven themes with an F1 between 96.9% (democracy) and 99.9% (fictional metaphors), for a micro-averaged F1 of 97.7%. The macro-averaged F1 is lower, at 70.0%, but this reflects the much harder and very sparse *subcategory* labels, not theme detection, which is what the IDS-NRx relies on. We therefore treat the teacher's theme-level labels as a sound basis for distillation. Per-theme figures are reported in Appendix A (Table 3).

2.3.4 The distillation corpus

We then turn the teacher's labels into a training set. Across the annotated pool, themes prevalence ranges from 5.3% for nationalism down to 0.1% for fictional metaphors, as Appendix A (Table 2) reports in full. On data this skewed, a classifier could score well simply by almost always predicting "no". To prevent that, we rebalance the data to a 1:3 positive-to-negative ratio. We keep every

theme-positive sentence and sample three theme-negative sentences for each, drawn so as to preserve channel and length diversity. The result is a multi-label training file of 58,824 French sentences, which the trainer expands into eleven independent one-vs-all problems, for a total of 24,198 positive and 72,594 negative theme decisions, each split 80/20 into training and validation. Rebalancing changes only the base rate the classifier sees while learning; it leaves untouched the human audit of §2.3.3, against which calibration is ultimately judged.

2.3.5 Training and evaluating the student classifiers

We then train each theme with its own binary (one-vs-all) classifier, a design we found more robust on these sparse, heterogeneous themes than a single shared multi-label head. The backbone is CamemBERT-base, a French RoBERTa encoder that we selected after comparing several French encoders and finding it clearly stronger on these themes than deeper but less well-adapted alternatives. We fine-tune with AdamW (learning rate 2×10^{-5}, batch size 16, weight decay 0.01), truncate sentences at 256 tokens, and keep, for each theme, the epoch that performs best on the validation split. Training runs through the LLM Tool's Training Arena, which records every configuration, data split, and metric so that the run can be reproduced.

The distilled classifiers transfer the teacher's competence well. Across the eleven themes the mean positive-class F1 is 86.8%, ranging from 80.6% for authority to 93.4% for ecology, with a two-class macro-F1 between 87.1% and 95.6%; the five NRx themes that carry the analysis are

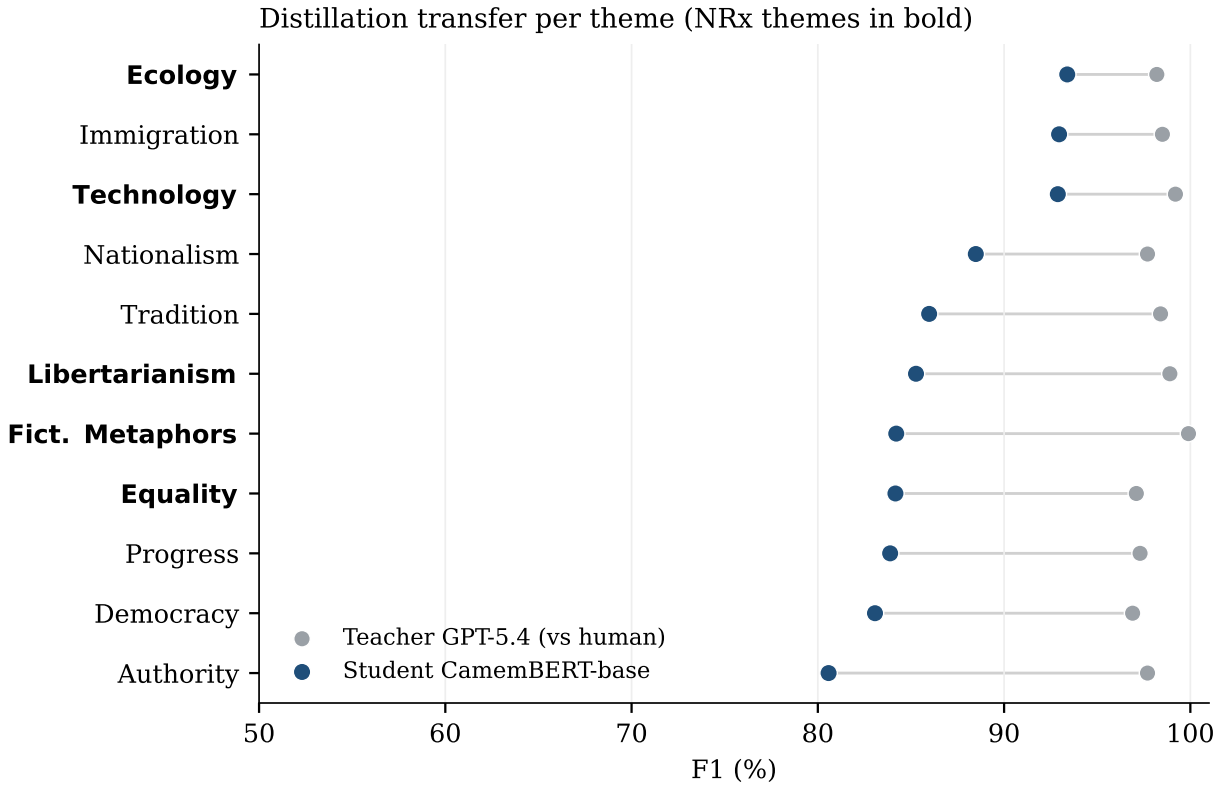

Figure 4: Distillation transfer per theme: F1 of the GPT-5.4 teacher (against human consensus, $N = 1{,}000$) and of the distilled CamemBERT-base student (validation), sorted by student F1. NRx themes in bold. The students recover most of the teacher's competence at a fraction of the inference cost.

detected at 84 to 93% F1. Figure 4 compares teacher and student F1 theme by theme, and the full per-theme student metrics appear in Appendix A (Table 4). These eleven classifiers form the operational IDS-NRx that we apply to the full corpus in Section 3.

2.4 Empirical strategy

We work at the level of the video and, for reception, at the level of its comment thread. Let P_{ik} be the classifier probability that political sentence i carries theme k, and let video v contain n_v political sentences.

Prevalence. We summarise a theme in a video in two ways. The thresholded share is the fraction of the video's political sentences assigned to the theme above the classifier's own decision boundary,

$$s_{kv} \;=\; \tfrac{1}{n_v} \textstyle\sum_{i \in v} \mathbf{1}[P_{ik} \geq \tau]\,, \qquad \tau = 0.5,$$

and the probability-weighted prevalence is the mean predicted probability, $\bar{p}_{kv} = \frac{1}{n_v} \sum_{i \in v} P_{ik}$, which uses the full classifier output and needs no threshold. For the neoreactionary bundle we also use the sentence-level indicator that at least one of the five themes is present, $\mathbf{1}[\max_{k \in \mathrm{NRx}} P_{ik} \geq \tau]$, and its yearly mean. The thresholded and probability-weighted measures agree throughout (Appendix B), so the substantive claims do not depend on the threshold.

Trends. For each theme we regress the per-video measure on upload time,

$$y_{cv} \;=\; \alpha + \beta\, \mathrm{time}_v + \gamma_c + \varepsilon_{cv},$$

by ordinary least squares with standard errors clustered on the channel. With the channel fixed effects γ_c included, β is the within-channel trend, the change in a given channel's own content over time; without them it is the population trend, which also reflects the changing composition of active channels. Because the panel is unbalanced we report both and treat their comparison as substantive rather than selecting one.

Engagement and reception. For count outcomes (views, likes, comments, and the number of political comment sentences) we use Poisson pseudo-maximum-likelihood (Silva and Tenreyro, 2006), which is consistent for the conditional mean under arbitrary overdispersion and admits zeros, with channel and upload-year fixed effects and the log video duration as a control. The theme regressors are the per-video probability-weighted prevalences $\bar{p}_{kv}$, standardised, so each coefficient is a rate ratio per one standard deviation of a theme and the eleven themes are comparable on one scale. For reception the outcome is the count of political comment sentences with an offset on the log number of comment sentences, so the model estimates the rate at which a comment sentence is political as a function of the video's theme mix. Videos that a channel posts in the same quarter share topical context and audience, so the natural unit of statistical dependence is the channel-quarter; we cluster the standard errors at that level throughout the engagement and reception analyses. This absorbs the dominant within-channel-within-period correlation while keeping the cluster count of order a few hundred, so the normal approximation is sound and we still use $t(G-1)$ critical values. Clustering on the eighteen channels alone is the familiar small-G regime, anti-conservative for some coefficients and over-conservative for others; we re-estimate every focal effect at three levels (channel, channel-year, channel-quarter) in Appendix B (Table 8) and report which conclusions depend on the choice.

Classifier-independent check. Independently of the classifier, we represent each sentence by its production Qwen3 embedding (available in the YouPol dataset) and score its cosine similarity to a prototype built from an external seed of canonical neoreactionary passages. A placebo prototype, the corpus centroid, guards against mistaking a generic movement of the embedding space for a drift toward the concept, and theme-specific Rocchio axes isolate each direction after the global component is removed (Appendix B).

2.5 Close reading and lexical evidence

The classifiers give breadth, but they do not by themselves say what a detected theme *means* in context, and they can attach one label to rather different content. We therefore pair every quantitative result with a qualitative reading of the material behind it. For each theme we read a sample of the transcript sentences the classifier assigns to it, together with the comment threads beneath the videos most dense in that theme, and we interpret them in light of the political sequence of the period (e.g., the elections, the pandemic, the 2020 statue removals, the June 2023 riots and the 2024 immigration law, etc.) so that a rise or fall in a curve can be read as a recognisable discursive shift rather than a bare number. For the reception analysis in particular we read individual videos against their own comment threads, setting what a host says beside how that video's audience replies, which is what makes the production-to-reception relationship of Section 3 legible. This reading is what licenses the substantive claims of Section 3, in particular that much of what the classifier flags as neoreactionary is the French far right's own repertoire in a borrowed idiom: the classifier locates the content, and the reading establishes its sense. We complement both with simple lexical counts, the yearly frequency of diagnostic expressions such as *grand remplacement*, *woke* and *civilisation*, which are transparent, independent of any model, and corroborate the classifier-based trends. Every excerpt we quote is drawn verbatim from the corpus and translated from the French, with the source channel and year in parentheses.

3 Results

Section 1.4 set out a comparative expectation: if the neoreactionary current is diffusing into the French far right, its themes should gain ground on the classical repertoire over the window, so that the distance between the two narrows. This section tests that expectation and then turns to reception. The empirical strategy, including the trend, engagement and reception models, is set out in Section 2.4 and is not repeated below; each subsection states what it tests and reads the relevant coefficients.

3.1 No secular neoreactionary drift

We begin with the neoreactionary bundle as a whole and test whether the share of content carrying at least one of the five themes grows. Figure 5 plots that yearly share, with the quarterly trajectory smoothed in the background. The series sit near 27% in 2016, dips in 2018, rises to about 31% across the 2019 to 2021 culture-war and pandemic window, and then declines to about 22% by 2024, below its starting level. Fitted as a linear trend, the population trajectory is essentially flat, with a slope of -5.8 percentage points per decade that is not distinguishable from zero ($p = 0.49$). Holding the composition of channels constant through channel fixed effects, the within-channel trend is a significant decline of -14.2 points per decade ($p = 0.02$), so once we follow the same influencers over time their content becomes less neoreactionary rather than more. Read against the classical

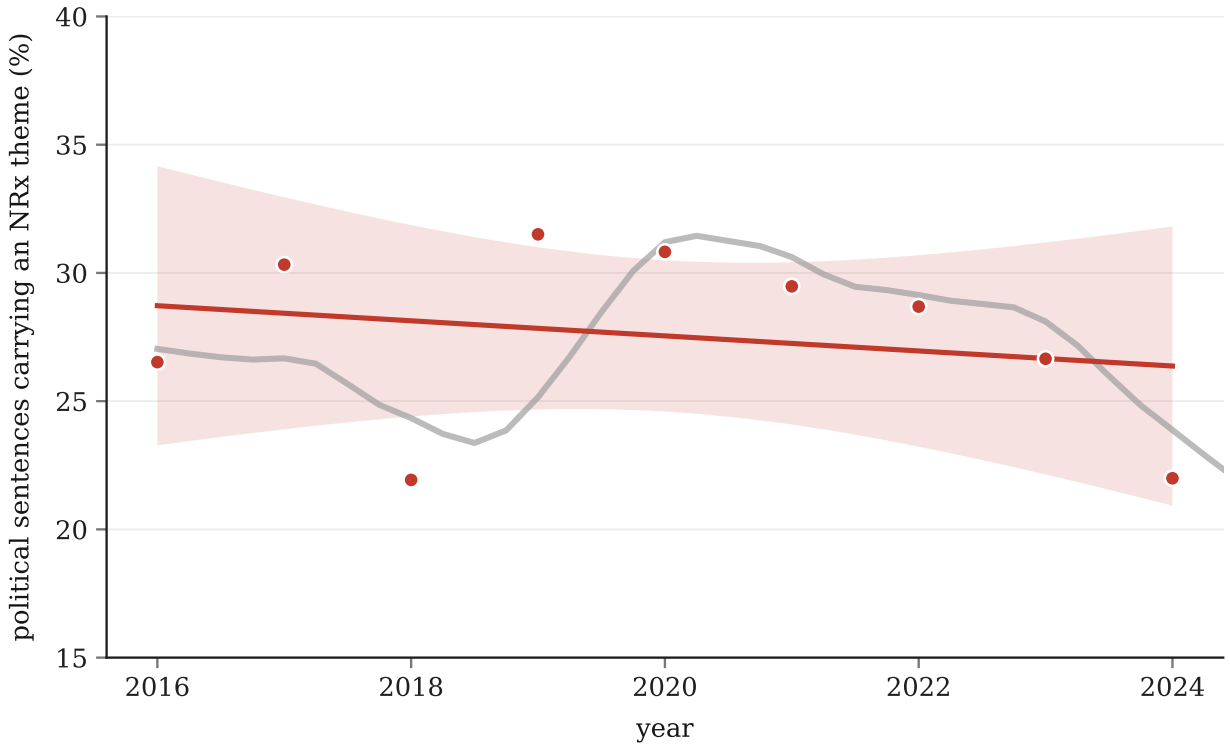

Figure 5: Yearly share of political sentences carrying at least one of the five neoreactionary themes (points), with the population least-squares trend and its 95% confidence band (red), over the quarter-by-quarter smoothed trajectory (grey). The share crests in the 2019 to 2021 window and recedes thereafter; the population trend is flat and the within-channel trend declines (Section 3).

far-right themes of the next subsection, which climb over the same years, this share moves in the opposite direction, so the distance that the diffusion hypothesis expects to narrow in fact widens. A semantic measure built from sentence embeddings appeared at first to rise, but a placebo prototype unrelated to neoreaction rises by as much; once the global component is removed no theme direction increases. The check is reported in Appendix B and we do not rely on it.

3.2 The classical repertoire of the far-right intensifies

The eleven themes move on different schedules, each tracking the events of its moment (Figure 6). What unifies their recent movement is a re-nationalisation organised around immigration, the steepest riser of the set: its probability-weighted prevalence climbs from about three and a half percent of political content in 2016 to ten percent in 2024, a trend of roughly nine points per decade across the window that steepens to about twenty over the recent years (Figure 6). The register hardens alongside the volume, moving from a managerial discussion of migration policy in 2016 to 2018 toward an existential and demographic frame thereafter, and the literal phrase *grand remplacement* (great replacement) appears only a few times a year in 2017 and 2018 and then exceeds a hundred occurrences a year from 2021, tracking the candidacy of Éric Zemmour and the consolidation of Reconquête. The classified sentences speak for themeselves: *"how long before France becomes an African France, under the combined effect of demography and uncontrolled immigration?"* (Frontières / Livre Noir, 2022) and *"the great replacement is potentially already*

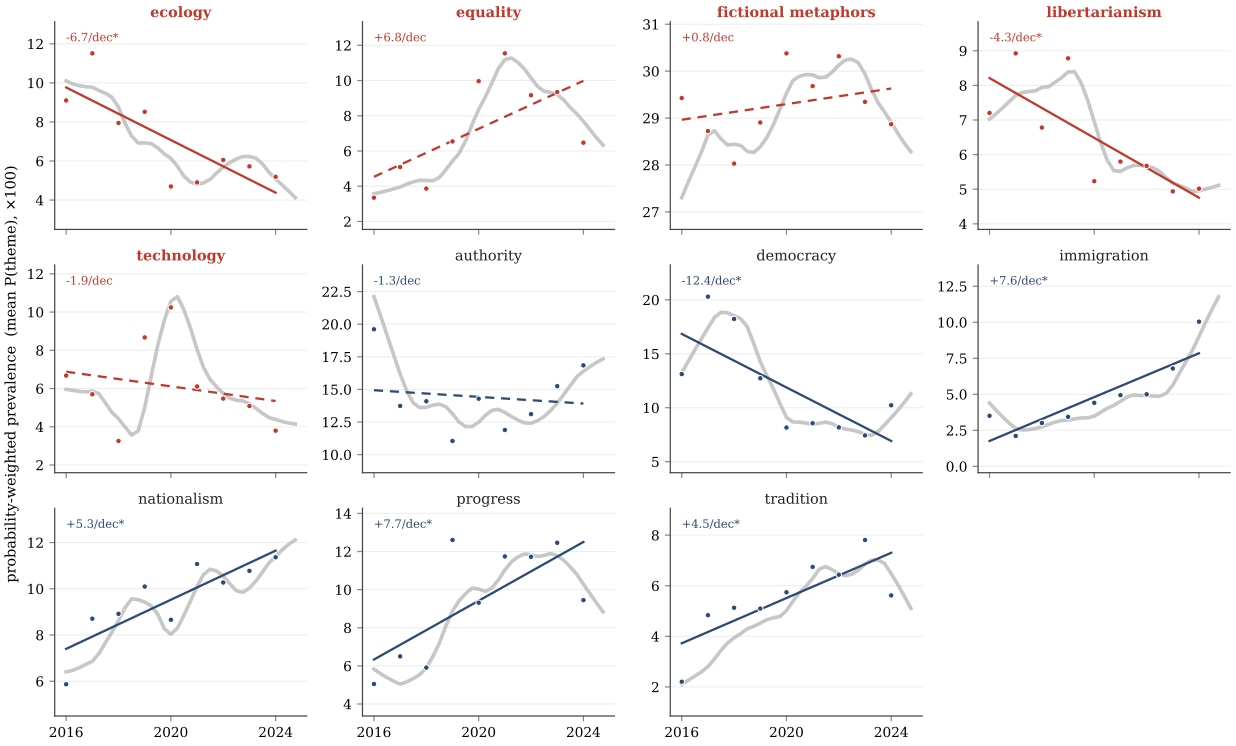

Figure 6: Probability-weighted yearly prevalence of each of the eleven themes (mean predicted probability, ×100), with the quarter-by-quarter smoothed trajectory (grey), the least-squares trend (line, solid where significant) and its slope per decade. Neoreactionary themes in red, classical far-right themes in indigo. The classical themes intensify over the recent years; the neoreactionary themes crest in the 2019 to 2021 window and recede.

complete; in twenty years it is over" (Frontières / Livre Noir, 2023), and after the riots of June 2023 and during the 2024 immigration law it fuses with a security sub-frame that conflates immigration and delinquency, *"firm measures must be taken to put these foreign delinquents permanently beyond the capacity to do harm"* (Psyhodelik, 2024).

The other classical themes are also increasing, and the figure shows each of them holding or gaining ground over the window (Figure 6). Authority is among the most prevalent themes of the corpus, near twelve percent on average; it is flat across the decade as a whole but rebounds sharply over the recent years, by about twenty-two points per decade, as an enforcement arm increasingly fused with immigration after 2023, *"more sentences, harsher sentences, a toughening of sentences, and more resources for the gendarmerie and the forces of order"* (Frontières / Livre Noir, 2023). Nationalism rises by some seven points per decade and shifts from a thin civic register toward a defence of heritage and a Christian-civilisational identity, reacting to the 2020 wave of statue removals, *"they want to pull down statues, they want to revisit our history"* (Frontières / Livre Noir, 2021) and *"how can you give people a country to love when you make everything that founded our history disappear?"* (VA Plus, 2024).

Tradition rises more modestly, by about three points per decade, migrating from the family-morality politics of the Manif pour tous years toward the same civilisational nativism, with the word *civilisation* climbing from a couple of dozen occurrences in 2016 to several hundred by 2023 as identity displaces faith, *"I am not Christian; I simply defend Christian civilisation"* (Le Fil d'Actu, 2021). The critique of progressivism climbs by about six points per decade across the window and marks the corpus's sharpest lexical discontinuity, the word *woke* being practically absent before 2020 and then exceeding four hundred occurrences a year from 2021 and peaking in 2024, the clearest sign that the French far right imported and nativised an Anglo-American culture-war vocabulary, *"these people want to bring segregation back, except that instead of calling it segregation and racism, they call it safe space, anti-racism, progressivism"* (Psyhodelik, 2021). Democracy is the most stable theme, roughly flat across the window though climbing by about nine points per decade in the recent years, an anti-elite populism already formed in 2016 that pours every grievance into the figure of a treasonous oligarchy, *"today's elected officials are puppets that have been put in place"* (VA Plus, 2023), and that acquires a new electoral-legitimacy object after the June 2024 dissolution.

3.3 The neoreactionary themes stagnate, and are largely domesticated

The five neoreactionary themes share a single shape, and none of them gains ground over the window (Figure 6). Each is near flat through 2016 to 2019, rises into the 2019 to 2021 culture-war and pandemic window, and recedes thereafter: technology falls by about ten points of prevalence per decade and fictional metaphors by about six, equality and libertarianism decline more gently, by roughly two and three points, and only ecology shows no significant trend, while the bundle as a whole declines by some fourteen points per decade within channels followed over time. A close reading of the sentences the classifier assigns to each theme shows why these levels matter less than they might seem, since most of that content is the French far right's own repertoire expressed in a register that resembles the neoreactionary one without sharing its content.

The technology theme, the steepest faller of the five at about ten points per decade, is carried, during its 2020 to 2021 peak, by a critique of digital surveillance and lost sovereignty rather than by techno-capitalist acceleration. It warns of biometric identity and a Chinese-style social-credit order, *"rejecting the European digital identity; the first step toward a Chinese-style social credit system"* (Psyhodelik, 2024), and when transhumanism appears it is named in order to be refused, *"I am against transhumanism and against all that nonsense"* (BAPTISTENTX, 2022). The

libertarianism theme, down about three points per decade, is dominated by French sovereigntism and decentralisation and by the discussion of libertarian and anarcho-capitalist ideas, taxation, private property and the minimal state, rather than by the principle of exit; the recurrent posture is to capture and defend the state rather than to leave it, *"what is a state for if not to organise the security of its citizens?"* (GEORGES, 2024). The few genuinely neoreactionary fragments are reported as foreign curiosities, *"it is a topic that comes up very often among libertarians, the secession of a deep-state state"* (Psyhodelik, 2024), or re-pointed at the populist cleavage between an elite and a people.

Ecology, the one neoreactionary theme without a significant trend, is overwhelmingly an anti-Green and pro-nuclear culture war, with the rare climate-sceptic strand reframing warming as a natural cycle independent of human action, *"there have always been periods of warming and periods of cooling of the planet, so even without humans the warming would occur"* (Frontières / Livre Noir, 2023). The equality theme, which falls by about two points per decade, fires mainly on manosphere anti-feminism and on anti-woke grievances about race and privilege; its clearest natural-hierarchy line is framed by gender, *"yes, there is a natural hierarchy, and the left cannot stand the idea of hierarchy"* (Thaïs d'Escufon, 2024). The fictional-metaphors theme, down about six points per decade, is the weakest and noisiest class; where the red-pill metaphor appears it is the manosphere red pill rather than the neoreactionary one, and the French awakening lexicon is the native *dissidence* and *grand remplacement* of Soral and Camus rather than the Cathedral.

Two structural features reinforce this reading. The themes are concentrated: a single high-volume essayist channel accounts for a large share of every count, so removing it thins the apparent mass considerably, which is the opposite of a movement diffusing across producers. And the genuine neoreactionary fragments, the secessionist exit, the accelerationist horizon, the Dark Enlightenment mythology, are present only homeopathically and as scattered citation of Anglophone discourse. The vocabulary travels as surface and irony during the culture-war window and is absorbed into the native repertoire rather than consolidating into an indigenous current.

3.4 Engagement profiles

We turn from what producers say to how audiences respond, asking whether platform incentives privilege any theme. Figure 7 relates the density of each theme in a video to its views, likes and comments, holding the channel, the upload year and the video duration constant, through a Poisson pseudo-maximum-likelihood model (Silva and Tenreyro, 2006) with standard errors clustered on the channel-quarter (Appendix B, Table 8). The effects are heterogeneous and theme-specific, and several themes, classical and neoreactionary alike, draw strong engagement; what differs is the currency. Comments are the contested one, and four themes raise them, each through a register of its own.

Immigration draws about six percent more comments per standard deviation (1.06, $p = 0.02$, robust at every level of clustering) by provoking political argument: beneath an hour-long interview on Franco-Algerian relations, ninety thousand views and some fourteen hundred comments split into a sustained debate, one side carrying the host's sovereigntist frame, *"Macron has literally destroyed all of French foreign policy"* (VA Plus, 2023), the other answering from an anti-colonial position, *"Driencourt is a colonial nostalgic... long live Algeria!"* (VA Plus, 2023). Equality raises comments even more (1.07, $p = 0.002$), through grievance and a contest to reframe: a montage of Black speakers rejecting anti-racism produces a thread led by a single inversion liked seven thousand times, *"is it just me, or is a white woman trying to prove to a Black woman that she should feel oppressed?"* (VA Plus, 2020), while the manosphere videos of an ex-identitarian host fill with men

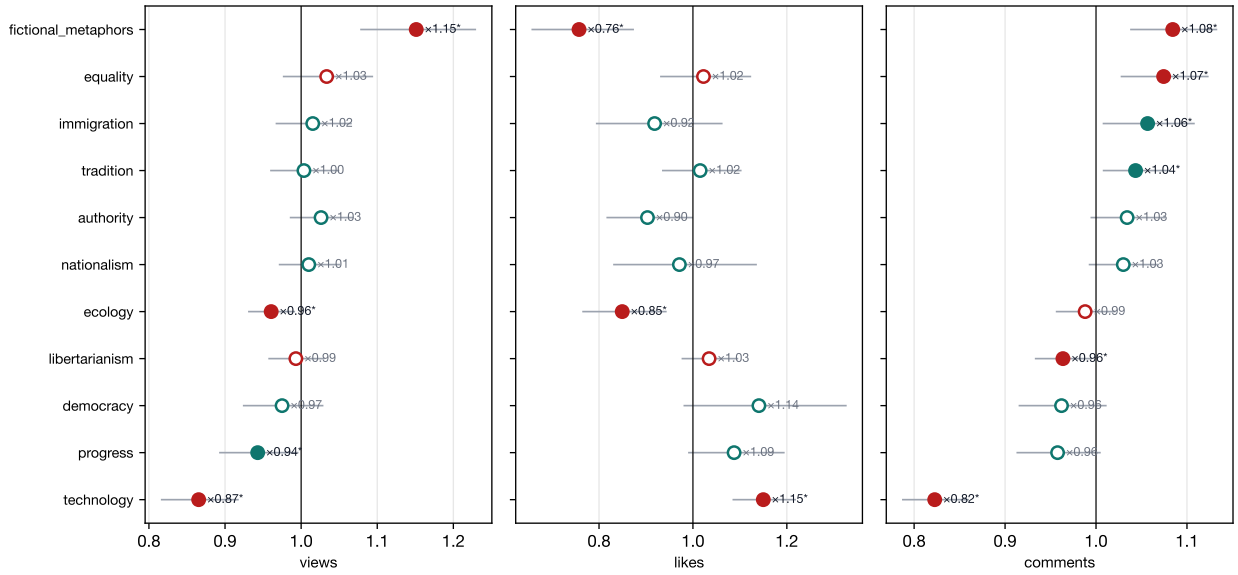

Figure 7: Theme to engagement: rate ratios per +1 standard deviation of each theme's per-video prevalence (PPML, channel and upload-year fixed effects; standard errors clustered on the channel-quarter, $t(G-1)$ critical values). Neoreactionary themes in red, classical themes in teal; filled markers denote a 95% interval excluding one. Effects are heterogeneous and no neoreactionary theme holds a systematic advantage; immigration draws more comments, while no theme carries a robust view premium.

volunteering their own divorce and custody stories. Tradition raises them (1.04, $p = 0.02$) as an identity ritual rather than an argument, the comments swearing allegiance to an embattled heritage, *"Zemmour, the man who keeps transmitting the values of our ancestors, while others trample them"* (Frontières / Livre Noir, 2021). And fictional metaphors raise them most of all (1.08, $p < 0.001$) through play: an essay casting Greta Thunberg as a modern Cassandra turns its own metaphor into a comment-thread game, *"I can see the manipulations with the naked eye"* (Tatiana Ventôse, 2019).

The other currencies separate the themes further. Fictional metaphors reach the widest audience, gaining fifteen percent more views per standard deviation (1.15, $p < 0.001$) while losing likes (0.76, $p < 0.001$): the ironic, pop-cultural wrapper carries a partisan claim to viewers who would never click on a tract, as in a comedic breakdown of "woke" cinema watched nearly a million times whose host opens by promising to *"relax with a compilation of the worst woke bad-buzz, and laugh about it, because laughter is the best weapon"* (VA Plus, 2023); the draw is the entertainment itself, and on a three-hour Psyhodelik exposé a viewer admits, *"I do not know any of the people involved, but it is extremely entertaining"* (Psyhodelik, 2020). Technology trades reach for loyalty, drawing fewer views (0.87, $p < 0.001$) but the most likes (1.15, $p < 0.001$): in these channels it attaches to gaming and consumer-technology content, and its small, devoted audience rewards expertise and humour rather than arguing politics, *"so, did the PlayStation 4 really flop in Japan?"* (Psyhodelik, 2020).

The engagement the neoreactionary themes attract is therefore real but tilted, toward reach, play and niche loyalty rather than the political argument that immigration provokes. This is already a sign of the pattern the next section develops: the audiences these themes mobilise come to be entertained, to compete at wordplay or to affirm an identity, more than to contest a political claim.

3.5 A neoreactionary signature in the comments

We test whether the themes change how audiences engage, not only what producers say. The one neoreactionary-specific pattern in our data concerns this reception side rather than prevalence (Figure 8). Using a Poisson model with an offset for the log number of comment sentences per video and channel and year fixed effects, we estimate how the density of each theme in a video relates to the rate at which its comment sentences are classified as political. Holding the channel constant, the classical themes raise that rate: democracy and nationalism increase it most, each by about seven percent per standard deviation, immigration by about four percent, and authority slightly, all significantly. Among the neoreactionary themes, technology and fictional metaphors lower it strongly, by about thirteen and eleven percent, and ecology modestly, while equality and libertarianism are flat; the critique of progressivism also lowers it slightly. The within-channel estimate is therefore a clean depoliticisation signature attached to the strongest neoreactionary themes, and it holds under every level of clustering (Appendix B, Table 8).

The descriptive record agrees with the within-channel estimate: across the videos most dense in each theme, about ten percent of comment sentences are political under technology and fictional metaphors, against twenty to twenty-three percent under immigration, nationalism and democracy, with the corpus baseline near seventeen percent (Appendix B, Table 7). Reading a video against its own comment thread gives the signature its content, and three cases, each drawn from the videos most dense in a theme, make the mechanism legible.

Where a video makes a concrete political demand, the thread takes it up and argues it. In a 2018

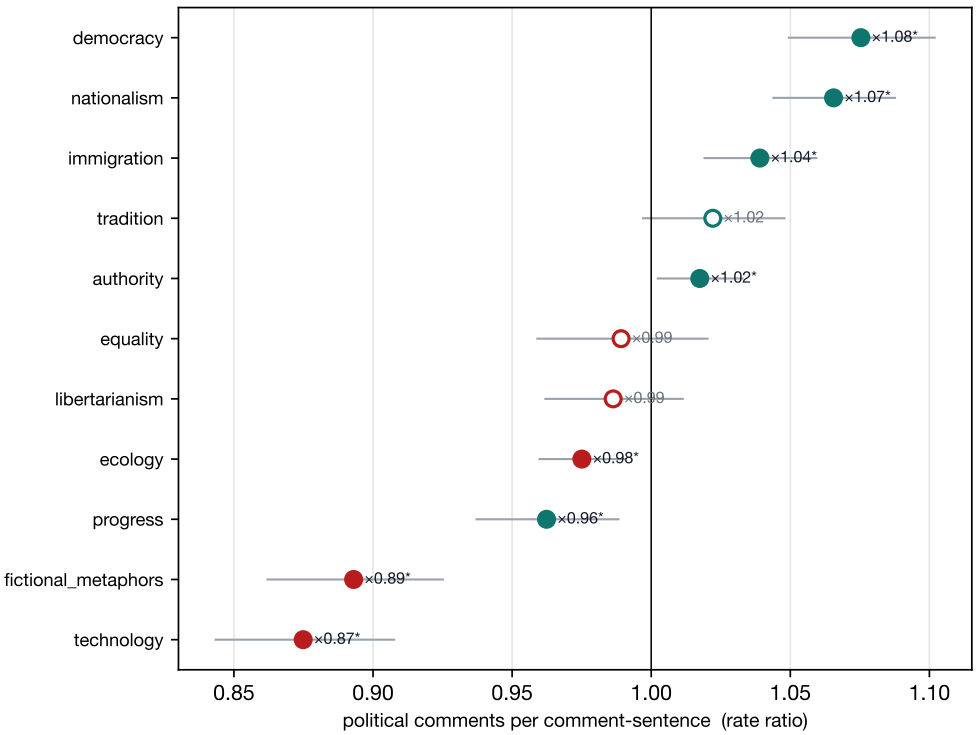

Figure 8: Theme to comment politicisation: political comment sentences per comment sentence (PPML with an offset on the log number of comment sentences, channel and year fixed effects; standard errors clustered on the channel-quarter, $t(G-1)$ critical values). Technology and fictional metaphors depress politicisation most strongly; the classical themes that move (immigration, nationalism, democracy) raise it.

video on the Mayotte crisis the host presents mass clandestine immigration as the cause and the end of birthright citizenship as the cure, *"we must challenge birthright citizenship so that the children of illegal migrants can no longer obtain French nationality, and so dry up the flow of immigration"* (La Chronique Politique, 2018), and the most liked comments then take up that exact measure and amplify it, *"birthright citizenship should be abolished in Mayotte"* (La Chronique Politique, 2018) and *"the government will do nothing: gutlessness and political correctness oblige, and Macron could not care less about Mayotte"* (La Chronique Politique, 2018). The political influencer discusses immigration and a policy, then amplified in the comments.

Where a video carries a neoreactionary-coded theme, the relationship is different. In these channels the technology theme attaches largely to gaming and consumer-technology content, and a 2021 video arguing that the loot boxes of a major game publisher amount to gambling, *"the fight against loot boxes continues, a study tying them to gambling"* (Psyhodelik, 2021), draws the least political thread of our cases, a hobbyist debate over games economics and consumer protection, *"give us back 2001, when you bought a game and got the entire game with no strings attached"* (Psyhodelik, 2021). The audience likes heavily and argues about publishers and hardware, not parties or policy. Where the register is fictional or literary the thread turns aesthetic: a 2023 Le Hussard video wraps a self-published novel in a noir framing narrative, *"in the cold, impersonal evening of the big city... a lone man writes"* (Le Hussard, 2023), and its viewers respond as a readership, even praising the writing *"is it me, or is the writing of this video even better than the others?"* (Le Hussard, 2023).

One qualification accompanies the result. The political-comment detector uses a deliberately broad definition of the political (§2.1), so the lower rate is not an artifact of an instrument tuned to party vocabulary: the gaming, ironic and literary threads are genuinely less engaged with current affairs, social issues and power. What the signature captures is a real shift in the *register* of engagement, toward affect, fandom and connoisseurship and away from political argument, rather than only a change in its quantity. Where neoreactionary content is not simply the French repertoire in a borrowed idiom, it attaches to objects, games, gadgets, films and novels, that pull audiences that way. We take up the interpretation in Section 4 and report the supporting checks in Appendix B.

4 Discussion

4.1 Coexistence rather than hybridisation

Over the observation window the French far-right channels do not reorganise their discourse toward the neoreactionary grammar. The share of content carrying any neoreactionary theme crests in 2019 to 2021 and then recedes, the population trend is flat, and the within-channel trend declines, so that the same influencers become less neoreactionary as they age rather than more (Section 3). What intensifies instead is the classical repertoire of the far-right, and it does so as a coherent bundle organised around immigration, with nationalism, tradition, the anti-progressivism polemic and authority moving with it. The relationship between the established far right that the IDS-ED measures and the neoreactionary current that the IDS-NRx measures is therefore one of coexistence rather than of the accretion or hybridisation that adjacent work has documented for other reactionary settings (Hemmila and Perliger, 2024; Mirrlees, 2025; Boursier et al., 2026). This qualifies, for the French YouTube case, the radicalisation-by-absorption dynamic that Behr and Gianoncelli (2025) describe in the conservative intellectual press, and it shows that the doctrinally fluid illiberal universe of Laruelle (2024) need not produce convergence of its currents in every national and platform setting.

4.2 Why the neoreactionary core does not take root

The substantive contribution of the paper is an account of why the neoreactionary programme fails to diffuse into a far right that shares its illiberalism. The close reading of Section 3 suggests that the obstacle is structural rather than circumstantial. The neoreactionary grammar is techno-capitalist, post-national, oriented toward exit rather than voice, and ironically indifferent to Christianity. The French far right in this corpus is statist, national, oriented toward the capture of a protective state, and Christian-civilisational. These are clearly opposed on the very axis that constitutes each of them, the relation to the nation-state. The consequence is legible theme by theme. Where neoreaction offers techno-capitalist acceleration, the political influencers offer a defence of the body and the nation against surveillance, and an explicit refusal of transhumanism. Where neoreaction offers the Cathedral and the red pill, the political influencer already holds a richer native mythology of decline, the *grand remplacement* and the *dissidence*, that occupies the same rhetorical slot.

The Anglophone vocabulary therefore is domesticated rather than adopted. Its surface, the irony, the contrarian posture, the citation of foreign figures, travels easily during the culture-war window and supplies a handful of fragments; its content is absorbed into a pre-existing French repertoire whose roots in the Nouvelle Droite's genealogies of difference (Gianoncelli, 2021) and in republican and nativist idiom (Schir and Laruelle, 2025; Boursier, 2022) are deeper than the imported lexicon. The French specificity, a dense reactionary tradition and a strong-state political culture, is what makes the far right permeable to the affect of neoreaction but resistant to its programme. This locates a boundary condition on accounts of straightforward transnational diffusion (Ganesh, 2025; Abrahamsen and Williams, 2024): a shared opposition to liberal modernity does not entail the circulation of a common doctrine when the receiving tradition is organised around the very institution, the nation-state, that the doctrine seeks to dissolve.

4.3 A depoliticisation signature and the anti-political grammar of neoreaction

Within a channel, denser neoreactionary themes are followed by less politicised comment threads, while the classical themes that move are followed by more politicised ones. The neoreactionary posture dissolves political claims rather than contesting them, recasting normative questions as matters of technical governance, jurisdictional choice or aesthetic and mythological representation, so that content built on technology, irony and literary allusion produces commentary that engages the world through affect and competence rather than through partisan argument. The case studies of Section 3 make this concrete and add a qualification. In these channels the technology theme attaches largely to gaming and consumer-technology talk and the fictional register to cinephile irony and literary essay, so audiences arrive as hobbyists and readers rather than as a political base, and the depoliticisation is as much a property of the objects these themes attach to, games, gadgets, films and novels, as it is a deliberate anti-political strategy. This is the metapolitical effect of Section 1 in its analytical sense (Boursier, 2026; Schilk, 2023), the operation on common sense that the scholarship on online far-right culture has described (Nagle, 2017; Marwick and Lewis, 2017; Tuters and Hagen, 2020), and it is consistent with the reading of neoreaction as an anti-political formation (Dall'Agnola, 2024; Miranda, 2026).

4.4 Limits

Several limits frame the interpretation. The corpus is concentrated: a single high-volume channel contributes a large share of the sentences, so volume-weighted statistics partly describe one voice,

which is why we lead on the within-channel and channel-balanced readings and report leave-one-channel-out checks in Appendix B. The classifier labels are calibrated against a human audit at the theme level, but their alignment with the theoretical neoreactionary grammar is imperfect, most visibly for the fictional-metaphors theme, a sparse and low-confidence class, and for the equality theme, which fires on manosphere and anti-woke content as much as on a philosophy of natural hierarchy; the close reading of Section 3 is the necessary complement to the counts. The reception result depends on a political-comment detector trained for classical language, and the cross-video descriptive shares are confounded by channel and format, so the depoliticisation claim rests on the within-channel estimate and on the qualitative register evidence rather than on raw shares.

Another limitation concerns the boundaries of the corpus itself. While our findings suggest that the reception of neoreactionary ideas in the French far-right ecosystem does not primarily occur through the major far-right YouTube channels included in the sample, this should not be interpreted as evidence of an absence of circulation. Rather, NRx themes may diffuse through a broader constellation of secondary actors that fall outside the scope of our corpus. Actors such as the *Rage Culture* collective, for example, have devoted substantial attention to neoreactionary authors and concepts while remaining peripheral to the ecosystem captured by YouPol. Similarly, some of the most influential intermediaries in the French reception of NRx ideas are not included in our dataset. Figures such as Julien Rochedy have played a visible role in introducing and translating Anglophone neoreactionary references for French-speaking audiences, often through formats and platforms that escape the boundaries of our sample. The absence of these actors means that our results speak primarily to the domestication of NRx within the mainstream far-right YouTube sphere rather than to the entirety of its French-language circulation.

A last limitation is temporal. Our analysis ends in 2024, just before a major phase in the public visibility of neoreactionary ideas. The reelection of Donald Trump and the renewed attention paid to figures associated with the American tech right generated a substantial wave of interest in NRx within the French far-right online sphere. In the months that followed, dozens of videos explicitly devoted to neoreactionary thought have appeared across French far-right YouTube. Because these developments fall outside our observation window, our findings should be interpreted as describing the period preceding this episode of intensified media attention. Future research will be needed to determine whether this post-2024 surge represents a temporary moment of curiosity or the beginning of a more durable integration of neoreactionary themes into the French far-right repertoire.

Conclusion

This paper asked whether neoreactionary themes occupy a growing place in the discursive production of French far-right YouTube influencers. The answer, formulated with the caution that the measurement demands, is negative. A theme classifier and a classifier-independent semantic test agree that the corpus does not drift toward the neoreactionary grammar over the 2016 to 2024 window. Holding the composition of channels constant, neoreactionary content declines rather than grows. What intensifies is the classical far-right repertoire, organised around immigration and carrying nationalism, the critique of progressivism, tradition and authority with it, in a re-nationalisation that tracks the political sequence of the period, from the woke debate and the Zemmour candidacy to the riots of 2023 and the immigration law and dissolution of 2024. The neoreactionary themes that the classifier detects are in large part the French far right's own concerns expressed in a register that resembles neoreaction without sharing its content, and the genuine neoreactionary grammar appears only at the margins and fails to consolidate.

The paper makes three contributions. Substantively, it documents a national reactionary tradition that intensifies on its own terms rather than converging on an imported programme, and it explains the limit on diffusion structurally: a far right organised around the nation-state cannot absorb a doctrine organised around dissolving it, however much the two share an opposition to liberal modernity. It also locates the neoreactionary footprint in the register of political engagement rather than in theme prevalence, through a within-channel depoliticisation signature that gives the anti-political grammar of neoreaction an empirical fingerprint. Methodologically, it shows that a distillation-based classifier can track fine-grained ideological content at scale while remaining auditable against close reading, and it reports a cautionary result for embedding-based drift measurement, where a cosine-to-prototype score can rise for reasons unrelated to the target concept. The same instruments and design are transposable to the other linguistic spaces of the YouPol infrastructure, beginning with Quebec, where whether the neoreactionary core diffuses more strongly than it does in France remains an open and testable question.

Appendices

A Detailed IDS-NRx training metrics

This appendix gives the per-theme detail behind Section 2.3. Table 2 reports the training-data distribution: how often each theme occurs in the 100,000-sentence annotated pool, and the size of the one-vs-all training set after rebalancing to a 1:3 positive-to-negative ratio. Table 3 reports the agreement of the GPT-5.4 teacher with the two-expert human consensus on the 1,000-sentence audit. Table 4 reports the validation performance of the eleven distilled CamemBERT-base classifiers. All three tables, and the figure in Section 2.3, are regenerated from the released artifacts by the table-building script noted above.

Table 2: Training-data distribution. *Pool* columns: positive sentences and prevalence among the 100,000 GPT-5.4-annotated political sentences. *Training* columns: the one-vs-all set after rebalancing to a 1:3 positive-to-negative ratio (all positives kept, three negatives sampled per positive).

	Annotated pool		One-vs-all training		
Theme	Pos.	Prev. (%)	Pos.	Neg.	Total
Nationalism	5,328	5.33	5,328	15,984	21,312
Authority	3,213	3.21	3,213	9,639	12,852
Equality†	3,052	3.05	3,052	9,156	12,208
Immigration	2,836	2.84	2,836	8,508	11,344
Democracy	2,739	2.74	2,739	8,217	10,956
Ecology†	2,317	2.32	2,317	6,951	9,268
Tradition	1,935	1.94	1,935	5,805	7,740
Progress	1,332	1.33	1,332	3,996	5,328
Technology†	700	0.70	700	2,100	2,800
Libertarianism†	646	0.65	646	1,938	2,584
Fict. Metaphors†	100	0.10	100	300	400
Total	—	—	24,198	72,594	96,792

†NRx theme.

Table 3: Teacher validation: agreement of the GPT-5.4 annotator with the two-expert human consensus on a held-out audit of 1,000 sentences (per-theme binary detection). Inter-annotator agreement between the two human coders was Light's $\kappa = 0.979$.

Theme	F1	Precision	Recall
Fict. Metaphors†	99.9	99.9	100.0
Technology†	99.2	98.8	99.7
Libertarianism†	98.9	98.1	99.8
Immigration	98.5	97.9	99.2
Tradition	98.4	97.6	99.3
Ecology†	98.2	97.8	98.7
Authority	97.7	97.1	98.4
Nationalism	97.7	97.0	98.4
Progress	97.3	95.5	99.2
Equality†	97.1	95.8	98.5
Democracy	96.9	95.5	98.3
Macro avg	70.0	73.2	72.0
Micro avg	97.7	97.0	98.5

†NRx theme. All values in %.

Table 4: Student classifiers: per-theme validation performance of the eleven one-vs-all CamemBERT-base detectors distilled from the GPT-5.4 teacher (positive class; 80/20 train/validation split). $F1_{macro}$ is the two-class macro-average.

Theme	F1	Precision	Recall	$F1_{macro}$	n_+
Ecology†	93.4	91.1	95.8	95.6	450
Immigration	93.0	92.2	93.7	95.3	556
Technology†	92.9	92.5	93.2	95.3	133
Nationalism	88.5	86.6	90.5	92.3	1,040
Tradition	86.0	84.8	87.2	90.6	376
Libertarianism†	85.3	84.6	85.9	90.2	128
Fict. Metaphors†	84.2	88.9	80.0	89.5	20
Equality†	84.2	84.6	83.7	89.4	603
Progress	83.9	80.1	88.1	89.1	260
Democracy	83.1	77.6	89.4	88.5	527
Authority	80.6	80.6	80.6	87.1	623
Mean	**86.8**	85.8	88.0	91.2	—

†NRx theme. All values in %.

B Robustness

This appendix records the checks behind the two substantive claims of Section 3, the absence of a neoreactionary drift and the depoliticisation signature. All are regenerated by the released scripts from the committed data.

The two prevalence measures agree. Every per-theme result holds whether a theme is measured by the thresholded share of sentences above the classifier boundary or by the probability-weighted mean prevalence, which uses the full classifier output and fixes no threshold. The figures in the main text are drawn with the probability-weighted measure and the trends with the thresholded one, and the two coincide in sign and ordering, so the conclusions do not depend on the cut.

No drift under any specification. Table 5 re-estimates the yearly NRx-any trend across a grid that crosses three activation thresholds (0.3, 0.5, 0.7), three minimum-volume filters (≥ 5, 10, 20 political sentences) and the two specifications, population and within-channel. No cell of the grid rises significantly; the population trend is flat throughout, the within-channel trend is flat to significantly negative, and the eighteen slopes range from -24.5 to -5.8 percentage points per decade. The within-channel decline is sensitive to the dominant essayist channel: under leave-one-channel-out the within-channel slope ranges from -16.8 to 0.0 points per decade and narrows to flat when that channel is removed, while never turning positive. The safe reading of the trend is therefore flat to declining, and never rising, on every cut of the data.

A classifier-independent semantic test. Table 6 reports the same video-level question with sentence embeddings rather than the classifier. The raw cosine to the external-seed neoreactionary prototype does rise over the window, but the cosine to a placebo prototype, the corpus centroid, rises by as much, which reveals a generic component: over the window the embeddings concentrate, so the average sentence moves toward the centre of the space and therefore toward any fixed prototype, including prototypes unrelated to neoreaction. The theme-specific Rocchio axes, which subtract the corpus centroid and isolate each direction, do not rise and decline for four of the five themes. A cosine-to-prototype drift score, common in text-as-data work, therefore requires a placebo and a discriminant axis before a rise can be read as a drift toward the concept of interest.

Inference for engagement and reception. Videos a channel posts in the same quarter share topical context and audience, so we cluster the engagement and reception standard errors on the channel-quarter, which keeps the cluster count of order a few hundred and avoids the small-G problem that clustering on the eighteen channels alone would create. Table 8 re-estimates every focal coefficient at three levels, the channel, the channel-year and the channel-quarter; the point estimates are identical and only the inference changes. The depoliticisation signature (technology and fictional metaphors lowering comment politicisation) and the immigration effects on comment politicisation and on comment volume are significant at every level. The small engagement premia, immigration on views and libertarianism on likes, are significant only when clustering on the eighteen channels, the anti-conservative small-G regime, and vanish once the cluster count is large, so we do not treat them as effects. Clustering on the eighteen channels is in turn over-conservative for the immigration politicisation effect, which it leaves just short of significance, a reminder that the small-G correction errs in both directions.

Table 5: Robustness of the no-drift result. The yearly NRx-any trend (percentage points per decade) under a grid crossing three activation thresholds, three minimum-volume filters and two specifications (population vs. within-channel), and under leave-one-channel-out. No cell rises significantly. $^{*}\,p < 0.05$.

Check	Value
Population trend (preferred spec)	-5.8 pp/dec
Within-channel trend (preferred spec)	-14.2 pp/dec*
Specification grid (18 cells): cells rising significantly	0
cells falling significantly	6
slope range across the grid	[-24.5, -5.8]
Leave-one-channel-out (within-channel): range	[-16.8, +0.0]

Table 6: Are videos becoming more neoreactionary? Video-level trend in NRx intensity (FR far-right YouTube videos with ≥ 10 political sentences; upload-time OLS, cluster-robust SE on channel). Three independent measures agree: once the global component (placebo) is removed, no NRx-specific measure rises. $^{*}\,p < 0.05$.

	Measure	Within-channel	Population
Classifier	share of sentences with $\geq$1 NRx theme	-14.19 pp/dec*	-5.84 pp/dec
Semantic (raw)	cosine to NRx prototype (centroid)	+3.07 %/dec*	+0.95 %/dec
	cosine to corpus centroid (placebo)	+1.41 %/dec*	+5.18 %/dec
Semantic (theme-specific)	Rocchio axis: equality	+0.26 %/dec	+24.62 %/dec
	Rocchio axis: technology	-7.26 %/dec	+62.31 %/dec
	Rocchio axis: libertarianism	-41.81 %/dec*	+131.42 %/dec
	Rocchio axis: ecology	-12.16 %/dec*	+10.09 %/dec
	Rocchio axis: fictional metaphors	+187.47 %/dec*	-320.96 %/dec

Reception, controlled and raw. Table 7 places the within-channel rate ratio next to a raw, uncontrolled quantity: the share of political comment sentences under each theme's top-decile densest videos. The two agree. The themes with a within-channel rate ratio below one are also those whose densest videos draw the fewest political comments, about ten percent under technology and fictional metaphors against twenty to twenty-three percent under immigration, nationalism and democracy, with the corpus baseline near seventeen percent. The channel and format confound appears only when the densest set is narrowed to a handful of videos dominated by one essayist channel, which is why the main text reads the signature off the within-channel estimate and the close reading rather than off the densest few videos.

Classifier calibration and its limits. The teacher-student validation of Appendix A detects each theme at high F1, which licenses the prevalence counts. Two themes are weaker proxies for the theoretical neoreactionary grammar. Fictional metaphors is a sparse, low-confidence class with an unusually high baseline probability across the corpus, so its level should not be read as a prevalence; and equality fires on manosphere and anti-woke content as much as on a philosophy of natural hierarchy. For both, the close reading of Section 3 is the necessary complement to the classifier, and the substantive claim about the neoreactionary themes does not rest on either of these two classes alone.

Table 7: Comment politicisation, two ways. *Within-channel*: the rate ratio from the PPML of Figure 8 (channel and year fixed effects). *Raw share*: the share of comment sentences classified political under each theme's top-decile densest videos, with no controls (baseline 16.9%). The two agree: the themes with a within-channel rate ratio below one are also the ones whose densest videos draw the fewest political comments, a convergent-validity check on the depoliticisation signature. The channel/format confound appears only when the densest set is narrowed to a handful of videos dominated by one essayist channel. $^{*}p < 0.05$ (channel-quarter cluster-robust). NRx themes marked †.

Theme	Within-channel RR	Raw share political (%)
technology†	0.87*	10.2
fictional metaphors†	0.89*	10.3
progress	0.96*	12.1
ecology†	0.98*	15.6
libertarianism†	0.99	16.0
equality†	0.99	11.7
authority	1.02*	17.8
tradition	1.02	16.2
immigration	1.04*	20.6
nationalism	1.07*	22.2
democracy	1.08*	22.5

†NRx theme.

Table 8: Inference robustness. Focal PPML coefficients (rate ratio per +1 SD of theme prevalence) with standard errors clustered at three levels: the eighteen channels, channel-year, and channel-quarter (the preferred level, used in Figures 7 and 8). Point estimates are identical across columns; only the p-value changes. The depoliticisation signature and the immigration politicisation effect hold at every level; the small engagement premia are significant only under the eighteen-channel clustering, the anti-conservative small-G regime. $^{*}p < 0.05$ (t($G-1$)).

Outcome	Theme	RR	p, clustered on		
			channel (G=18)	ch.-year (G≈90)	ch.-quarter (G≈290)
views	technology	0.87	0.008*	0.000*	0.000*
	fictional metaphors	1.15	0.084	0.000*	0.000*
	immigration	1.02	0.020*	0.603	0.546
likes	technology	1.15	0.003*	0.001*	0.000*
	libertarianism	1.03	0.015*	0.169	0.255
comments	immigration	1.06	0.010*	0.042*	0.024*
politicisation	technology	0.87	0.000*	0.000*	0.000*
	fictional metaphors	0.89	0.000*	0.000*	0.000*
	immigration	1.04	0.072	0.002*	0.000*
	nationalism	1.07	0.000*	0.000*	0.000*

References

Abrahamsen, R. and Williams, M. C. (2024). Radical right dystopias in the global culture wars. *Global Studies Quarterly*, 5(1):ksae091.

Behr, V. and Gianoncelli, E. (2025). The radicalization of 'intellectual conservatism' in the illiberal moment: The case of The European Conservative. *Journal of Illiberalism Studies*, 5(3):59–74.

Boursier, T. (2022). White supremacism on YouTube: How to rewrite history from a racist point of view. In Kaiser, C., Khiari, O., and Lühr, V. S., editors, *Temporalities of Diversity*, pages 65–87. Waxmann.

Boursier, T. (2026). Métapolitique d'extrême droite : Usages et limites d'un concept. *Canadian Journal of Political Science*, pages 1–21.

Boursier, T. and Lemor, A. (2025). Mesurer la pénétration des idées d'extrême droite dans les discours gouvernementaux français. une analyse des déclarations de politique générale (1959–2024). *Revue Française de Science Politique*, 75(2):261–291.

Boursier, T., Pronovost, V., and Lacroix, O. (2026). Intersectionalising hatred: How multiple antagonisms reinforce reactionary coalitions. *Journal of Gender Studies*.

Bures, E. (2023). The intellectual as culture warrior: Metapolitics and the european new right. *Fascism*, 12(1):1–26.

Carter, E. (2018). Right-wing extremism/radicalism: Reconstructing the concept. *Journal of Political Ideologies*, 23(2):157–182.

Dall'Agnola, J. (2024). Illiberal technologies: Linking tech companies, democratic backsliding, and authoritarianism. *Journal of Illiberalism Studies*, 4(3):1–10.

Demias-Morisset, R. (2023). La démocratie illibérale, un concept héritier des apories de la critique anti-démocratique du populisme ? In Blanc-Noël, N. and Dauphin, T., editors, *Vers un nouvel âge des extrêmes ? Populismes et transformations sociales*, pages 39–56. L'Harmattan.

Demias-Morisset, R. (2024). The paradoxical sources of illiberalism: A synoptic approach to the genealogies of illiberalism. *Journal of Illiberalism Studies*, 4(2):1–19.

Ganesh, B. (2025). The western far right and digital technology: Fuzzy collectivity from translocal whiteness to networked metapolitics. *Sociology Compass*, 19(2):1–15.

Gianoncelli, E. (2021). The unification of the 'new right'? on europe, identity politics and reactionary ideologies. *New Perspectives*, 29(4):364–375.

Hemmila, T. and Perliger, A. (2024). Hybridization or salad bar ideology? testing ideological convergence within the american violent far right. *Crime & Delinquency*.

Laruelle, M. (2022). Illiberalism: A conceptual introduction. *East European Politics*, 38(2):303–327.

Laruelle, M. (2024). Introduction: Illiberalism studies as a field. In Laruelle, M., editor, *The Oxford Handbook of Illiberalism*. Oxford University Press, 1 edition.

Lemor, A. and Boursier, T. (2026). YouPol: A collaborative research infrastructure and database for political content on YouTube and TikTok. Preprint, SocArXiv.

Lemor, A., Dinan, S., and Gilbert, J. (2026). LLM Tool: A hybrid pipeline for automated large-scale text annotation using local language models and BERT classifiers. Preprint, SocArXiv.

Marwick, A. E. and Lewis, R. (2017). Media manipulation and disinformation online. Technical report, Data & Society Research Institute, New York.

Miranda, A. (2026). *Les Lumières sombres*. Gallimard, Paris.

Mirrlees, T. (2025). The alt-right's discourse on 'cultural marxism': A political instrument of intersectional hate. *Atlantis*, 39(1):49–69.

Nagle, A. (2017). *Kill All Normies: Online Culture Wars from 4chan and Tumblr to Trump and the Alt-Right*. Zero Books, Winchester.

Norocel, O. C. (2023). Research bricolage on far-right metapolitics: Superordinate intersectionality perspectives on digital identities. *Innovation: The European Journal of Social Science Research*, 24(5):1–17.

Schilk, F. (2023). Metapolitics as programmatic storytelling. *Zeitschrift für Diskursforschung*, (2):339–347.

Schilk, F. (2025). The metapolitics of crises: How the new right weaponises narratives to mainstream far-right ideology. *International Journal of Politics, Culture, and Society*.

Schir, P. and Laruelle, M. (2025). Social and discursive capital as illiberal enabler: A tale of two far-right fictions in france. *Politics and Governance*, 13.

Silva, J. M. C. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Tuters, M. and Hagen, S. (2020). (((they))) rule: Memetic antagonism and nebulous othering on 4chan. *New Media & Society*, 22(12):2218–2237.