societies

MDPI

Article

Algorithmic Sacredness and Algorithmic Pluralism: Content Moderation and the Symbolic Dispossession of Contemporary European Paganism

Giuseppe Maiello [1,2,*] and Ondřej Roubal [1]

[1] Department of Social Sciences, University of Finance and Administration, 100 00 Prague, Czech Republic; ondrej.roubal@vsfs.cz
[2] Department of Humanities and Cultural Heritage, University of Campania "Luigi Vanvitelli", 81100 Caserta, Italy
* Correspondence: giuseppe.maiello@mail.vsfs.cz

Abstract

This paper examines mainstream platform moderation as it encounters contemporary pagan religious practice in Europe and reads what it finds as a symptom of a wider condition: the algorithm functions less as a neutral tool than as a productive instrument of an extractive political economy whose characteristic operation is the appropriation, classification, and revenue-conditioned filtering of human expression. Minority religious traditions are structurally exposed within it: too small to constitute an accommodated market, too polysemic for classifiers trained on majority devotional and Anglo-American extremism corpora, and too fragmented to extract policy concessions. Empirically, the paper draws on netnography of pagan online communities across Europe, informal conversations with 97 practitioners across nine European jurisdictions, and a corpus of 247 takedown notices and appeal exchanges (2019–2025). Three recurrent rationale-clusters—devotional content classified as occult, as extremism-adjacent, and as unsafe activity—are read as predictable outputs of the system's cost structure rather than as ordinary classifier errors. Two concepts are proposed. Algorithmic sacredness names the transfer of gate-keeping functions previously held by ecclesiastical, state, and editorial actors, routed through Bourdieu's meta-capital as extended to platforms. Algorithmic pluralism names a programmatic direction toward infrastructures whose governance is not capital's.

Keywords: contemporary paganism; Rodnoverie; platform capitalism; data colonialism; content moderation; digital religion; epistemic exclusion; meta-capital; algorithmic sacredness; algorithmic pluralism

Check for updates

Academic Editors: Gabriella Punziano, Felice Addeo, Angela Delli Paoli and Giuseppe Masullo

Received: 8 June 2026
Revised: 31 July 2026
Accepted: 3 August 2026
Published: 5 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

A Beltane fire-jumping video disappears overnight. A photograph of a Veles altar is restricted as "occult". A livestreamed Wheel-of-the-Year liturgy is interrupted because its multilingual chant fails an automated language-policy check. A Kolovrat emblem posted by an explicitly eclectic, self-ironic Czech Viking group is treated as belonging to the same symbolic class as identical-pixel posts by ethnonationalist Heathen channels in another European country. An appeal is submitted; a category label is returned without specifics; the post is restored or not; the practitioner adjusts. Multiply by years. Multiply by jurisdictions. Multiply across a movement whose communal life, since the late 1990s,

has been conducted on infrastructures built and owned by Anglo-American technology corporations.

The platforms on which contemporary European pagans live their communal lives do not exist to provide a religious public sphere; they exist to convert human expression into measurable behavior, behavioral prediction into targetable attention, and targetable attention into revenue. Couldry and Mejias [1] have named this data colonialism: an emerging order in which human life itself becomes a direct factor in capitalist production [1]; Couldry and Mejias's more recent Data Grab [2] sharpens the analogy with historical colonialism and proposes a three-register framework for resistance [3]. Within that order, moderation is not a downstream filter on a stable expressive world. It is the operating layer at which the political economy meets specific kinds of human expression and decides—by classifier, by policy, by appeal-routing logic, by regulatory pressure mediated back through the same proprietary infrastructure—what kinds of expression remain inside the revenue-generating perimeter.

Contemporary paganism in Europe—encompassing Wicca, Heathenry and Ásatrú, Druidry, eclectic Goddess spirituality, the Slavic Native Faith milieu (Rodnoverie, Rodzimowierstwo, Rodná víra), the Baltic native faiths, Hellenic reconstructionism, and the loosely affiliated milieu of Norse and Slavic-themed reenactors and "Vikingomania"—adjacent communities that Maiello and Cirklová [4] have analyzed under the heading of *Northilization*—is a diagnostic case for this layer because it falls on the wrong side of every term in the platform's optimization. The population is too small to be a market segment. Its iconography is polysemic in ways that an Anglo-American extremism corpus, on which classifiers are predominantly trained, cannot disambiguate. Its ritual forms do not fit the definition of "organized religious worship" in any of the categories the platform's policy team understands. Its internal pluralism—eclectic Slavic Native Faith circles, Czech Vikings' "Švejking"/"Lardcake" self-irony [4], gender-balanced reconstructionist groups, Goddess-centered communities—is invisible to classifiers whose categories were not built to register it. And the platform's political economy gives it no leverage to extract policy concessions. There is no advertiser pulling spend over Beltane removals, no national regulator escalating on Wheel-of-the-Year demotions, no corporate diversity council whose remit covers Rodnoverie.

Within this configuration, the platform offers two faces. The first face is censorship—the takedowns, demotions, age gates, and shadow bans whose rationales fill the 247-item corpus this study draws on. The second face is compromise—appeal forms, transparency reports, contextual policy carve-outs, trusted-flagger pipelines, "religious content" exemptions, and, most recently, the systemic risk assessments mandated under the EU's Digital Services Act. Both faces, we argue, are moves within the same game. The censorial face removes expression that does not fit; the compromising face re-admits a re-shaped fraction of that expression on conditions that leave the underlying optimization untouched. Compromise within the rules of the dominant game is not, on our reading, pluralism in the strong sense developed in Section 5; it is a governance arrangement whose register obscures the extent to which the algorithm—or, more precisely, the political economy that owns and deploys it—sets the terms of the game.

The paper is organized as follows. Section 2 develops the theoretical framework by combining critical algorithm studies, the data-colonialism program, and a platform extension of Bourdieusian meta-capital. Section 3 describes the materials and methods. Section 4 reports empirical findings. Section 5 proposes the conceptual pair algorithmic sacredness/algorithmic pluralism. Section 6 discusses the implications for post-digital religion under platform capitalism. Section 7 concludes.

2. Theoretical Framework: Critical Algorithm Studies, Data Colonialism, and Platform Meta-Capital

Three converging literatures inform the analysis.

Critical algorithm studies [5–13] have established two results we rely on. First, classifier outputs are never epistemically neutral: training data, labeling conventions, evaluation metrics, and engineering judgment all encode and reproduce the assumptions of the institutions that produce them. Second, the failures of these systems cluster on populations that are sparse in training data, polysemic in representation, or already coded as deviant by the wider culture from which the data are drawn [6,9]. We take both results as background. The narrower argument of this paper does not require their re-establishment.

Data colonialism [1,2], with the political-economic framing brought into focus by recent critical reviews [3,14], supplies the political-economic horizon. The operative thesis is that platform extraction is not a clean break with prior capitalism but a perpetuation of colonial patterns of appropriation in a new register, with symbolic and economic violence taking the place of physical violence. For our purposes, the implication is that platform moderation is not external to the extractive apparatus; it is one of its operating instruments. On the reading we develop, what is moderated can be safely extracted, and what is removed tends to be an expression that cannot be safely extracted within the revenue model.

A complementary discipline comes from Couldry's [15] earlier work on the "myth of big data". Couldry argues that the term *big data* performs ideological work: it implies a different order of knowledge that bypasses humans' "meager attempts at self-understanding through interpreting the local details of what they think, say, and do" [15]. He calls this the myth of big data. He argues for a project of de-reification—refusing the myth as such, interrogating specific claims, spreading the technical literacy that opening the black box requires. The disciplinary lesson is direct: "algorithmic moderation", in its strong form, is at risk of doing to platform classifiers what "big data" does to data processing—reifying a partly mechanical process into a thing that knows, converting the system's outputs into a new order of binding judgment, and suppressing the human interpretive labor that is in fact still constitutive of every step. We retain the term *algorithmic moderation* but refuse the reification: algorithms do not have magisteria; the proprietary infrastructures that deploy them, under the political economy of platform capitalism, do.

Bourdieusian meta-capital, extended to platforms. Bourdieu's late account of the state describes it as wielding meta-capital over the various fields of cultural production: the capacity to set the rate of exchange between forms of capital, the rules of play in specific fields, and the categories through which fields recognize their own internal hierarchies [16]. Couldry [17] argued, by analogy, that media institutions exercise a comparable meta-capital across specific cultural fields: a power exercised not within any one specialized field but across fields, influencing what counts as symbolic capital in each and circulating general categories within which downstream actors in those fields conduct their own internal contests. We extend the same move one step further. Contemporary digital platforms exercise meta-capital across, among other fields, the religious field. They define what counts as recognizable devotional capital within Slavic Native Faith, Heathenry, Wicca, Druidry, Baltic native faiths, and Goddess-centered groups. They condition the rates at which symbolic capital from any of those traditions can circulate to a wider European public. They set the categorical vocabulary—"occult", "sensitive", "extremism-adjacent", "unsafe activity"—within which intra-field disputes are conducted. State meta-capital, as Bourdieu insisted and Couldry [17] developed, is exercised under forms of political accountability, however imperfectly; platform meta-capital is exercised under none of those constraints. The institution that determines what counts as a recognizable Beltane in 2026 is a Delaware-incorporated holding company whose board has no obligation to the European

pagan public, no regulator with jurisdiction over its theology of categories, and no historical analog in any prior political settlement.

The three literatures converge on a single framing claim: under platform capitalism, the meta-capital previously distributed across ecclesiastical, state, and editorial actors has been substantially concentrated in the moderation regimes of a small number of proprietary infrastructures. Whether the resulting redistribution of recognition is felt as "censorship" or "compromise" depends on where a particular tradition stands in the negotiation. The majority of traditions feel compromise. Minority traditions feel censorship. Neither feels meta-capital, because meta-capital, like all naturalized forms of symbolic power, works by becoming invisible.

A second strand of the framework, which the reviewed version of this paper did not develop in its own right, concerns the platform-governance literature that has emerged over the last decade at the intersection of law, communication studies, and science-and-technology studies. Klonick [18] has documented the process by which the major platforms have become, in effect, the *new governors* of online speech, developing internal norms, adjudicative procedures, and appeal architectures that in earlier constitutional settlements were the province of state institutions. Douek [19] has argued that the empirical scale at which contemporary moderation operates—hundreds of millions of decisions per day—has driven a shift from the older, item-by-item adjudicative model toward what she calls a *systems-thinking* approach in which the platform manages probability distributions of decisions rather than the correctness of any individual decision. Roberts [20] has documented the concealed labor regime that sustains this apparatus: a globally distributed and structurally precarious workforce of commercial content moderators whose invisibility is not accidental but is a design feature of the platform's public presentation of moderation as an algorithmic operation. Together, these three references establish two propositions that inform our reading of the empirical material. First, moderation is not primarily a technical operation on which sociological analysis has been overlaid—it is a governance regime whose technical apparatus is itself an artifact of specific labor arrangements, legal environments, and commercial incentives. Second, the opacity of that regime, particularly toward small and stigmatized populations, is not a residue of technical immaturity but a stable structural feature of how platform-scale moderation has been organized. To this stream we add a further reference from the political-science literature on religious minorities: Fox [21] has assembled comparative cross-national evidence documenting that structural discrimination against minority religions is a persistent feature of contemporary states, including Christian-majority European states, and that discrimination clusters in the domains of visibility, recognition, and public expression. Fox's analysis operates at the state level; our contribution shows that the same pattern is reproduced at the platform level, and that where state discrimination has increasingly been contested through law, platform-level discrimination has hitherto lacked either the legal instruments or the transparency conditions on which contestation depends.

A third strand of the framework concerns the field's internal heterogeneity. Contemporary pagan studies, and especially the literature on Slavic Native Faith [22–29], emphasizes the dual character of the movement: on the one hand, a creative, pluralistic, eco-spiritual reconstruction of pre-Christian religiosity; on the other, a discursive landscape inflected by nineteenth-century philological inheritances—Indo-European ideology, Kossinna-style ethnoarchaeology, the Urheimat motif—which a fraction of practitioners refunctionalize as identity claims [28–31]. Any serious analysis of platform moderation of pagan content must reckon with this heterogeneity. Moderation systems trained to detect ethnonationalist iconography will struggle to distinguish a photograph posted by an ecologically inflected clan from visually similar material circulating in ethnonationalist channels.

3. Materials and Methods

The design is hybrid. Empirical material is drawn from three sources.

3.1. Netnography

Since 2018, GM has conducted non-participant and participant observation in Slavic Native Faith Facebook groups (Czech, Slovak, Polish, pan-Slavic); Telegram channels of Czech, Polish, and Russian-language pagan circles; two Discord servers; r/pagan and r/SlavicPaganism on Reddit; one VKontakte community; and several large Czech-language Facebook groups oriented to Norse Heathenry, Viking living history, and Germanic magic. The latter include the three groups dedicated to Vikings and Norsemen with memberships above two thousand—"Vikingové", "Vikingové cz", and "Čeští vikingové" (the last of which originated in May 2016 as the seed of what has since been described as the Czech Vikings phenomenon [4])—and the largest Czech-language group of this orientation, "Runy, pohanství a germánská magie" ("Runes, Paganism, and Germanic Magic"), which at the time of writing counts approximately 21,200 members. This is a non-trivial figure in the Czech online ecosystem, given that the 2021 national census recorded 2764 self-declared pagans, 189 druids, and 42 animists in the population [32].

Field notes record posting volumes, themes (bibliographic recommendations, standing-stone photography, archaeological objects, festival announcements, goddess imagery, living-history images, self-irony), takedown events, shadow-banning complaints, and migration patterns across platforms. Earlier collaborative work on this terrain [4,33] supplied both the entrée to several of these groups and a baseline against which the moderation events of 2023–2026 could be read.

3.2. Informal Interviews

Sampling. Recruitment combined purposive and snowball strategies stratified along three axes: tradition (Slavic Native Faith, Greek and Roman polytheism, Baltic native faiths, Nordic Heathenry and Ásatrú, Wicca, eclectic and Goddess-centered groups), jurisdiction (nine European countries, listed below), and role within the community (leaders and ritual organizers vs. rank-and-file practitioners without an organizational role). Initial contacts came from GM's long-standing insider trajectory in the European pagan milieu (Section 3.5) and from the netnographic sites described in Section 3.1; onward introductions followed the snowball chain within and across jurisdictions. Inclusion criteria were three: (i) adult self-identified pagan practitioner; (ii) active on at least one of the digital platforms whose moderation this study concerns; and (iii) willing to discuss personal experience of takedowns, demotions, appeals, or the felt opacity of platform moderation. Interlocutors who declined the third criterion were excluded from the interview corpus even where they participated in the netnographic observation. No interlocutor was compensated. The final corpus is $N = 97$.

Setting. Conversations with these 97 pagan practitioners across the Czech Republic, Italy, Greece, Slovakia, Lithuania, Latvia, Russia, Denmark, and Sweden took place at ritual gatherings, community festivals, after-circle discussions, in private homes, on long walks, and over messaging applications when in-person contact was not possible. They were unstructured in the sense that a fixed question schedule did not guide them; they were nonetheless purposeful, returning across encounters to a recurring set of concerns about platform visibility, takedown experience, symbolic recognition, and the felt opacity of moderation. The interlocutors included community leaders, ritual organizers, and active content creators on the one hand, and ordinary practitioners with no organizational role on the other.

The interlocutors span different traditions—Slavic Native Faith (Czech Slované, Slovak Rodný kruh, and the Russian Union of Slavic Communities of the Slavic Native Faith), Greek and Roman Gentiles (Greek Supreme Council of Ethnic Hellenes, Italian Pietas), Baltic native faiths (Lithuanian Romuva, Latvian Dievturī), Nordic Heathenry and Ásatrú (Forn Sed circles in Sweden), Wicca, and eclectic and Goddess-centered groups—as well as different platform repertoires and different generational profiles within the European contemporary pagan milieu.

The choice of an informal register has both epistemic and ethical rationales. Epistemically, formal interview protocols would have foregrounded the practitioners' post hoc rationalizations of moderation events at the expense of the more diffuse, more affectively textured material that emerges in conversation—the half-finished anecdotes, the off-hand comparisons, the disputes about whether a particular takedown was really about the symbol or the comment thread, the moments at which a practitioner remembers a takedown only because someone else mentions a similar one. The dynamic the paper analyzes (opacity, asymmetry, prophylactic self-censorship) is the kind of dynamic that does not survive translation into a structured questionnaire; an informal register preserves the diffuseness that is part of the phenomenon. Ethically, in a small, internally surveilled (by adjacent religious groups and, for some sub-populations, state actors) and accustomed to misrepresentation, the imposition of a formal interview frame can foreclose disclosure rather than enable it.

In practice this meant that conversations were rarely recorded; field notes and reconstructive memos written shortly after each encounter, sometimes corrected by the interlocutor in subsequent exchanges, are the primary documentary trace. Where interlocutors explicitly authorized audio recording, recordings exist and were transcribed by GM in the language of the conversation. Where they did not, only the field notes and memos exist—this asymmetry is discussed as a methodological feature, not an incidental gap, in Section 3.5.

Coding. The corpus of notes, memos, and (where available) transcripts was coded thematically in two passes by GM using an open coding scheme that was iteratively refined during fieldwork. The first pass identified emic categories (the terms practitioners themselves used to describe moderation events: *smazaný* [erased], *stopnutý* [demoted], *pod kloboučkem* [shadow-banned]); the second pass mapped these onto etic categories aligned with the takedown-notice coding scheme described in Section 3.3, so that interview material and notice material could be read together. Because the corpus is qualitative and single-coder, formal inter-coder reliability metrics (e.g., Cohen's κ) are not applicable. In their place, three validation procedures were used: (i) *member-checking*, in which reconstructive memos were returned to a purposive subset of interlocutors for correction and comment before coding was finalized; (ii) *triangulation* across the three empirical strands (netnographic observation, interview memos, takedown-notice corpus), so that a claim entering the findings section is supported by material from more than one strand; and (iii) *reflexive validation* through periodic discussion of ambiguous cases with OR, whose analytical distance from the field served as a check on GM's insider familiarity. The specific asymmetry between leaders and rank-and-file practitioners, and the distinction between events directly experienced and events recounted second-hand, was preserved as a coding dimension rather than aggregated away.

The Russian-language material requires a specific methodological note. Since the 2022 designation of Meta as an "extremist organization" in the Russian Federation and the subsequent restrictions on Facebook and Instagram, in-country access to practitioners on those platforms has substantially decreased; in parallel, the documented intensification of platform pressure on minority religious content within Russia [22] has made open identifi-

cation a security concern. Russian-language interlocutors were therefore reached through long-standing contacts established before 2022 within the European Congress of Ethnic Religions network and through diaspora practitioners residing outside the Russian Federation; conversations took place via Telegram, which at the time of fieldwork was already subject to access restrictions inside the Russian Federation but remained reachable for users familiar with the wide ecosystem of VPN tools through which Russian users routinely bypass such blocks. The informal register suited this segment of the work particularly well: it allowed conversations to drift in and out of the topic over weeks, without the documentary footprint of a formal protocol. An additional layer of anonymization (composite biographical details, no platform-specific identifiers) was applied to the resulting notes. The choice of Telegram is itself a small datum for the argument: a population whose religiosity has been pushed off Meta-controlled platforms by both state action and platform-side moderation conducts its conversations on an infrastructure that is itself officially restricted, with usability conditional on the practitioners' own circumvention competence.

A second strand of Russian-language material came from offline contact with a Russian-speaking contemporary pagan community based in the Czech Republic. This circle is markedly more eclectic than the quasi-ethnonationalist and openly ethnonationalist groups GM encountered during earlier visits to Russia: it draws on Slavic and Baltic symbolic repertoires without anchoring them in a primordialist ethnopolitics, accepts members of mixed background, and is more open to gender-balanced ritual leadership. The contrast between the diasporic eclectic milieu in the Czech Republic and the ethnonationalist currents observed in some other Slavic-speaking countries in the past is itself analytically useful: it suggests that what circulates as "Slavic Native Faith" on a given platform is not a single object but a distribution of practices whose moderation as one homogeneous category—routinely "occult" or, more recently, "extremism-adjacent"—erases precisely the internal pluralism that distinguishes a Veles altar in a Prague flat from a Rodnover gathering whose iconography genuinely overlaps with far-right repertoires.

3.3. The Takedown-Notice Corpus

With explicit consent (Section 3.5), GM obtained 247 takedown notices and appeal correspondence from a subset of informants, covering 2019–2025. The corpus was shared through encrypted channels; notices were anonymized before coding by stripping platform-side account identifiers, unique appeal-reference numbers, and any accompanying comments that would identify the originating group. Each notice was coded along six dimensions: (i) *platform of origin* (Meta family—Facebook, Instagram, Threads; YouTube; TikTok; VKontakte; other); (ii) *content type* (text, still image, video, livestream, immersive event); (iii) *stated rationale* invoked by the platform (occult/witchcraft; hate symbol/extremism-adjacent; unsafe activity; sexual or nudity policy; misinformation or manipulated media; other); (iv) *appeal trajectory* (no appeal filed/appeal filed and denied/appeal filed and content restored); (v) *jurisdiction of the account* (as self-reported by the informant); and (vi) *year* (2019–2025). A summary of the corpus by rationale, platform of origin, and appeal outcome is presented in Section 4.1 (Table 1).

Table 1. Takedown-notice corpus by stated rationale, platform of origin, and appeal outcome. (N = 247), 2019–2025.

Stated Rationale	*N* (%)	Meta Family	YouTube	TikTok	VKontakte	Appeal Filed (% of *N*)	Restored on Appeal (% of Appeals Filed)
Sexual/nudity policy	102 (41.3%)	90	7	4	1	2 (2.0%)	0 (0.0%)

Table 1. *Cont.*

Stated Rationale	*N* (%)	Meta Family	YouTube	TikTok	VKontakte	Appeal Filed (% of *N*)	Restored on Appeal (% of Appeals Filed)
Occult/witchcraft/ dangerous content	71 (28.7%)	32	18	0	21	3 (4.2%)	0 (0.0%)
Hate symbol/extremism-adjacent	38 (15.4%)	19	10	1	8	0 (0.0%)	0 (0.0%)
Misinformation/ manipulated media	13 (5.3%)	3	7	0	3	5 (38.5%)	2 (40.0%)
Unsafe activity/sensitive content (ritual events)	12 (4.9%)	6	0	0	6	1 (8.3%)	0 (0.0%)
Other/platform-specific	11 (4.5%)	4	3	1	3	5 (45.5%)	2 (40.0%)
Total	247 (100%)	154 (62.3%)	45 (18.2%)	6 (2.4%)	42 (17.0%)	16 (6.5%)	4 (25.0%)

Note: Rows are sorted by frequency. Column percentages for platform of origin are computed against $N = 247$ in the totals row; row-level platform counts sum by row to the corresponding *N*. *Appeal filed* is a subset of *N* (percentage of that row's notices for which an appeal was submitted). *Restored on appeal* is a subset of *Appeal filed* (percentage of that row's appeals in which content was restored). "Meta family" aggregates Facebook, Instagram, and Threads (see the paragraph below the table for the rationale of the aggregation).

3.4. Genealogy of the Inquiry

The questions pursued here did not arise in the abstract. They are the conceptual continuation of three earlier strands of fieldwork.

First, in 2018, GM conducted an online ethnographic study of nude culture and contemporary Czech paganism, prepared for the SIEF 2019 Congress in Santiago de Compostela [34]. The fieldwork was conducted inside closed Facebook groups accessible only by recommendation, in which practitioners discussed and documented the place of bodily nudity within ritual and communal practice. This discussion has been part of contemporary paganism since its modern reconstructions, and, in relevant practitioner contexts, is neither erotic nor pornographic. A few months into the study, Facebook implemented a wave of platform-wide content actions that removed both the images and, more consequentially, the written comments through which practitioners had been articulating their reasoning. The corpus on which the research had been built vanished. Until that point, the major digital platforms had been assumed to be places of effectively permanent preservation—that once a post had passed the filter of moderation, it would remain accessible. The 2018 episode disabused that assumption and made the infrastructural contingency of contemporary pagan online life a researchable problem in its own right.

Second, ethnographic work on the decline of the first generation of Czech Slavic-pagan organizations [35] and on transgender experience within the Czech pagan community [33] had already exposed how digital media reshape the boundary between in-group recognition and external visibility, and how minority configurations within an already minority religion become particularly difficult to render publicly without simultaneously rendering them as targets.

Third, the collaborative study of the Czech Vikings [4] tracked an explicitly digital-native community whose existence as a recognizable formation is dated to the launch of the Facebook page Čeští vikingové in May 2016, and whose practices unfold across Facebook, Instagram, TikTok, and YouTube, in tandem with offline reenactment, LARP, and living-history activity. That study showed that the same iconographic repertoire—runes, hammers, wolf imagery—is mobilized by communities ranging from explicitly eclectic and self-ironic Czech Viking groups to ethnonationalist Heathen currents elsewhere in Europe; the moderation regimes of mainstream platforms, however, encounter only the surface of that repertoire.

3.5. Reflexivity, Ethical Framework, and Validation

Positionality. GM has been a participant in the European pagan milieu since 1998: co-founder of the Czech group Radhošť and the registered association Rodná víra (officially registered 13 April 2001); co-founder of the World Congress of Ethnic Religions; current board member of the European Congress of Ethnic Religions; and the author of autoethnographic and historical-comparative work on Slavic paganism [28,36]. This trajectory grants emic access—to private chats, to ritual circles, to disputes that practitioners would not narrate to a stranger—that purely etic methods could not reach [37]. It also requires methodological vigilance against the affective predispositions that long familiarity produces. Throughout, GM treats personal historical participation as data and as bias, and triangulates insider observations with the testimony of practitioners with whom there is no prior relationship. OR has no comparable insider relationship to the field; his contribution is to the conceptual framework, the analytical reading, and the two concepts proposed in Section 5, and his analytical distance functions as one of the validation checks discussed above.

Ethical framework. The study was conducted in accordance with the ethical guidelines of the American Anthropological Association and the European Association of Social Anthropologists, and in accordance with Regulation (EU) 2016/679 (GDPR) and Czech Act No. 110/2019 Coll. on the Processing of Personal Data.

Informed consent. Verbal informed consent was obtained from every interlocutor at the outset of every conversation, in the interlocutor's preferred language (GM conducted each conversation in the interlocutor's mother tongue). The full script of the verbal informed consent is filed with the editorial office as a supplementary document. It covered, in each case: (i) GM's identity as a researcher and his long-standing affiliation with the European pagan milieu; (ii) the purpose of the conversation and of the broader research project; (iii) the voluntary nature of participation and the right to withdraw or redirect the conversation at any time; (iv) the anonymization regime described below; (v) the documentary form the conversation would take (notes and reconstructive memos as default; audio recording only where the interlocutor explicitly authorized it); and (vi) the offer of member-checking. Where sensitive material was subsequently broached (Russian-language conversations under post-2022 legal risk; conversations about takedowns of extremism-adjacent iconography), consent was reconfirmed orally before the specific topic was addressed. No signed consent forms were used: in a small, internally surveilled, and frequently misrepresented population, the imposition of a signed documentary trace would have foreclosed disclosure and would, for the Russian-language segment, have constituted a security risk for interlocutors.

Anonymization and data-security safeguards. All interlocutors are anonymized in the published record. First names are replaced with tradition-appropriate pseudonyms where a name is required for narrative continuity; specific groups smaller than a national community are referred to only by tradition and jurisdiction (e.g., "a Czech Rodnover circle") rather than by name; platform-specific handles, group URLs, event dates, and image content that would permit re-identification are not reproduced. Where the combination of tradition, jurisdiction, and biographical detail would be individually identifying in a small community, composite biographical details are used and are flagged as such in the field notes. Field notes, reconstructive memos, and (where they exist) recordings are stored on an encrypted local drive under GM's exclusive control; the takedown-notice corpus is stored separately, also encrypted, with account-identifying metadata stripped before coding. The Russian-language material carries additional safeguards, as detailed in Section 3.2: no in-country recording, no platform-specific identifiers in the notes, and composite biographical details throughout.

Limits of access, and the researcher's structural position. A second axis of reflexivity concerns the limits of access. Critical algorithm studies have long noted that the proprietary nature of moderation pipelines blocks the standard methods of social science [7,10,12,20]. We cannot inspect the classifiers; we can only observe their effects, the rationales they generate, and the testimony of those who build or police them. This asymmetry is not an inconvenience to be remedied through better access requests. It is one of the principal mechanisms by which platform meta-capital is maintained, and we treat it accordingly.

4. Findings

The corpus and the field material are read in four parts. Throughout, we distinguish between empirical findings drawn directly from the corpus, theoretical interpretations offered in dialogue with the critical-algorithm-studies and data-colonialism literatures, and stronger *critical hypotheses* about platform cost structures that go beyond what the outside-only methodology of this study can directly demonstrate. The reader is invited to weigh each register accordingly.

4.1. The Three Faces of Misclassification

Corpus overview. Table 1 summarises the 247-notice corpus by stated rationale, platform of origin, and appeal outcome. On the rationale axis, the largest single class is *sexual/nudity policy*, which accounts for 102 of 247 notices (41.3%); the empirical prominence of this class is not an accident of the sample but is continuous with the earlier 2018 SIEF ethnography of naturist paganism (Section 3.4), and it indicates that the moderation category most frequently applied to contemporary pagan content in this corpus is one that classifies devotional and ritually motivated bodily exposure through a policy framework designed for adult content. Three further analytically salient classes account for the bulk of the remainder: content classified as *occult/witchcraft/dangerous content* (71 notices, 28.7%), as *hate symbol/extremism-adjacent* (38 notices, 15.4%), and as *unsafe activity/sensitive content* covering ritual events (12 notices, 4.9%). A residual is distributed across *misinformation/manipulated media* (13 notices, 5.3%) and *other/platform-specific* (11 notices, 4.5%).

On the platform axis, the corpus is dominated by Meta-family notices—Facebook, Instagram, and Threads together—accounting for 154 of 247 notices (62.3%); YouTube contributes 45 notices (18.2%), VKontakte 42 (17.0%), and TikTok 6 (2.4%). No notices originated from Reddit, Discord, or Telegram public channels in the final corpus.

A final aggregate observation, which we treat as itself a finding, concerns the appeal trajectory. Appeals were filed in only 16 of 247 cases (6.5%); of these, 4 (25.0% of filed appeals, 1.6% of the full corpus) resulted in content restoration. The very low appeal-filing rate is one of the most striking features of the corpus. We return to it in Section 4.2, where we read it as a downstream effect of the opacity regime rather than as evidence of practitioners' agreement with the moderation outcome.

Analytical scope of Section 4.1. The three subsections that follow—Section 4.1.1 on content classified as *occult*, Section 4.1.2 on content classified as *hate symbol/extremism-adjacent*, and Section 4.1.3 on live ritual events classified as *unsafe activity*—deliberately do not exhaust the six rationale-classes reported in Table 1. The selection tracks an analytical distinction, not a distributional one. The three classes discussed in depth share a common structural property: they are cases of *misclassification driven by symbolic polysemy*, in which the classifier encounters a symbol, an iconographic complex, or a ritual context that does not sit cleanly within its category system and assigns it to a category belonging to an entirely different symbolic universe (a Perun-rune emblem to "occultism," a Kolovrat to "hate symbols," a fire-jumping rite to "dangerous activity"). The largest class in the corpus, *sexual/nudity policy* (102 notices, 41.3%), is analytically distinct: the classifier does not

misidentify the content, it correctly detects bodily exposure and then applies a nudity policy that makes no accommodation for devotional or ritually motivated bodily practice. This is a *policy-fit* problem rather than a *category-conflation* problem, and it has already been narrated as such in Section 3.4 as the structural pattern first identified in the 2018 SIEF ethnography of naturist Czech paganism [34] and now, on the aggregate evidence of Table 1, confirmed as the dominant single class in the 2019–2025 corpus. In the political-economic reading developed in Sections 4.3, 4.4 and 5, the two kinds of problems converge on the same underlying arrangement; at the level of empirical findings, they warrant separate treatment.

A note on the platform-of-origin categories is also warranted here. "Meta family" aggregates Facebook, Instagram, and Threads because these three services share a common Community Standards regime, a common moderation pipeline, and a common appeals architecture under Meta Platforms, Inc.; a notice issued on Instagram is, for the purposes of the analytical dimensions on which this study codes, an instance of the same moderation apparatus that issues notices on Facebook and Threads. Where the analysis in what follows turns on features that differ across the three services—for example, the specific interaction of an image-recognition model with the Instagram feed rather than the Facebook one—the individual service is named in the text.

With these distributions on the record, three rationale-classes recur across the corpus in ways that warrant closer analytical attention. In what follows we distinguish, for each class, between (a) the empirical finding (what the corpus and field material contain), (b) the theoretical interpretation we develop with the critical-algorithm-studies and data-colonialism literature, and (c) the stronger *critical hypothesis* we advance about the platform's cost structure. The three are not equivalent in epistemic weight, and the reader is invited to hold them apart.

4.1.1. Devotional Content Classified as "Occult"

Seventy-one of 247 notices (28.7%) invoked an undefined notion of occultism—"occult content", "witchcraft", "dangerous content", or platform-specific equivalents—to take down photographs of altars, ritual objects, and emblem-style images. The asymmetry with majority devotional imagery is structural: a livestreamed Catholic Mass with comparable visual intensity (incense, candles, ritual blade) does not trigger comparable flags.

Theoretical interpretation. We read this asymmetry as consistent with a training distribution in which the Catholic Mass functions as the unmarked case of "religious content" and pagan altars as the marked case, in line with what critical algorithm studies have documented about the cultural inheritance of majority/minority categories in training corpora assembled from Anglo-American materials [6,9]. This interpretation is well supported by literature and by the aggregate pattern in Table 1.

Critical hypothesis. We further hypothesize—more cautiously, and without direct access to the platforms' internal reasoning—that the asymmetry is sustained by a cost structure under which the platform incurs no advertiser-side cost for reproducing the marked/unmarked distinction: advertisers do not pull spend over pagan removals, national regulators do not escalate on them, and corporate diversity councils, where they exist, have not to our knowledge been chartered to consider the case. This is a *hypothesis about incentives*, not a proof of intent, and we present it as such. It fits the pattern; it is not deduced from the pattern. Complementary evidence from platform-side sources—leaked internal documents, litigation discovery, whistleblower testimony—would be needed to strengthen it, and outside such evidence we retain the more cautious formulation.

4.1.2. Devotional Content Classified as "Extremism-Adjacent"

Thirty-eight cases (15.4%) cited "hate symbols" or "extremist content" to take down content the informants characterized as devotional. The symbolic vocabulary involved—Othala, Sonnenrad, Kolovrat, valknut, Mölnir, reconstructions of Perun's axe—is genuinely polysemic, and the corpus contains some content whose iconographic neighborhood overlaps with extremist appropriations. *Theoretical interpretation.* The reformist reading of this category is that the platform is doing its best at a hard signal-extraction problem—the polysemic vocabulary is genuinely polysemic—and that a contextual classifier would in principle solve it. We accept the first half of that reading (the problem is genuinely hard) and disagree with the second (the solution is not, on the pattern the corpus displays, converging).

Critical hypothesis. Our more cautious formulation is that the platform is under regulatory pressure—notably the EU's Digital Services Act, whose Article 34 obligation on Very Large Online Platforms to assess systemic risks arising from the design and functioning of their services entered into force in 2024—to act on extremism-adjacent iconography, and that in the absence of countervailing pressure, minority-pagan content is among the least costly places to satisfy that pressure: over-removing a Kolovrat, on our reading, imposes no visible cost on the platform, because no advertiser, no national regulator, and no organized political constituency is currently in a position to object. The literature on religious discrimination and automated moderation [see below] documents comparable patterns of asymmetric over-classification affecting other minority religious populations, though the specific evidentiary link to internal cost structures remains, in every such case, an inference rather than a demonstration. We advance the incentives hypothesis in that same register: as a plausible reading of a pattern, not a claim to have opened the black box that would confirm it.

The Czech Vikings sub-corpus [4] is the clearest demonstration. Practitioners who go to considerable rhetorical lengths to distinguish their explicitly self-ironic, eclectic-spiritual practice from ethnonationalist Heathenry—practitioners whose own neologism "Švejking" signals the distance from the toxic-masculinity "peak male physique" Viking that Thykier Makeeff [38] has analyzed in Scandinavia, and whose Czech self-representations are re-described as "Lardcakes"—nonetheless see their content treated as belonging to the same symbolic class as content originating in genuinely ethnonationalist Heathen channels elsewhere in Europe. The classifier flattens what the community itself goes to considerable rhetorical lengths to distinguish—a flattening that, on the cost-structure hypothesis advanced above, the platform has no countervailing incentive to correct.

4.1.3. Live Ritual Events Classified as "Unsafe Activity"

Live-streamed or recorded ritual gatherings—outdoor sabbats, full-moon esbats, group invocations—were repeatedly demoted, age-gated, or removed under "unsafe activity" (for fire rites), "sensitive content", or, when multilingual chant audio failed automated language-policy checks, flagged for review and disrupted mid-broadcast. Informants who organize such events describe a chronic uncertainty: the rite is licit, the platform is permissive in policy, but the classifier intervenes anyway.

The downstream behavior is prophylactic self-censorship and migration to smaller, less searchable platforms. From the practitioner's perspective, the behavior is rational. From the platform's perspective, it is—on our critical hypothesis, and consistent with the aggregate pattern in Table 1—low-cost: the displaced audience is too small to register in growth metrics and, in any case, is replaced by content from larger, more advertiser-friendly verticals. On this reading, the moderation regime produces a quiet attrition that no quarterly report will mark; the interpretive weight the reader places on the reading is a function of how much of the surrounding literature on platform economics one accepts.

4.2. Opacity as Governance

Across the 247 takedowns, only 19 (7.7%) included an explanation specific enough to permit a meaningful appeal. The remainder offered category labels—"violates community standards" and "sensitive content"—without reference to the specific feature that triggered the action. This is consistent with the broader literature on platform explanation generation [10–12], which has documented that explanation is itself a downstream production whose outputs need not be faithful to the upstream classifier's reasoning.

The downstream consequence of this opacity is visible in the appeal trajectory reported in Table 1. Only 16 of 247 notices (6.5%) elicited a formal appeal from the affected practitioner. This rate is strikingly low, and it is not, we argue, evidence that practitioners accept the moderation outcomes. The interview material consistently records the opposite: sustained perceptions of unfairness, misclassification, and religious discrimination, together with a settled expectation that the appeal system will neither explain nor reverse the decision. Two features of the distribution in Table 1 support this reading. First, the 6.5% overall filing rate is close to the 7.7% of notices that included a specific rationale (§4.2 above): appeals cluster in cases where the practitioner has something concrete to contest, and the rest of the corpus—the great majority—offers no such grip. Second, the filing rate is markedly higher in the two rationale-classes where the platform's message is closest to a specific claim—*misinformation/manipulated media* (38.5%) and *other/platform-specific* labels (45.5%)—and drops to zero for *hate symbol/extremism-adjacent*, where the potential legal exposure of contesting an extremism designation almost certainly deters filing. Of the 16 appeals filed, 4 (25.0% of appeals filed; 1.6% of the full corpus) resulted in restoration; the two rationale-classes carrying the analytical weight of Sections 4.1.1 and 4.1.2—*occult/witchcraft* and *hate symbol/extremism-adjacent*—recorded zero restorations across the entire corpus. Opacity, in this respect, is not a passive absence of explanation; it operates as a lever of governance whose effect is to narrow the very appeal channel through which the platform's public commitments would be tested.

We read opacity not as an engineering deficiency but as a governance instrument. Specific explanations would expose the classifier's policy gradients to adversarial probing by both bad-faith actors and good-faith critics. Non-specific explanations exhaust the appellant's resources, externalize the cognitive cost of the moderation event onto the user, and preserve the platform's optionality to update the underlying policy without notice. On our reading, this is less a cost to the platform than a feature of the arrangement. The literature on opacity in algorithmic governance [7] treats it as constitutive of the form, and our corpus is consistent with that treatment.

A theological consequence follows, which we name without endorsing. Practitioners, confronted with repeated unexplained takedowns, interpret the pattern through frames already available to them—residual Christian-cultural prejudice, state-aligned suppression of minority faith, in some corners, the persecutory narratives that ethnonationalist factions use for recruitment. Opacity, in this respect, tends to consolidate the very interpretive frames that subsequently legitimate further moderation. The scholarship on religious discrimination and automated moderation [see the DSA and platform-governance literature engaged below] has begun to document analogous dynamics in other minority-religious populations and treats them as a systemic property of opacity-under-scale rather than as a set of one-off failures.

4.3. The Training-Data Problem as a Political-Economic Problem

We cannot inspect the training data of the major platforms. The published technical and policy literature [6,9,12] and the platforms' own Transparency Reports (2020–2025) make it clear, however, that classifiers in this domain are trained predominantly on Anglo-

American extremism corpora and on majority-religion devotional imagery—both because those corpora exist and because building corpora that adequately represent minority European devotional traditions would cost money that no platform has a commercial reason to spend. The training-data problem is therefore not a research problem awaiting a methodological fix. It is a political-economic problem. The corpora exist because someone paid for them. Nobody has paid for Veles altars. Nobody is going to.

This re-description matters. Calls for "training-data audits" in the reformist register implicitly assume that an audit followed by an investment in better data would close the gap. The political-economic reading observes that, absent a structural change in who pays for the audit and for the resulting corpus, the audit is a discovery procedure that surfaces a gap that the platform's optimization has no commercial reason to close. What sustains the reformist reflex, in this and adjacent cases, is what Morozov [39] has diagnosed as *technological solutionism*: the disposition to redescribe complex, historically sedimented social phenomena as neatly definable problems awaiting a computational fix, and to read the fix's failures as calls for further computational refinement rather than as symptoms of a mis-specified question. Applied to moderation, solutionism converts the political-economic question of who is entitled to co-author the terms of religious recognition into the engineering question of how to build a better classifier—a translation that keeps the underlying property arrangement invisible.

4.4. Infrastructural Contingency and Dispossession

A recurring concern across the interview corpus—expressed by leaders and ordinary practitioners alike—concerns the contingency of the infrastructures on which pagan communal life now depends. The 2024 closure of Microsoft's AltspaceVR was repeatedly cited as a cautionary case, even by practitioners who had not used it themselves. Group archives, ritual schedules, and member networks accumulated over the years can disappear with a single corporate decision taken elsewhere. The 2018 platform-wide action that erased the corpus of online ethnography of Czech naturist paganism [34] is part of the same pattern, observed at the level of a single researcher's archive.

The reformist reading of this calls for "infrastructural commitments—preservation, portability, succession". The political-economic reading observes that infrastructural commitments cost money and are made only when the cost of not making them exceeds the cost of making them. For pagan communities, that calculus is settled in advance. We are at this point describing not a moderation problem but a property problem: who owns the substrate of communal life, and on what terms can the community dwell there. The colonial-frame analogy [1–3,14] is, here, not stretched. The substrate has been appropriated; the dwellers can be evicted; the eviction is invisible until it happens.

5. Algorithmic Sacredness and Algorithmic Pluralism

5.1. Algorithmic Sacredness

By algorithmic sacredness, we name the now-routine transfer of the gate-keeping function in the religious public sphere to non-human computational systems. The transfer is real: classifiers materially decide which pagan symbols circulate, which rituals can be hosted, and which leaders can build audiences. These are precisely the functions that, in earlier media regimes, were performed by ecclesiastical authorities, state recognition procedures, and editorial decisions in religious publishing [40].

The term must be disciplined. The transfer is not a transfer to algorithms as new sacred objects in their own right. Algorithms do not want anything. That the transfer is nonetheless routinely misread as a transfer *to the algorithm* is itself a datum, and one for which Sánchez Capdequí, Gil-Gimeno, and Echeverría-Esparza's [41] account of *techno-functional neo-*

immanentism supplies the cultural-theoretical background. They describe a hegemonic contemporary configuration in which technology, while remaining formally within the this-worldly register of the technical and the efficient, absorbs functions—orientation, promise, adjudication of legitimacy—that in earlier arrangements were carried by explicitly transcendent frameworks [41]. Algorithmic sacredness, on our reading, is what happens when that neo-immanentist hegemony encounters populations whose practice is explicitly religious: the cultural condition under which the platform's adjudicative outputs *feel* binding in a register that was, until recently, the register of religious and public authority, even as the platform disclaims any such role. The transfer is to the political economy that owns the algorithms—to the proprietary infrastructures within which classifiers, policy teams, regulatory liaisons, and appeals architectures together produce binding decisions. To imagine the algorithm as sacred is to make the same naturalizing move that Couldry [15] describes for "big data": treating the system's outputs as a new order of knowledge that bypasses the interpretive practices of the humans actually involved. We name this risk and refuse it. What is sacralized, in the strong sense in which symbolic power misrecognizes its own arbitrariness [16,17], is not the algorithm but the unquestionability of the political economy that runs it. Practitioners do not appeal to the platform's right to make these decisions. They appeal individual decisions within the rules the platform sets, on infrastructures the platform owns, with vocabulary the platform supplies. The platform's prerogative is the sacred object. The algorithm is its priest.

This re-description has two further advantages. First, it grounds algorithmic sacredness in critical political economy and prevents it from sliding into the theological inflation that the digital-religion vocabulary, on its own, would invite. Second, it makes the concept legible across registers: as a religion-specific name for what is more generally analyzable as platform meta-capital under data colonialism [1,2,17].

5.2. Algorithmic Pluralism

Algorithmic pluralism is the normative counterpart of the diagnosis. In the reformist register in which platforms and regulators are happy to discuss the problem, algorithmic pluralism appears as a slate of design proposals: training-data audits that ensure minority religious traditions are present and correctly labeled; explanation interfaces calibrated to permit informed appeal, including by users without legal counsel; the development of contextual classifiers able to distinguish devotional from political uses of polysemic symbols; consultation mechanisms that bring minority-religious expertise into policy formation, on a model adapted from disability and language-rights advocacy; and infrastructural commitments—preservation, portability, succession—that recognize the dependence of religious communities on platforms.

Each proposal is a compromise within the rules of the dominant game, and each is sustained by the solutionist reflex diagnosed by Morozov [39]: the confidence that whatever the problem, the appropriate response lies within the same technical register that produced it. Training-data audits, conducted by the platform or by an auditor paid by the platform, produce findings whose remediation depends on the platform's commercial willingness to invest in corpora it has no advertiser-side incentive to invest in. Explanation interfaces designed by the platform raise the cost of explanation for the appellant relative to the platform's appellate threshold. Contextual classifiers relocate the same hard signal-extraction problem to a slightly more complex model without changing the cost structure that produced the bias. Consultation mechanisms admit a small number of minority-religious experts to a process whose outputs the platform retains the right to ignore. Infrastructural commitments are voluntary undertakings on the part of corporations whose

prior voluntary undertakings—AltspaceVR, every shuttered messaging product, every deprecated archive—have not been honored at the moment of corporate restructuring.

Each compromise is a move within the algorithm's game. Each leaves the platform's political economy untouched. Each can be made and unmade by the same corporate body that made and unmade the moderation regime that produced the problem. To take any of these compromises as a solution is to mistake the conditions of one's exposure for its resolution.

Couldry and Mejias's three-register resistance framework [2,3] supplies the vocabulary for the move we want to make next. Algorithmic pluralism in the reformist register is the within-the-system register. It is a real option, and a partial one. The stronger reading of algorithmic pluralism—which we propose here—requires the beyond-the-system register: federated and community-owned infrastructures (Mastodon-family protocols, Matrix-family communications, IPFS-style preservation), funded outside the data-colonial logic, governed by communities whose religious life they host, with portability across providers as a structural rather than promised feature. This is not a complete proposal; it is a programmatic direction. It is, however, the only register in which "algorithmic pluralism" names anything other than a compromise inside the dominant game.

5.3. The Pair and the Post-Digital Condition

Held together, algorithmic sacredness and algorithmic pluralism form a diagnostic-prescriptive pair. Algorithmic sacredness names the political-economic concentration of meta-capital in proprietary moderation infrastructures, and the misrecognition of that concentration as a property of the technology. Algorithmic pluralism, in its stronger reading, names a programmatic direction toward infrastructures whose governance is not capital's. Both presuppose that algorithms are not external instruments operating on a stable religious world but co-constituents of that world, embedded in a political economy whose operations make their decisions binding. Both treat the post-digital condition not as a rupture to be either celebrated or lamented, but as a terrain on which the sociology of religion must be rebuilt with new conceptual tools.

For contemporary pagans specifically, the pair clarifies the stakes of a movement that has, almost from its modern inception, lived online. The same infrastructures that allowed the Ethnic religions to coalesce across borders in the late 1990s—the mailing lists, the early forums, and eventually Facebook and VKontakte—now condition its visibility through classifiers it cannot inspect, deployed by corporations whose interests it cannot affect. The movement's historical pluralism, its principle of unity in diversity adopted at the First World Congress of Ethnic Religions in 1998, faces a counterpart challenge in the algorithmic register: whether the political economy of platform capitalism can be made to host, rather than to flatten, that pluralism—or whether the answer must be sought elsewhere.

6. Discussion: Post-Digital Religion Under Platform Capitalism

Three implications follow from the analysis.

First, the methodological asymmetry between the researcher and the infrastructure is reproduced in the practitioner's position. The practitioner cannot inspect the classifier any more than the researcher can; the practitioner cannot escape the infrastructure, because the infrastructure is where the community is; the practitioner cannot extract policy concessions, because the population is below the threshold of commercial visibility. The post-digital condition, for minority traditions, is the experience of being subject to a power one cannot see, cannot leave, and cannot negotiate with. It is the digital-religion analog of the more general condition that Couldry and Mejias [1,2] name as data colonialism.

Second, the privileged status of the pagan case is not a function of pagan-specific exceptionalism. It is a function of the case's structural position relative to the platform's optimization. Other minority configurations—minority-language speakers in jurisdictions with majority-language moderation regimes, indigenous communities whose iconography is treated as costume by classifiers trained on majority cultural materials, queer communities whose explicit content is removed under policies that do not similarly police majority sexual content—share the structural position. The same analysis travels. The pagan case is instructive because the asymmetries are unusually visible, not because they are unusually severe.

Third, the available response registers are not equivalent. Within-the-system reforms (DSA systemic risk assessments, religious-content carve-outs, contextual classifiers, transparency reporting) are real but conditional, and the conditions are set by the actors whose game these reforms are part of. Against-the-system tactics (boycotts, regulatory complaints, public-interest litigation, mass migration off a given platform) raise the cost of business-as-usual at the margin but, by themselves, do not alter the political economy that produces it. Beyond-the-system construction (federated, community-owned, religiously competent infrastructures, funded outside data-colonial extraction) is the only register that addresses the underlying property problem. Still, it is also the register with the highest coordination cost and the most precarious resource base. We do not pretend to know how to thread these three registers. We do claim that algorithmic pluralism is not a within-the-system proposal, and that any reading of the term that takes it as one has already conceded the game.

The argument also opens questions that this paper cannot fully resolve. Among them: how generative AI systems, increasingly used to produce devotional content (AI-generated icons, AI-generated prayers, AI-mediated divinatory practices), interact with the moderation regime, and whether their integration intensifies or partially short-circuits the data-colonial dynamic; whether quantum-informed methodological frameworks [42,43] might offer a way of theorizing the entangled, indeterminate co-production of sacredness across human and non-human actors; and how the comparative perspective—Catholic-dominant Italy, strongly secularised Czech Republic, post-Soviet milieux—refines or qualifies the general claim.

7. Conclusions

The encounter between contemporary European pagans and platform moderation is not, on closer inspection, a story about classifier errors or about minority-religion sensitivity training. It is a story about the political economy in which a small number of proprietary infrastructures have, in their moderation regimes, concentrated the meta-capital previously distributed across ecclesiastical, state, and editorial actors. Within that political economy, censorship and compromise are the two faces of the same operation. The censorial face removes expressions that do not fit the optimization; the compromising face re-admits a reshaped fraction of them on conditions that leave the optimization untouched. The algorithm is the dominant player not because it is intelligent or sacred, but because, on our reading, the political economy that owns it has been positioned to set the rules of the game.

We have proposed algorithmic sacredness and algorithmic pluralism as analytic instruments for engaging this condition. The first names the concentration of meta-capital and the misrecognition of that concentration as a property of the technology. The second names a programmatic direction toward infrastructures whose governance is not capital's. Contemporary pagans are a small population; the diagnostic value of their experience is large.

The wider stake is not whether a Kolovrat passes the next classifier update. It is whether the social sciences can describe—and the publics that read them act on the description—the fact that, under platform capitalism, the algorithm is not a tool for ac-

commodating religious diversity. We have argued that it operates as an instrument of an extractive order whose tendency, when religious diversity costs the system nothing, is to suppress it. The compromises offered within that order are real and partial. They are not a solution to the problem. They are the form the problem takes in the register of platform governance as currently practiced.

Author Contributions: Conceptualization, G.M. and O.R.; Methodology, G.M. and O.R.; Formal analysis, G.M. and O.R.; Investigation, G.M.; Writing—original draft, G.M. and O.R.; Writing—review & editing, G.M.; Supervision, G.M. and O.R. All authors have read and agreed to the published version of the manuscript.

Funding: The result was created in solving the student project n. 7427/2025/01 "Online and Offline Activities of Local communities in the Czech Republic" using objective oriented support for specific university research of the University of Finance and Administration.

Institutional Review Board Statement: This study has been granted an ethics exemption for the following reasons: 1. This study is a non-interventional, qualitative social science research project that collects data through online ethnographic observation, informal interviews, and platform content analysis. It does not involve any medical, clinical, biological, or psychological interventions, nor does it involve minors or individuals lacking legal capacity. 2. Under Czech law, the obligation to obtain prior approval from an ethics committee applies only to biomedical and clinical research (as stipulated in Act No. 372/2011 Coll. and Act No. 378/2007 Coll.); there are no equivalent legal requirements for non-interventional social science research; Furthermore, the authors' home institution, the University of Finance and Administration in Prague, does not require prior ethical approval for this type of research, and its Research and Development Department has reviewed and confirmed that the study complies with institutional guidelines regarding informed consent, anonymization, and data protection. Based on the aforementioned legal provisions, institutional policies, and methodological grounds, this study is exempt from prior ethical review by an ethics committee.

Informed Consent Statement: Informed consent was obtained from all subjects involved in the study.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available because they consist of anonymized takedown notices and informal-interview material, the disclosure of which could compromise participants' confidentiality and safety.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Couldry, N.; Mejias, U.A. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*; Stanford University Press: Stanford, CA, USA, 2019.
2. Mejias, U.A.; Couldry, N. *Data Grab: The New Colonialism of Big Tech and How to Fight Back*; The University of Chicago Press: Chicago, IL, USA, 2024.
3. Wintola, T.O. Review of Data Grab: The New Colonialism of Big Tech and How to Fight Back, By Ulises A. Mejias and Nick Couldry. *Insight Turk.* **2025**, *27*, 429–431. [CrossRef]
4. Maiello, G.; Cirklová, J. Northilization in Central Europe: Examining Czech Vikings. *Preprint* **2024**. [CrossRef]
5. Beer, D. *The Data Gaze: Capitalism, Power and Perception*; Sage: London, UK, 2019.
6. Noble, S.U. *Algorithms of Oppression: How Search Engines Reinforce Racism*; NYU Press: New York, NY, USA, 2018.
7. Pasquale, F. *New Laws of Robotics: Defending Human Expertise in the Age of AI*; Harvard University Press: Cambridge, MA, USA, 2020.
8. Bonini, T.; Treré, E. *Algorithms of Resistance: The Everyday Fight Against Platform Power*; MIT Press: Cambridge, MA, USA, 2024.
9. Crawford, K. *Atlas of AI*; Yale University Press: New Haven, CT, USA, 2021.
10. Gillespie, T. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*; Yale University Press: New Haven, CT, USA, 2018.
11. Suzor, N.P. *Lawless: The Secret Rules That Govern Our Digital Lives*; Cambridge University Press: Cambridge, UK, 2019.
12. Gorwa, R.; Binns, R.; Katzenbach, C. Algorithmic content moderation: Technical and political challenges in the automation of platform governance. *Big Data Soc.* **2020**, *7*, 1–15. [CrossRef]

13. Van Dijck, J.; Poell, T.; de Waal, M. *The Platform Society*; Oxford University Press: Oxford, UK, 2018.
14. Ticona, J. Review of The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism, by N. Couldry and U. A. Mejias. *Contemp. Sociol. A J. Rev.* **2021**, *50*, 196–198. [CrossRef]
15. Couldry, N. The myth of big data. In *The Datafied Society: Studying Culture Through Data*; Schäfer, M.T., van Es, K., Eds.; Amsterdam University Press: Amsterdam, The Netherlands, 2017; pp. 235–239.
16. Bourdieu, P. Rethinking the state: Genesis and structure of the bureaucratic field. *Sociol. Theory* **1994**, *12*, 1–18. [CrossRef] [PubMed]
17. Couldry, N. Media meta-capital: Extending the range of Bourdieu's field theory. *Theory Soc.* **2003**, *32*, 653–677. [CrossRef]
18. Klonick, K. The new governors: The people, rules, and processes governing online speech. *Harv. Law Rev.* **2018**, *131*, 1598–1670.
19. Douek, E. Content moderation as systems thinking. *Harv. Law Rev.* **2022**, *136*, 526–607.
20. Roberts, S.T. *Behind the Screen: Content Moderation in the Shadows of Social Media*; Yale University Press: New Haven, CT, USA, 2019.
21. Fox, J. *Thou Shalt Have No Other Gods Before Me: Why Governments Discriminate Against Religious Minorities*; Cambridge University Press: Cambridge, UK, 2020.
22. Aitamurto, K. *Paganism, Traditionalism, Nationalism: Narratives of Russian Rodnoverie*; Routledge: London, UK, 2016.
23. Aitamurto, K.; Simpson, S. *Slavic Native Faith*; Cambridge University Press: Cambridge, UK, 2026.
24. Mačuda, J. Ideologie Rodnověří v Názorech a Praxi Jeho Čelních Představitelů. Ph.D. Thesis, Masaryk University, Brno, Czech Republic, 2012.
25. Shnirelman, V. Ancestral wisdom and ethnic nationalism: A view from Eastern Europe. *Pomegranate* **2007**, *9*, 41–61.
26. Simpson, S. Only Slavic gods: Nativeness in Polish Rodzimowierstwo. In *Cosmopolitanism, Nationalism, and Modern Paganism*; Rountree, K., Ed.; Palgrave Macmillan: New York, NY, USA, 2017.
27. Stasulane, A. A reconstructed indigenous religious tradition in Latvia. *Religions* **2019**, *10*, 195. [CrossRef]
28. Maiello, G. Dea vobiscum sit: Considerazioni su paganesimo contemporaneo, ideologia, scienza e ricerca dell'Urheimat protoindoeuropea. In *Deus Vobiscum? Esperienze Religiose del Postmoderno*; Santi, C., Salvati, A., Eds.; Mimesis Edizioni: Milan, Italy, 2025; pp. 203–219.
29. Galtsin, D. Climbing Europe: Celticity in Russian pagan and nativist movement (1990s–2010s). *Pomegranate* **2019**, *20*, 208–233. [CrossRef]
30. Demoule, J.-P. *Mais où Sont Passés les Indo-Européens? Le Mythe D'origine de l'Occident*; Seuil: Paris, France, 2014.
31. Arvidsson, S. *Aryan Idols: Indo-European Mythology as Ideology and Science*; Motilal Banarsidass: New Delhi, India, 2025.
32. Czech Statistical Office. *Religious Beliefs: Population by Religious Belief*; Czech Statistical Office: Prague, Czech Republic, 2021.
33. Maiello, G.; Velkoborská, K. On being a transgender man in the Czech pagan community. *Relig. Soc. Cent. East. Eur.* **2021**, *14*, 15–25. [CrossRef]
34. Maiello, G. Nude Culture and Contemporary Czech Paganism: An Online Ethnography Interrupted. In Proceedings of the 14th SIEF Congress, Santiago de Compostela, Spain, 14–17 April 2019.
35. Maiello, G. On the agony of Czech Slavic paganism and the representation of one's own funeral among contemporary Czech pagans. *Pomegranate Int. J. Pagan Stud.* **2018**, *20*, 137–156. [CrossRef]
36. Maiello, G. Contested Homelands and Reconstructed Ancestors: Indo-European Imaginaries and Indigeneity in Contemporary Slavic Paganism. *J. Relig. Eur.* 2026, *in press*.
37. Rountree, K. Neo-paganism, Native Faith and indigenous religion. *Soc. Anthropol.* **2014**, *22*, 81–100. [CrossRef]
38. Thykier Makeeff, T. Power from the pagan past: The Viking warrior as template for (toxic) masculinity. In *Uses and Abuses of Power in Alternative Spiritualities*; Harvard Divinity School: Cambridge, MA, USA, 2023.
39. Morozov, E. *To Save Everything, Click Here: The Folly of Technological Solutionism*; PublicAffairs: New York, NY, USA, 2013.
40. Campbell, H.A.; Tsuria, R. (Eds.) *Digital Religion: Understanding Religious Practice in Digital Media*, 2nd ed.; Routledge: London, UK, 2022.
41. Sánchez Capdequí, C.; Gil-Gimeno, J.; Echeverría-Esparza, P. The Hegemonic Character of Techno-Functional Neo-immanentism and its Relationship with Culture Wars. *Religions* **2023**, *14*, 943. [CrossRef]
42. Barad, K. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*; Duke University Press: Durham, NC, USA, 2007.
43. Fuchs, C. *Digital Humanism: A Philosophy for 21st-Century Digital Society*; Emerald: Bingley, UK, 2022.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.