Why and How: Public-Sphere Persuasion and Institutional Legitimacy*

Laure Athias†

July 1, 2026

Abstract

Norms acquire institutional legitimacy when recognized as rationally justifiable. I develop a model of public-sphere persuasion in which individuals evaluate candidate norms along two dimensions: merit and implementability. Endorsement is bundle-based: a norm must be judged desirable and institutionally feasible. Because Weberian legal–rational bureaucracy gives state agents a comparative advantage in implementability arguments, peer and state-agent persuasion need not operate symmetrically. Exposure scales channel-specific signals' precision, hence persuasive force. I characterize the equilibrium persuasion architecture: when each channel bundles merit and implementability arguments, peer and state-agent exposures are substitutes; under specialization, cross-channel fit makes exposures locally complementary. I test the model using geocoded Swiss panel data on support for equal institutional treatment of in-group and out-group members. I construct channel-specific content-weighted exposures by combining predetermined canton-level reach with canton-year variation in each channel's universalism tilt. Endorsement increases with peer and state-agent exposure, while their interaction is negative, consistent with within-channel bundling. A workplace-switch design holds canton-year discourse fixed: private-to-public non-mission moves increase endorsement by 10 percentage points. Mechanism tests reveal the argumentative content: peer exposure reduces welfare-spending support, while public-sector entry reduces satisfaction with democracy. I further leverage two predetermined shifters of the bundling mechanism: cantonal church–state separation, which opens the normative *why* to public justification, and exposure to French republican state-building, which strengthens merit–implementability bundling. Across both empirical designs, the persuasion architecture moves toward within-channel bundling. Overall, cultural change is an endogenous process of norm legitimation whose pace is governed by the persuasion architecture.
***JEL* Codes:** D02, D63, D71, D73, D83, P16, Z13.

*I am grateful to James A. Robinson for discussions that have nurtured this paper and for providing an intellectual home. I also thank Marius Brülhart, Martin Castillo-Quintana, Jacob Hariri, Leander Heldring, Rachel Kranton, Eduardo Montero, Raul Sanchez de la Sierra, Judith Scheele, Mathias Thoenig, Bruno Ventelou, Nicolas Werquin, Giorgio Zanarone, and Yanos Zylberberg for helpful comments at different stages of the project. I also thank the participants in the HEC Lausanne Public Economics and Policy Seminar, University of Chicago Substantivist Economics and Societies conference, 13th World Congress of the Econometric Society, 29th Annual SIOE Conference, 23rd session of the Institutional and Organizational Economics Academy. I declare no research support, financial relationships, or other potential conflicts of interest.

†Laure Athias, University of Lausanne (Email: laure.athias@unil.ch).

"If spirit is something like the achievement of a collective normative like-mindedness, then there are such norms and we are subject to them only because we have legislated such norms," and yet they can retain "real, genuinely binding, and justifiable normative force."
Robert B. Pippin (2008), p. 17

"The validity of the claims to legitimacy may be based on rational grounds—resting on a belief in the 'legality' of patterns of normative rules and the right of those elevated to authority under such rules to issue commands. [...] In the case of legal authority, obedience is owed to the legally established impersonal order."
Max Weber, (1978), pp. 215–216

"The modern idea of order, in contradistinction to the medieval Christian ideal, was seen from the beginning as for the here and now. But it definitely migrates along a path, running from the more hermeneutic to the more prescriptive."
Charles Taylor (2004), p. 7

1 Introduction

Institutional legitimacy matters because it affects institutional performance: it shapes voluntary compliance and enforcement frictions (Tyler (2006)) and, more broadly, conditions whether materially similar institutions work across environments (Moscona et al. (2026)). In economics, legitimacy—when it appears—typically enters either through cultural beliefs and values that sustain institutional equilibria (Greif (1994), Tabellini (2008), Acemoglu and Robinson (2025)), or through direct preferences over rights, liberty, or democratic procedures (Besley and Persson (2019), Gratton and Lee (2024)). This paper studies the prior process through which legitimacy is formed: how a candidate norm comes to be recognized as worthy of institutionalization. A long tradition in political philosophy and anthropology argues that formal institutions are not merely technical devices. They give normative meanings institutional form, making the rules they enact intelligible and potentially legitimate (e.g., Castoriadis (1987), Taylor (2004), Malinowski (1922), Sahlins (1976), Sahlins (2000)). In modern societies, unlike premodern orders grounded in inherited, sacred, or customary authority, norms can claim institutional legitimacy only insofar as they are justifiable (Habermas (1990), Pippin (2008)). Legitimacy therefore turns on whether individuals can recognize a candidate norm as justified. Norms cannot then be treated only as inherited background conditions; they must also be modeled as objects of deliberation, evaluated in terms of both their normative appeal and their institutional implementability.

This paper develops a tractable theory of how institutional legitimacy emerges through persuasion in the public sphere—the arena of public reasoning in which shared judgments about public matters are formed[1]. Because the object of persuasion is a candidate norm for public institutionalization rather than private approval, the relevant question is whether the norm is worthy of institutionalization. Citizens therefore evaluate it along two latent dimensions: *merit* (m), capturing whether the norm is normatively desirable, and *implementability* (h), capturing whether it can be

[1] Taylor (2004) argues that the public sphere is central to institutional legitimacy in modern Western societies, and defines it as a common space in which members of society meet, both face-to-face and through various media, to discuss matters of common interest and "form a common mind about these" (Taylor (2004), p. 83). Habermas (1989) traces the emergence of this rational and critical discursive space of public justification to 18th-century bourgeois society, particularly in Europe, through institutions such as coffeehouses and salons, literary journals and newspapers, and learned societies and clubs, and locates it between the private realm and the state.

administered, enforced, and sustained in practice[2]. Receivers endorse when the posterior expected value of institutionalizing the norm exceeds a threshold. This expected value can be purely material or include non-material components; what matters is that receivers evaluate the consequences of realizing a candidate norm as a rule of public action. Endorsement is thus bundle-based: persuasion is strongest when citizens receive matched reasons that answer both *why* the norm should hold and *how* it can be carried out. I abstract from the extensive margin of competition across candidate norms and focus on the intensive margin: conditional on a norm being on the public agenda, I ask how persuasion shapes endorsement of it as a legitimate rule of public action.

A central source of asymmetry between producers in the public sphere is institutional role. Legal-rational bureaucracy in the sense of Weber (1922) gives state agents a comparative advantage in implementability arguments: because their authority is attached to office rather than person, and is valid only within legally defined competences, state agents are structurally required to justify what can be done, by which rules, and through which administrative means[3]. These role-based obligations spill over into the public sphere: because state agents routinely justify action through rules, competences, and administrative means, their marginal cost of producing precise implementability arguments is lower than that of peers. This motivates distinguishing two persuasion channels in the public sphere—peers and state agents—that need not operate symmetrically. I model each channel as an aggregate producer of public reasons, summarizing the content supplied by the relevant set of actors. Exposure scales informational content: greater reach makes a channel's reasons more precise and therefore more persuasive in citizens' posteriors. Each channel chooses how much precision to supply in each dimension, trading off convex effort costs against an internalized legitimacy payoff that combines instrumental gains from compliance with the normative value of living under rules they recognize as publicly justified.

This yields an equilibrium persuasion architecture. When channels supply sufficiently complete bundles internally, peer and state-agent exposures are substitutes: because each channel already provides both merit and implementability reasons, additional exposure to one channel reduces the marginal value of exposure to the other. When channels specialize across dimensions, by contrast, receivers must assemble matched argumentative bundles across sources. Specialization therefore creates cross-channel fit and complementarity at low exposure. This architecture also has dynamic implications. Because endorsement is bundle-based, implementability specialization makes persuasion contingent on cross-channel completion: state-agent exposure supplies the institutional *how* but leaves the normative *why* incomplete, while peer exposure supplies the normative *why* but leaves institutional implementability incomplete. Consensus is therefore harder to build when one type of reason is missing, but dynamic feedback can generate tipping-point dynamics once merit and implementability arguments jointly circulate, even if the normative *why* remains weakly disciplined. Without specialization, cultural change is smooth.

I take the model to geocoded individual-level panel data from the Swiss Household Panel

[2]The model does not presume that citizens explicitly debate the primitives (m, h). Public-sphere exchanges typically proceed through concrete cases, policy controversies, and institutional proposals. Such discourse indirectly conveys information about desirability and implementability, and endorsement updates as public reasons accumulate. Formally, (m, h) are sufficient statistics for the informational content of these arguments in the learning problem.

[3]The Weber epigraph is useful because legal-rational authority contains both dimensions emphasized in the model. Its legitimacy rests on the legality of normative rules, which speaks to the normative *why*, and on the right of office-holders to issue commands under those rules, which speaks to the institutional *how*. Bureaucratic role specialization makes state agents especially likely to supply the latter type of reason. This does not give state agents superior standing in the public sphere: participants have equal justificatory standing, but not identical informational positions. State agents matter because legal-rational office makes certain implementability reasons more available to them, insofar as these reasons are publicly intelligible as rule-bound claims rather than expressions of private interest or discretionary power.

(SHP). Switzerland is a useful setting because linguistic segmentation and strong decentralization generate within-country variation in culture and institutions, allowing the persuasion architecture to be identified across distinct environments. Switzerland is also an informative case: despite its stable, direct-democratic institutions, its long-run increase in affective polarization is among the steepest in the OECD, comparable to the U.S. trend (Boxell et al. (2024)). I study endorsement of universalism as a principle that can justify state action: whether Switzerland should grant foreigners the same opportunities as Swiss citizens or instead give Swiss citizens preferential opportunities. This question is a direct empirical counterpart to the model's object, since it asks respondents to evaluate a candidate rule of public action rather than report a private attitude in the abstract. Universalism—the moral principle of equal treatment across in-group and out-group members, in contrast to context-dependent moral orders in which obligations are stronger toward kin or in-groups (Enke et al. (2022))—is also a substantively important case. It has been characterized as a central feature of WEIRD moral psychology, rooted in the transformations associated with the Western Church that weakened kin-based particularism and expanded impersonal obligations (Henrich (2020)), and can be situated within the expansion of the modern Western moral order rooted in the new theories of Natural Law (Taylor (2004)). Yet Western societies differ sharply in how universalistic principles are institutionalized (Esping-Andersen (1990)), and whether foreigners should receive equal institutional treatment remains one of the most salient contemporary tests of universalism.

The empirical strategy is designed to identify how endorsement responds to variation in the persuasive force of public justification. The SHP provides annual measures of endorsement, allowing fixed-effects designs. I first operationalize the model's notion of channel-specific effective precision by combining canton-year variation in each channel's universalism tilt with predetermined canton-level reach proxies. The peer channel is built from private-sector workers, with reach proxied by association density; the state-agent channel is built from public-sector workers, with reach proxied by the public-employment share. In this exposure design, I restrict the sample to non-employed individuals so that respondents do not mechanically enter the construction of the content-weighted exposure measures, and to canton non-movers so that variation in the discourse environment is not confounded by endogenous relocation across cantons. Conditional on fixed effects, trends, and rich controls, identification comes from within-canton, over-time changes in each channel's universalism tilt, scaled by predetermined reach. These changes capture variation in the informational content circulating through state-agent and peer channels that is not explained by common shocks, canton-specific trajectories, or observed political and demographic conditions. I interpret this variation as shifts in channel-specific effective precision. The baseline specification estimates a model-disciplined reduced form in which endorsement responds to peer and state-agent effective precision and their interaction, thereby identifying both own-channel persuasiveness and the persuasion architecture. The main identification concern is that these within-canton changes in channel-specific universalism tilt may partly reflect unobserved canton-year shifts in the broader demand for universalism or in the local discourse environment, rather than channel-specific persuasion alone. I therefore complement this design with a second strategy that delivers sharper individual-level exposure shocks while absorbing the canton-year discourse environment through canton-year fixed effects: within-person switches between private and public employment. To address concerns about dynamic self-selection and mission-driven sorting, I focus on private-to-public switches within non-mission industries, excluding health, education, and social care. In these industries, I find no evidence of differential pre-trends in endorsement prior to the switch.

The empirical results support the model along two margins: the own effect of each channel and the interaction between them. In the baseline content-weighted exposure regressions, endorsement rises with both peer and state-agent effective precision, while their interaction is negative. In the

model, this pattern is consistent with within-channel bundling and diminishing returns, implying substitutability across channels. In magnitude, peer effective precision has the larger marginal effect, but state-agent effective precision also has a sizeable effect conditional on the peer environment and the full set of fixed effects and controls: a one-standard-deviation increase in peer content-weighted exposure raises endorsement by 5.5 percentage points, evaluated at mean state-agent exposure, while the corresponding increase in state-agent content-weighted exposure raises it by 4.4 percentage points. The positive own effects rule out a pure crowd-out interpretation, while the negative interaction rules out a purely additive diffusion process: exposure matters not only because individuals encounter more universalist agents, but because peer and state-agent channels carry partly overlapping bundles of normative *why* and institutional *how* arguments.

The workplace-switch design provides a sharper exposure shock. Because treatment varies at the individual-year level, I can include canton-year fixed effects and hold fixed the local discourse environment. Moving from private to public employment within non-mission industries raises endorsement by roughly 10 percentage points. This larger within-person estimate is consistent with attenuation in the baseline canton-level reach proxies and with the model's prediction that endorsement rises with effective exposure to persuasive public reasons. The switch design also reveals a sharp institutional-tier pattern: the implied net effect of federal public employment relative to private employment is close to zero, while the effect is concentrated among cantonal and communal state agents. This reinforces the choice of the canton-year as the empirical boundary of the public sphere in this setting. In the model's terms, the relevant persuasive state-agent channel is not generic public employment status, but exposure to state agents operating at the level of public authority where the merit of a candidate norm can be bundled with the rules, procedures, and administrative means through which that norm is publicly justified and institutionally implemented. Finally, the reverse-switch estimate is small and statistically insignificant. This pattern is consistent with persistence under learning: once beliefs have shifted, removing exposure need not mechanically restore prior beliefs absent offsetting new information.

Adjacent-outcome mechanism tests both rule out generic ideological diffusion or occupational conformity and help identify the argumentative content through which universalism is made institutionally legitimate. In the baseline content-weighted exposure design, peer exposure does not simply shift respondents toward more expansive egalitarian preferences: it significantly reduces support for welfare spending. This pattern indicates that persuasion in favor of universalism is not simply a shift toward pro-social or left-wing preferences. Instead, peer discourse appears to carry a more specific argumentative content: equal institutional treatment can be endorsed as a rule of public action while remaining compatible with limits on public provision. Similarly, in the workplace-switch design, moving into public employment does not generate broad institutional loyalty. If anything, it lowers satisfaction with democracy while increasing endorsement of institutional universalism. This pattern is difficult to reconcile with simple imitation or generic public-sector socialization. It is instead consistent with state agents making visible the gap between universalistic rules and existing democratic performance. Exposure to state agents can therefore make individuals more supportive of universalistic rules while also more critical of how democratic institutions currently function.

I then identify the bundle-based persuasion mechanism by leveraging predetermined cantonal environments that should shift the persuasion architecture ex ante. Because the argumentative content circulating within each channel is not directly observed, the empirical strategy is not to claim that these environments directly measure merit or implementability. Rather, I use theoretically and historically grounded shifters that alter whether merit and implementability are more likely to travel together within channels or to be assembled across channels. The key test is cross-design coherence: shifters that move the content-weighted exposure design toward within-channel bundling should also moderate the workplace-switch effect in the predicted direction, even though the two

designs rely on different sources of variation.

I first use cantonal church–state separation. In modern Western secular orders, common institutions are no longer grounded as directly in sacred or ecclesiastical authority; their legitimacy must instead be articulated through public reason within a political order of popular sovereignty (Taylor (2004), Taylor (2007)). Cantonal church–state separation measures how far this secular logic is carried into institutional design. Where separation is stronger, the common institutional order is less inherited as an already-given sacred order and more open to public justification as a set of principles to be institutionally realized. The state is therefore not merely an instrument for implementing given norms; it becomes part of the process through which norms are publicly justified and institutionally realized. Church–state separation thus opens the normative *why* to justification in the public sphere. In the model, this increases the supply of complete justificatory bundles within channels: because merit must also be articulated, a given channel is more likely to provide both the normative *why* and the institutional *how*. I operationalize this variation using cantonal church-tax regimes: in some cantons, the state is required to levy a church tax, whereas in others it is not. Consistent with this mechanism, cantonal church–state separation shifts the exposure–endorsement mapping toward within-channel bundling in the baseline design and increases the private-to-public switching effect. This pattern is not driven simply by secular receiver composition. When I control for canton-level support for church–state separation, measured by cantonal vote shares in a 1980 referendum on state secularity, and interact this support with peer exposure, state-agent exposure, and their interaction, the institutional church–state shifter continues to move the conditional exposure–endorsement mapping.

The second shifter, exposure to French republican state-building, operates through a different margin. Whereas church–state separation shifts the regime of justification by detaching common institutions from inherited ecclesiastical authority, French republican state-building shifts the logic through which normative claims are connected to institutional realization. In the republican tradition associated with Rousseauian political thought, legitimate law is not an external command imposed on individuals; it is binding because citizens can understand it as an expression of collective self-legislation, or the general will (Rousseau (1923)). Under the Third Republic, this state-centered grammar was institutionalized by tying public administration to substantive civic and moral ends (Weber (1976)). This grammar makes the link between normative merit and legal-political realization especially salient: claims about what is right are more likely to be articulated together with claims about how public authority can institutionalize them. I use French-speaking Swiss cantons, and French-speaking identity in the workplace-switch design, as proxies for exposure to this cultural grammar of state legitimacy[4]. The empirical pattern matches the mechanism. In the content-weighted exposure design, French-speaking environments strengthen the peer-exposure effect and make the peer–state-agent interaction more negative, consistent with greater within-channel bundling and increased substitution between channels. The pattern is not absorbed when the exposure–endorsement mapping is allowed to vary with receiver-side support for state secularity or predetermined cantonal ideology. The switch design provides the key cross-design check: because it compares workers within the same canton-year environment, it holds fixed contemporaneous cantonal institutions and local discourse while asking whether the same direct exposure to state agents has greater persuasive force for French-speaking individuals. It does: the public-employment effect is about three times larger for French-speaking workers than for German-speaking workers.

Taken together, these results support a bundle-based persuasion mechanism and the model's view that predetermined institutional and cultural environments shift the mapping from effective precision

[4]I use language regions as cultural exposure, not as institutional treatment. This usage follows a literature exploiting the Swiss language border to identify cultural differences while holding formal institutions fixed (e.g., Eugster et al. (2011), Eugster and Parchet (2019)).

into endorsement—the persuasion architecture—rather than merely the level of universalism.

The mechanism is not specific to Switzerland, but its portability depends on scope conditions. The model suggests three such conditions. First, norms must be contestable: they must be open to public questioning rather than treated as inherited background. Second, public reasons must circulate through social networks and media, so that individuals can encounter arguments about both merit and implementability. Third, producers of reasons must internalize positive legitimacy payoffs, so that supplying persuasive arguments is consequential. Where these conditions hold, institutions can shape persuasion not only by changing incentives or policies directly, but by altering the architecture through which citizens come to see a norm as both desirable and institutionally realizable. External validity therefore depends less on Swiss institutions per se than on whether other settings generate similar conditions for public contestation, channel reach, and merit–implementability bundling. The Swiss setting is useful because it combines decentralized authority and rich subnational variation. At the same time, it is not a case in which direct democracy mechanically guarantees broad political participation: turnout has historically been modest despite extensive democratic rights (Hosli et al. (2025)). This makes Switzerland a demanding setting for the argument. If institutions shape endorsement, they do so not simply by increasing formal opportunities to participate, but by structuring the public-reason environment in which institutional legitimacy is formed. The framework also helps interpret a broader tension. If public discourse increasingly separates the normative *why* from the institutional *how*, consensus may become harder even in a stable direct-democratic system. This offers a perspective on an apparent Swiss paradox: a country known for institutional stability and direct democracy can nevertheless experience a steep long-run rise in affective polarization (Boxell et al. (2024)).

This paper contributes to the literature on culture and institutions by formalizing how norms become recognized as worthy of institutionalization. Norms are not only inherited cultural background conditions or informal constraints (North (1990), Greif (1994), Tabellini (2008), Bursztyn et al. (2020)); in modern societies, they can claim institutional legitimacy insofar as they can be justified and realized through institutions (Pippin (2008)). This paper defines institutional legitimacy as the recognition of a candidate norm as both normatively desirable and institutionally feasible. Legitimacy is therefore bundle-based: persuasion must answer not only why a norm should hold, but also how it can be administered, enforced, and sustained in practice. The argument therefore moves from persuasion within bounded social spheres to persuasion over norms that claim general public validity. Families, markets, associations, religious communities, and firms all institutionalize norms, but such norms are typically justified within bounded communities. Norms acquire public institutional legitimacy only when they claim general validity and can be defended as rules of public action. The public sphere is the arena in which such claims to general validity are tested (Habermas (1989), Taylor (2004)). This implies a social rather than purely individual account of agency: individuals appear not only as bearers of private preferences, but as participants in public justification over rules whose legitimacy matters both instrumentally, by sustaining voluntary compliance (Tyler (2006)), and non-instrumentally, by shaping the norms under which they collectively live (Pippin (2008)). This perspective contrasts with social-capital accounts, which emphasize trust and cooperation as the channels through which civic engagement and associational life sustain democratic performance (Putnam et al. (1993), Putnam (2000)). Here, civic engagement and associational life matter instead as infrastructure for the public sphere: they sustain the free and informal circulation of public reasons. As in Bayesian persuasion models, persuasion operates by changing beliefs rather than preferences (Kamenica and Gentzkow (2011)); in this setting, the relevant beliefs concern the merit and implementability of institutionalizing a candidate norm. Consequently, private preferences over the norm may diverge from support for its institutionalization (Athias and Ventelou (2026)). This distinction also sharpens the paper's contrast with accounts that locate Western distinctiveness in a

specific bundle of private values or psychological traits (Henrich (2020)). Western societies may not be distinctive by virtue of sharing a fixed bundle of values, but by their capacity to justify and institutionalize candidate norms. While Acemoglu and Robinson (2012) stress the role of inclusive institutions in long-run prosperity, this paper therefore highlights a prior legitimacy condition for their effectiveness: the continued public justification of the inclusive norms they institutionalize.

This paper highlights persuasion in the public sphere as a channel through which culture can become fluid: public persuasion allows normative claims to be justified, revised, and translated into institutional change. It thereby contributes to the literature on cultural change by formalizing cultural change as an endogenous process of norm legitimation in the public sphere. The pace of cultural change is determined by the persuasion architecture through which norms acquire institutional legitimacy, with discontinuous change arising when persuasion generates tipping dynamics (Acemoglu and Jackson (2015), Acemoglu and Robinson (2025)). This dynamic structure parallels accounts of economic growth in which sustained change depends on the accumulation of useful knowledge linking new ideas to the means by which they can be made actionable (Mokyr (2010), Mokyr (2016)). The paper identifies an analogous structure for the expansion of the moral order associated with modern natural-law theories (Taylor (2004)): moral expansion requires not only publicly disciplined claims about what is right, but also credible accounts of how such claims can be realized through institutions.

The link between positive persuasion and normative justification connects the paper to welfare economics and social choice, where a central question is how impartial principles can be justified (Harsanyi (1955), Rawls (1971), Rawls (1997)). The state is central in this process because it is the institutional form through which candidate norms can be scaled into general, enforceable, and administratively sustained rules of public action. In this sense, the state is not merely a contract, an apparatus of enforcement and coercion, or a mechanism for aggregating private interests (Besley and Persson (2011), Besley (2020), Tilly (1975), Tilly (1985), Arrow (1951)); it is the institutional form through which publicly justified norms can become actual as rules of public action (Pippin (2008)). This paper shows that state secularity strengthens the mechanism. When public authority is intertwined with religious authority, the state implements a given sacred order (Graeber and Sahlins (2017), Acemoglu and Robinson (2025)) rather than mediating among publicly contestable reasons. State secularity therefore loosens religious control over public authority and opens the normative *why* to the public sphere, making the state a site in which collective norms can be justified and realized as public rules. This resonates with Taylor (2007)'s account of secular modernity: secularity is not the subtraction of religion from public life, but a transformation in the conditions under which claims about moral order become contestable and require justification. The paper also clarifies why direct-democratic authorization is not a sufficient foundation of institutional legitimacy. Authorization makes public justification consequential, but legitimacy first requires candidate norms to be publicly justifiable and then realized through institutions capable of sustaining them as rules of public action. This distinction connects to Castoriadis (1996)'s account of autonomy as collective self-institution rather than the mere presence of elections, rights, or constitutional form. It also helps make sense of Luo and Xue (2025), who show that liberal principles can diffuse through moral philosophy and generate political consequences even in autocratic settings. This creates a distinct form of state capture: even when a norm has acquired public legitimacy, political elites may prevent it from becoming actual by blocking, delaying, or selectively implementing the institutional conditions required to realize it. This implication connects to the political-economy literature on state capture, de facto power, institutional persistence, and captured democracy (Acemoglu and Robinson (2006), Acemoglu and Robinson (2008)), as well as to work on state capacity (Besley and Persson (2009), Besley and Persson (2011)), but shifts attention from control over rents, political power, or enforcement capacity to control over the institutional *how* through which publicly justified

norms become actual.

Finally, the paper speaks to the literature on bureaucracy by identifying a double-edged role for bureaucratic organization in cultural change. The economics literature has emphasized how bureaucratic organization shapes state capacity, implementation, and public-sector performance (Besley et al. (2022), Aneja and Xu (2024)). This paper shows that bureaucratic roles also have public-sphere spillovers. When public-sector roles are organized around Weberian legal–rational authority (Weber (1922)), state agents acquire practical knowledge about what can be done, by which rules, and through which administrative procedures. This institutional role makes them especially important in public persuasion: even outside formal decision-making, their public-sphere discourse can make candidate norms appear more or less institutionally realizable. This connects the paper to models of competition in persuasion, which study how persuasion changes when receivers hear from more than one source (Gentzkow and Kamenica (2017)). Here, however, the relevant sources are not merely competing senders choosing information structures; they are socially distinct channels—peers and state agents—whose institutional roles shape the kinds of reasons they can credibly supply. At the same time, because the state is also an institution through which norms are realized, cultural and historical environments shape whether state agents transmit implementability arguments alone or bundle them with merit reasons (Ekeh (1975), Weber (1976)). State-agent spillovers can therefore discipline legitimacy formation when implementability is tied to a well-articulated normative *why*, but they can also amplify whichever merit claim prevails when they supply institutional *how* without strengthening the informational discipline of the merit case itself. This distinction speaks to a classic tension in accounts of rationalization: legal–rational bureaucracy can sustain publicly contestable justification, but implementability-oriented discourse detached from substantive merit can also facilitate rapid legitimation of weakly disciplined or authoritarian normative claims (Horkheimer and Adorno (1944), Habermas (1989)). In this sense, the paper highlights a potential dark side of legal–rational bureaucratic capacity (Heldring (2026)).

The remainder of the paper is structured as follows. Section 2 presents a model of institutional legitimacy formation in the public sphere, in which persuasion operates over a candidate norm for institutional realization, and derives the conditions under which cultural change is smooth or rapid. Section 3 maps the model to empirical counterparts and lays out the identification strategies. Section 4 contains details on the data. I report the baseline results and the workplace-switch evidence in Section 5, and provide mechanism tests in Section 6. Section 7 concludes.

2 A Model of Public-Sphere Persuasion

This section formalizes institutional legitimacy as the recognition of a candidate norm as both justified and realizable through institutions. Conditional on a norm being on the public agenda, I ask how persuasion shapes endorsement of it as a rule of public action. The object of persuasion is therefore not a norm in the abstract, but a norm proposed for institutional realization. Citizens must assess not only whether the norm is desirable, but also whether it can be realized in practice. I capture this by assuming that citizens evaluate candidate norms along two dimensions: *merit*, which captures whether the norm is desirable and may reflect both material and non-material considerations, and *implementability*, which captures whether the norm can be administered, enforced, and sustained in practice. Endorsement is therefore bundle-based. The model does not presume that citizens explicitly debate the primitives (m, h). Public-sphere exchanges typically proceed through concrete cases, policy controversies, and institutional proposals. Such discourse indirectly conveys information about desirability and implementability, and endorsement updates as public reasons accumulate. Formally, (m, h) are sufficient statistics for the informational content of these arguments in the

learning problem. A useful illustration is the Zucman proposal for a coordinated minimum taxation standard on ultra-high-net-worth individuals (Zucman (2024)). What is at stake is not merely support for a tax instrument, but endorsement of a norm for institutionalization: that the ultra-rich should not be governed by a different fiscal principle from everyone else. Its persuasive force depends both on merit—equal civic standing, tax fairness, and the material consequences of the proposed fiscal rule—and on implementability—the claim that this principle can now be realized through workable domestic instruments and international administrative coordination. The model formalizes this persuasion problem.

2.1 Model Setup

A continuum of citizens i in a public sphere c at date t considers whether to endorse a norm as a rule of public action. Let $y_{i,c,t} \in \{0,1\}$ denote the observed public stance, and

$$s_{i,c,t} \;\equiv\; \mathbb{E}[y_{i,c,t} \,|\, \mathcal{F}_{i,c,t}] \in [0,1]$$

be the endorsement propensity given citizen i's information set $\mathcal{F}_{i,c,t}$, generated by the public arguments to which she is exposed. Public arguments are reasons circulating in publicly contestable form. I define the *public sphere* in cell (c,t) as the set of venues and interactions in which reasons about common affairs are produced and exchanged. Concretely, it includes (i) mass and local media; (ii) open civic forums, such as associations, clubs, and assemblies; and (iii) informal interpersonal communication. In the model, citizens observe only the subset of arguments to which they are exposed; I formalize this below through the information set $\mathcal{F}_{i,c,t}$. I abstract from censorship: any agent in principle contribute to and access this discourse, though exposure varies across citizens.

Endorsement concerns institutionalization, so two latent dimensions matter for persuasion: a *merit* dimension m, which captures whether the norm is desirable and may reflect material consequences, normative appeal, or both, and an *implementability* dimension h, which captures whether the norm can be administered, enforced, and sustained in practice so that its promised consequences can be realized. Citizen i's information set $\mathcal{F}_{i,c,t}$ induces posterior beliefs about both dimensions. Let

$$M_{i,c,t} \equiv \mathbb{E}[m \mid \mathcal{F}_{i,c,t}], \qquad H_{i,c,t} \equiv \mathbb{E}[h \mid \mathcal{F}_{i,c,t}]$$

denote citizen i's posterior assessments of merit and implementability. I assume that endorsement depends on a latent index

$$\Lambda_{i,c,t} = \ell(M_{i,c,t}, H_{i,c,t}),$$

where $\ell_M > 0$, $\ell_H > 0$, and $\ell_{MH} > 0$. This index represents the posterior value of institutionalizing the norm. Citizen i endorses the norm when this value exceeds an endorsement threshold. Equivalently, the observed endorsement propensity satisfies

$$s_{i,c,t} = \mathbb{E}[y_{i,c,t} \mid \mathcal{F}_{i,c,t}] = u(\Lambda_{i,c,t}),$$

where $u(\cdot)$ is increasing, with the endorsement threshold normalized into u. The complementarity captures the bundle-based nature of endorsement: learning that a norm is desirable has a larger effect when citizens also believe it can be implemented, and learning that it can be implemented has a larger effect when citizens also believe it is desirable.

Priors are Gaussian and independent:

$$m \sim \mathcal{N}(\mu_m, \sigma_m^2), \qquad h \sim \mathcal{N}(\mu_h, \sigma_h^2).$$

The supply of public reasons is costly. Producing persuasive arguments requires effort because reasons must be formulated in ways that are intelligible, contestable, and informative about

either merit or implementability. Merit arguments and implementability arguments draw on different knowledge bases. Merit arguments concern the normative and material desirability of the candidate norm. Implementability arguments concern the institutional *how*: whether the norm can be translated into rule-bound, administratively feasible, enforceable, and sustainable public action. A central source of heterogeneity among producers is institutional role. In particular, legal-rational bureaucracy in the sense of Weber (1922) gives state agents a comparative advantage in implementability arguments. Because their authority is attached to office rather than person, and is valid only within legally defined competences, state agents are routinely required to justify public action by reference to rules, competences, procedures, and administrative means. These role-based obligations spill over into the public sphere: relative to other citizens, state agents have a lower marginal cost of producing precise reasons about what can be done, by which rules, and through which administrative instruments. This does not give state agents superior standing in public justification. Participants have equal justificatory standing, but not identical informational positions.

This motivates distinguishing two aggregate producer channels in the public sphere. The first is state agents, denoted $j = S$: public employees whose public-sphere discourse spills over from their institutional roles. The second is peers, denoted $j = P$: other citizens participating in the same public sphere. Each channel summarizes the aggregate content supplied by the relevant set of actors. I do not model coordination among producers: unlike in families, firms, or organizations, public-sphere discourse is not governed by a single authority coordinating the content or precision of the reasons supplied. Channels therefore choose the content and precision of their arguments independently, both within and across channels.

I fix a public sphere (c, t) and suppress (c, t) subscripts in producer choices. Producing precision π_j^k on dimension k entails a convex reduced-form cost $K_j^k(\pi_j^k)$, with $K_j^{k\prime}(\pi) > 0$ and $K_j^{k\prime\prime}(\pi) \geq 0$. I impose the role-based implementability advantage as follows.

Assumption 1 (Role-based implementability advantage) *There exists $\delta > 0$ such that for all $\pi \geq 0$,*

$$K_S^{h\prime}(\pi) \leq (1 - \delta) K_P^{h\prime}(\pi).$$

On the merit dimension, the two channels are symmetric:

$$K_S^{m\prime}(\pi) = K_P^{m\prime}(\pi) \qquad \forall \pi \geq 0.$$

The assumption captures the idea that state agents are not intrinsically more persuasive. Their comparative advantage is role-based and dimension-specific: legal-rational office makes implementability reasons more available to them as publicly intelligible, rule-bound claims rather than as expressions of private interest or discretionary authority.

Assumption 2 (Legitimacy has social value) *Let*

$$S_{c,t} \equiv \int s_{i,c,t}, di$$

denote aggregate endorsement of the candidate norm in public sphere (c, t). Public endorsement is valuable because it can improve institutional performance through higher voluntary compliance and lower enforcement burdens, and because individuals may attach non-instrumental value to living under rules they recognize as publicly justified. Higher legitimacy raises a common payoff $V(S_{c,t})$, with $V'(S) > 0$. For the baseline analysis, I use the normalization

$$V(S) = BS, \qquad B > 0.$$

Imperfect internalization by producers. Each channel $j \in S, P$ is modeled as a representative producer of public reasons. Its precision choices affect the information received by a positive mass of citizens through reach, and therefore aggregate endorsement $S_{c,t}$. Channel j internalizes a fraction $\beta_j \in [0,1]$ of the common legitimacy payoff $V(S_{c,t})$. Its objective therefore scales the social value of legitimacy by β_j, generating an internalization wedge when $\beta_j < 1$.

The parameter β_j is a reduced-form measure of the marginal producer's internalization of instrumental and non-instrumental legitimacy payoffs, given whatever extrinsic or intrinsic motives make the supply of public reasons individually worthwhile. The model therefore abstracts from within-channel collective-action problems. Institutional environments can shift β_j: direct-democratic rights, for example, make public persuasion more consequential for rule-making, and therefore increase the expected payoff from supplying public reasons.

For each channel j, individual exposure within (c,t) decomposes as

$$E^j_{i,c,t} = E^j_{c,t} + \eta^j_{i,c,t}, \qquad \mathbb{E}[\eta^j_{i,c,t} \mid c,t] = 0,$$

where $E^j_{c,t}$ is aggregate reach in public sphere (c,t) for channel j, and $\eta^j_{i,c,t}$ is idiosyncratic reach. I assume that $\eta^j_{i,c,t}$ has support such that $E^j_{i,c,t} \geq 0$ almost surely. Citizens condition on their realized exposure.

Assumption 3 (Within-period predetermined reach) *Within (c,t), channel reach $E^j_{c,t}$, for $j \in \{S,P\}$, is determined by many micro-broadcasts and information-environment constraints, such as communication technology and media-market reach. Any single citizen or individual sender is negligible, so aggregate reach $E^j_{c,t}$ is taken as predetermined within the period.*

Channel reach $E^j_{c,t}$ is therefore predetermined within the period. By contrast, each channel's precision choice is non-atomistic: changing π^k_j shifts the informativeness of the signals received by the positive mass of citizens exposed through reach, and therefore affects aggregate endorsement.

Assumption 4 (Exposure–signal aggregation) *For each channel $j \in \{S,P\}$ and dimension $k \in \{m,h\}$, citizen i observes an aggregated signal*

$$\tilde{x}^k_{j,i,c,t} = \begin{cases} m + \tilde{\varepsilon}^m_{j,i,c,t}, & k = m, \\ h + \tilde{\varepsilon}^h_{j,i,c,t}, & k = h. \end{cases}$$

If $\pi^k_j E^j_{i,c,t} > 0$,

$$\tilde{\varepsilon}^k_{j,i,c,t} \sim \mathcal{N}\Big(0,\ (\pi^k_j E^j_{i,c,t})^{-1}\Big).$$

If $\pi^k_j E^j_{i,c,t} = 0$, no informative signal on (j,k) is observed; equivalently, the posterior contribution of that channel-dimension pair is zero. Signals are independent across j and k conditional on (m,h).

Citizen i's information set in (c,t) is

$$\mathcal{F}_{i,c,t} \ = \ \sigma\Big(E^S_{i,c,t},\, E^P_{i,c,t},\, \tilde{x}^m_{S,i,c,t},\, \tilde{x}^h_{S,i,c,t},\, \tilde{x}^m_{P,i,c,t},\, \tilde{x}^h_{P,i,c,t}\Big),$$

where inactive channel-dimension signals are treated as null observations.

Microfoundations. The mean-zero signal structure reflects the public and contestable nature of the arguments. Public-sphere claims are citable, answerable, and open to counter-argument; with many micro-senders and heterogeneous views, aggregation washes out idiosyncratic distortions,

leaving no systematic drift in either dimension. The precision formula can be microfounded in two equivalent ways. First, citizen i receives $E^j_{i,c,t}$ independent repetitions of a basic Gaussian signal from channel j on dimension k, and averaging yields cumulative precision $\pi^k_j E^j_{i,c,t}$. Second, information arrives as a continuous-time flow with rate π^k_j over an exposure window of length $E^j_{i,c,t}$, again yielding cumulative precision $\pi^k_j E^j_{i,c,t}$. Thus, exposure and content precision are complements in the production of posterior precision.

2.2 Analysis

2.2.1 Static analysis

Posterior with exposure-scaled precision. Define prior precisions $\tau_m \equiv \sigma_m^{-2}$, $\tau_h \equiv \sigma_h^{-2}$, and the total effective precision on dimension $k \in \{m, h\}$,

$$Q_{k,i,c,t} \;\equiv\; \sum_{j \in \{S,P\}} \pi^k_j \, E^j_{i,c,t}.$$

Reach scales informativeness linearly: each additional unit of exposure to channel j adds π^k_j units of precision on dimension k.

Lemma 1 *Under the Gaussian priors and Assumption 4, posteriors decouple by dimension and are Gaussian with precision $\tau_k + Q_{k,i,c,t}$. In particular,*

$$m \mid \{\tilde{x}^m_{j,i,c,t}\}_j \sim \mathcal{N}\left(\frac{\tau_m \mu_m + \sum_j (\pi^m_j E^j_{i,c,t})\, \tilde{x}^m_{j,i,c,t}}{\tau_m + \sum_j \pi^m_j E^j_{i,c,t}},\ (\tau_m + \sum_j \pi^m_j E^j_{i,c,t})^{-1} \right),$$

$$h \mid \{\tilde{x}^h_{j,i,c,t}\}_j \sim \mathcal{N}\left(\frac{\tau_h \mu_h + \sum_j (\pi^h_j E^j_{i,c,t})\, \tilde{x}^h_{j,i,c,t}}{\tau_h + \sum_j \pi^h_j E^j_{i,c,t}},\ (\tau_h + \sum_j \pi^h_j E^j_{i,c,t})^{-1} \right),$$

where channel-dimension pairs with $\pi^k_j E^j * i, c, t = 0$ are omitted from the relevant sums.

Proof 1 (Sketch) *Conditional on each dimension k, signals are independent Gaussians with precisions $\pi^k_j E^j_{i,c,t}$. Gaussian conjugacy implies that precisions add and that the posterior mean is a precision-weighted average of the prior mean and observed signals. Independence across dimensions implies that the posteriors for m and h decouple.*

Benchmark: within-channel bundling (no cross-channel fit). Fixing (i, c, t) and suppressing subscripts, for each channel $j \in \{S, P\}$ and dimension $k \in \{m, h\}$, the effective precision is

$$Q^k_j \;\equiv\; \pi^k_j \, E^j \;\geq 0.$$

In the benchmark, endorsement requires a *matched pair* of reasons (m, h) that is assembled *within the same channel*. A convenient Poisson microfoundation is that, in channel j, opportunities to encounter a persuasive merit argument arrive with precision Q^m_j and opportunities to encounter a persuasive implementability argument arrive with precision Q^h_j, independently. The probability that channel j delivers a complete (matched) bundle is then

$$\Lambda_j(Q^m_j, Q^h_j) \;\equiv\; \left(1 - e^{-Q^m_j}\right)\left(1 - e^{-Q^h_j}\right).$$

Let $\Lambda \equiv \Lambda_S + \Lambda_P$ denote the total complete-bundle index across channels, and let endorsement be

$$s(E^S, E^P) \;=\; u(\Lambda),$$

where $u : \mathbb{R}_+ \to [0, 1)$ is $\mathcal{C}^2$, increasing and concave ($u' > 0$, $u'' \leq 0$). I normalize $u'(0) = 1$. A canonical case is $u(\Lambda) = 1 - e^{-\Lambda}$.

Proposition 1 (Within-channel bundling generates substitutability across channels) *Under the benchmark above, endorsement is weakly increasing in each exposure and submodular:*

$$\frac{\partial s}{\partial E^S} \geq 0, \qquad \frac{\partial s}{\partial E^P} \geq 0, \qquad \frac{\partial^2 s}{\partial E^S \, \partial E^P} \leq 0.$$

Proof 2 (Sketch) *Since $\Lambda = \Lambda_S(E^S) + \Lambda_P(E^P)$, we have $\partial s / \partial E^j = u'(\Lambda)\, \Lambda_j'(E^j) \geq 0$. Moreover,*

$$\frac{\partial^2 s}{\partial E^S \, \partial E^P} = u''(\Lambda)\, \Lambda_S'(E^S)\, \Lambda_P'(E^P) \;\leq\; 0,$$

because $u'' \leq 0$ and $\Lambda_j' > 0$ whenever $\pi_j^m + \pi_j^h > 0$ and $E^j > 0$.

Interpretation. Monotonicity ($u' > 0$) captures that additional reach increases the arrival of complete argument bundles and hence endorsement. Concavity ($u'' \leq 0$) reflects diminishing returns from saturation, limited attention, and audience overlap. In the benchmark, $\Lambda = \Lambda_S(E^S) + \Lambda_P(E^P)$ is additively separable across channels, so the cross–partial is negative only through the curvature of u: $s_{SP} = u''(\Lambda)\Lambda_S'(E^S)\Lambda_P'(E^P) \leq 0$. This is therefore a prediction about how complete persuasive bundles combine at the endorsement stage, not about exposure mechanics per se.

Producer choice in the benchmark. In the benchmark, complete argument bundles are assembled within channels. I therefore model precision choices $\{\pi_j^m, \pi_j^h\}_{j \in \{S,P\}}$ as determined by within-channel optimization. The purpose of this step is to show how reach, internalized legitimacy payoffs, and role-based marginal costs map into equilibrium precision choices, and hence into the completeness of each channel's argumentative bundle. In the subsequent specialization extension, I take these content mixes as given in order to isolate how cross-channel fit changes the sign of the exposure cross-partial.

Lemma 2 (Benchmark producer optimality conditions) *Fix a public sphere (c, t) and suppress (c, t) subscripts. Channel $j \in \{S, P\}$ has aggregate reach $E^j \geq 0$ and chooses dimension-specific precisions $\pi_j^m, \pi_j^h \geq 0$. Under within-channel bundling, the complete-bundle index supplied by channel j is*

$$\Lambda_j(\pi_j^m, \pi_j^h; E^j) \;=\; \big(1 - e^{-\pi_j^m E^j}\big)\big(1 - e^{-\pi_j^h E^j}\big).$$

Let

$$\Lambda = \Lambda_j + \Lambda_{-j}$$

denote the total complete-bundle index, taking the other channel's precision choices as given. Channel j chooses (π_j^m, π_j^h) to maximize its internalized legitimacy payoff net of precision costs:

$$\max_{\pi_j^m, \pi_j^h \geq 0} \; \beta_j B u(\Lambda) \;-\; \sum_{k \in m,h} K_j^k(\pi_j^k), \qquad j \in S, P.$$

For an interior solution, the first-order condition for dimension $k \in m, h$ is

$$\beta_j B u'(\Lambda) \frac{\partial \Lambda_j}{\partial \pi_j^k} \;=\; K_j^{k\prime}(\pi_j^k).$$

Equivalently,

$$\beta_j B u'(\Lambda) E^j e^{-\pi_j^m E^j} \left(1 - e^{-\pi_j^h E^j}\right) \;=\; K_j^{m\prime}(\pi_j^m),$$

and

$$\beta_j B u'(\Lambda) E^j e^{-\pi_j^h E^j} \left(1 - e^{-\pi_j^m E^j}\right) \;=\; K_j^{h\prime}(\pi_j^h).$$

Interpretation. The optimality conditions show that precision investment equates the marginal legitimacy benefit of making a channel's bundle more complete with the marginal cost of producing precision. The marginal return to precision is increasing in the internalized legitimacy payoff $\beta_j B$ and depends on reach E^j, which determines how many citizens receive the channel's arguments and how quickly the channel approaches saturation. The conditions also display within-channel complementarity: the marginal return to merit precision is higher when implementability precision is higher, and conversely. Lemma 2 shows how Assumption 1 shifts equilibrium precision choices in a simple way. Holding fixed reach, internalized legitimacy payoffs, and the marginal endorsement value of complete bundles, state agents face a lower marginal cost of implementability precision. Their precision choice is therefore tilted toward the institutional *how*: rule-bound, administratively intelligible arguments about what can be done, by which rules, and through which instruments. I use this result as the supply-side foundation for the specialization extension below. The point is not that state agents are intrinsically more persuasive, but that legal-rational office makes implementability reasons less costly for them to produce.

In what follows, I take the precision vector $\{\pi_j^m, \pi_j^h\}$ as given (determined by the benchmark within-channel production problem) and study how specialization shapes the mapping from exposures (E^S, E^P) to endorsement through cross-channel fit.

Extension: cross-channel fit under specialization. When channels specialize, citizens may assemble a matched pair across channels: a merit reason from one channel and an implementability reason from the other. I capture this possibility by adding a cross-channel bundle index

$$\Lambda_\times \;\equiv\; \left(1 - e^{-Q_S^m}\right)\left(1 - e^{-Q_P^h}\right) \;+\; \left(1 - e^{-Q_S^h}\right)\left(1 - e^{-Q_P^m}\right),$$

and setting $s(E^S, E^P) = u(\Lambda_S + \Lambda_P + \Lambda_\times)$. This extension isolates cross-channel fit: exposure to one channel can raise the marginal persuasive value of exposure to the other when each channel supplies a different component of the justificatory bundle.

Proposition 2 (Cross-channel fit generates local complementarity) *Fix (i, c, t) and suppress subscripts. Let $f(x) \equiv 1 - e^{-x}$ and $Q_j^k \equiv \pi_j^k E^j$ for $j \in \{S, P\}$ and $k \in \{m, h\}$. Let Λ_S and Λ_P be the within-channel complete-bundle indices defined above, and define the cross-channel complete-bundle index*

$$\Lambda_\times \;\equiv\; f(Q_S^m) f(Q_P^h) + f(Q_S^h) f(Q_P^m).$$

Let total bundle index be $\Lambda \equiv \Lambda_S + \Lambda_P + \Lambda_\times$ and endorsement be $s(E^S, E^P) = u(\Lambda)$, where $u : \mathbb{R}_+ \to [0, 1)$ is $\mathcal{C}^2$ with $u'(0) > 0$. Then

$$\left.\frac{\partial^2 s}{\partial E^S \, \partial E^P}\right|_{E^S = E^P = 0} \;=\; u'(0)\, \kappa, \qquad \kappa \;\equiv\; \pi_S^m \pi_P^h + \pi_S^h \pi_P^m.$$

In particular, if $\kappa > 0$ then there exists $\varepsilon > 0$ such that $\partial^2 s / \partial E^S \partial E^P > 0$ for all $(E^S, E^P) \in [0, \varepsilon]^2$.

Proof 3 (Sketch) *At $E^S = E^P = 0$, we have $Q_j^k = 0$ and $f(0) = 0$, $f'(0) = 1$. Thus Λ_S depends only on E^S and Λ_P only on E^P, so they contribute no cross-partial at the origin. The cross term satisfies*

$$\left.\frac{\partial^2 \Lambda_\times}{\partial E^S \, \partial E^P}\right|_0 = \pi_S^m \pi_P^h + \pi_S^h \pi_P^m = \kappa.$$

Moreover, $\partial\Lambda/\partial E^S\big|_0 = \partial\Lambda/\partial E^P\big|_0 = 0$, so by the chain rule, $s_{SP}|_0 = u'(0)\,\Lambda_{SP}|_0 = u'(0)\kappa$. Continuity of s_{SP} yields positivity on a neighborhood when $\kappa > 0$.

Interpretation. The benchmark assumes that citizens obtain a complete case for endorsement—a persuasive *why* (m) and a persuasive *how* (h)—from within the same channel. Specialization breaks this: if one channel is relatively better at supplying implementability and the other at supplying normative justification, then complete bundles may be assembled *across* channels. The term $\Lambda_\times$ captures precisely this possibility: $f(Q_S^h)f(Q_P^m)$ can be interpreted as the probability of encountering an implementability argument via the state-agent channel *and* a merit argument via the peer channel, while $f(Q_S^m)f(Q_P^h)$ captures the opposite pairing. Under strong specialization ($Q_S^m \simeq 0$, $Q_P^h \simeq 0$), most complete bundles arrive through cross-channel matching rather than within-channel bundling.

The parameter

$$\kappa \equiv \pi_S^m \pi_P^h + \pi_S^h \pi_P^m$$

summarizes the scope for cross-channel assembly: it is large when citizens can obtain merit from one channel and implementability from the other. Proposition 2 shows that, at low exposure, allowing such cross-channel assembly induces local complementarity in reach. Increasing E^S raises the marginal return to E^P because it makes it more likely that a citizen who encounters one side of the justificatory bundle through one channel also encounters the missing side through the other. Under strong specialization, for example when state agents primarily supply implementability and peers primarily supply merit, the off-diagonal product $\pi_S^h \pi_P^m$ dominates, so complete persuasive bundles are assembled mainly across channels rather than within channels.

Specialization is therefore a force multiplier, not a universal obstacle to legitimacy formation. It raises precision along each channel's comparative-advantage dimension: peer exposure may strongly increase perceived merit, while state-agent exposure may strongly increase perceived implementability. But because endorsement is bundle-based, one-dimensional precision is not sufficient. High merit without implementability does not by itself generate strong endorsement, and high implementability without merit does not either. Specialization therefore creates high-dimensional precision together with cross-dimensional incompleteness. Its effect depends on the structure of exposure. In integrated public spheres, where citizens encounter both channels, specialization can accelerate legitimacy formation because audiences receive the full justificatory bundle across channels. In fragmented public spheres, by contrast, specialization can impede consensus because citizens may receive only one amplified argument without the complementary reason needed for endorsement.

Together, the benchmark and the specialization extension generate the paper's central empirical distinction. When channels are internally complete, peer and state-agent exposures substitute for one another because each channel can independently deliver a persuasive (m, h) bundle. When channels specialize, exposure to one channel raises the value of exposure to the other because citizens assemble the justificatory bundle across channels. The sign of the exposure cross-partial therefore reveals the persuasion architecture: substitution indicates within-channel bundling, while complementarity indicates cross-channel fit.

2.2.2 Dynamics: Feedback and cultural change

The previous results characterize within-period persuasion. I now allow endorsement to feed back into future reach through public-sphere propagation, generating norm dynamics.

Let aggregate endorsement at date t be $S_t \in [0,1]$. Public-sphere propagation implies that reach in the next period is increasing in current endorsement:

$$E^j_{t+1} \;=\; E^j(S_t), \qquad j \in \{S,P\}.$$

Given reach (E^S_{t+1}, E^P_{t+1}), the total bundle index is $\Lambda_{t+1} \equiv \Lambda(E^S_{t+1}, E^P_{t+1})$, and aggregate endorsement evolves according to

$$S_{t+1} \;=\; u(\Lambda_{t+1}) \;\equiv\; \Phi(S_t), \qquad \Phi(S) \;=\; u\Big(\Lambda(E^S(S), E^P(S))\Big).$$

This dynamic formulation differs from preference-falsification accounts of rapid norm change. In such accounts, individuals change their public stance when the perceived social acceptability of a position shifts, generating threshold cascades ((Kuran (1991), Kuran (1995), Bursztyn et al. (2020))). A reduced-form representation is

$$y_{i,t} = \mathbf{1}\{\, p_i + \mu_t - c_i(S_{t-1}) + \varepsilon_{i,t} \geq 0 \,\}, \qquad S_t = \mathbb{E}[y_{i,t} \mid \mathcal{F}_t],$$

where p_i is the private stance, $c_i(\cdot)$ decreases in perceived support, and μ_t is an acceptability shock. In this class of models, rapid change operates through the social cost of public expression. By contrast, the mechanism here operates through persuasion: reach determines which arguments citizens encounter, producer roles shape the precision and content of those arguments, and bundling determines whether endorsement changes smoothly or accelerates through cross-channel completion.

Assumption 5 (Exposure maps) *For $j \in \{S,P\}$, $E^j : [0,1] \to [0,\infty)$ is twice differentiable and increasing, with ($E^{j\,\prime} \geq 0$) and bounded derivatives. Exposure may saturate, so $E^{j\,\prime}$ may decline at high endorsement levels.*

Lemma 3 (Curvature decomposition) *For $\Phi(S) = u(\Lambda(S))$,*

$$\Phi'(S) = u'(\Lambda)\,\Lambda'(S), \qquad \Phi''(S) = u''(\Lambda)\,[\Lambda'(S)]^2 \;+\; u'(\Lambda)\,\Lambda''(S).$$

In the specialization extension, $\Lambda''(S)$ contains the strictly positive cross-fit term

$$2\,E^{S\,\prime}(S)\,E^{P\,\prime}(S)\Big[\pi^m_S \pi^h_P\, e^{-\{Q^m_S(S)+Q^h_P(S)\}} \;+\; \pi^h_S \pi^m_P\, e^{-\{Q^h_S(S)+Q^m_P(S)\}}\Big],$$

which is absent in the benchmark without cross-channel fit.

Proof 4 (Sketch) *The first display follows from the chain rule. For the second, write $\Lambda_\times(S) = f(Q^m_S(S))f(Q^h_P(S)) + f(Q^h_S(S))f(Q^m_P(S)), \qquad f(x) = 1 - e^{-x}$. Differentiating twice yields, for each off-diagonal pair, a cross term $2\,f'(Q^\ell_S)\,f'(Q^{\ell'}_P)\,Q^{\ell\,\prime}_S(S)\,Q^{\ell'\,\prime}_P(S)$, with $f'(x) = e^{-x}$ and $Q^{k\,\prime}_j(S) = \pi^k_j E^{j\,\prime}(S)$. Summing over the two off-diagonal pairs gives the stated expression.*

Interpretation. The update map has two sources of curvature. The first is the concavity of u, which captures saturation in endorsement and pushes toward gradual adjustment. The second is the curvature of the bundle index $\Lambda(S)$. Cross-channel fit adds a positive curvature term because exposure to one channel raises the marginal value of exposure to the other: when one channel supplies

the normative *why* and the other supplies the institutional *how*, public-sphere propagation can accelerate endorsement. Specialization can therefore generate nonlinear cultural change. At low or intermediate endorsement levels, cross-channel completion may make the update map locally convex; at high endorsement levels, saturation in exposure and endorsement can restore concavity. This provides the mechanism through which persuasion can generate both smooth cultural change, when bundling is internal to channels or feedback is weak, and rapid cultural shifts, when cross-channel completion interacts with sufficiently strong feedback.

Illustrations. The model also helps organize historical episodes in which endorsement shifts appear nonlinear. These illustrations are not tests of the model, but mappings from the theory's primitives to well-known cases.

Interwar Germany: specialization under fragmented reach. Interwar Germany can be interpreted as an environment in which receptiveness to particularist and exclusionary narratives was already nontrivial, consistent with the long-run intellectual currents emphasized by Kohn (1950) and Mosse (1964). In the model, a bureaucracy with a strong legal-rational tradition (Arendt (1963), Heldring (2026)) can spill over into the public sphere through high implementability precision in the state-agent channel, π_S^h. When this institutional *how* is paired with moral content supplied through peer or party-political discourse, cross-channel fit increases. At the same time, public-sphere reach was fragmented rather than broadly integrative: despite dense associational life, civic engagement often occurred within segmented subpublics, and party politics remained polarized, limiting common arenas in which competing groups encountered shared reasons (Berman (1997)). Mass persuasion technologies, including radio propaganda (Adena et al. (2015)), can be interpreted as raising effective reach in this fragmented environment. The model then suggests a mechanism through which preexisting particularist merit claims, bureaucratic implementability discourse, segmented public-sphere structure, and a late increase in reach can jointly generate rapid endorsement shifts.

The French Revolution: broad reach and within-peer bundling. A contrasting illustration is late–1780s France. Peer discourse already bundled the normative *why* and the institutional *how*: claims about general interest, equality, and representation circulated together with concrete institutional proposals. Contributions such as Condorcet (1785) and Sieyès (1789) illustrate this coupling. In the model, this corresponds to substantial within-channel bundling in the peer channel, with high π_P^m and π_P^h, so that peer-delivered complete bundles can be generated without relying primarily on cross-channel matching. The pamphlet and newspaper boom, together with salons, clubs, and other forms of associational life, increased peer reach E^P by expanding the frequency and scale of public argument exchange (Habermas (1989), Darnton (1982), Popkin (1990)). As the Estates–General convened and parliamentary procedure emerged through the National and Constituent Assemblies, public endorsement became tied to formal political choice. In the model, this is an internalization shock: β_P and β_S rise because producing and relaying public reasons becomes more consequential for institutional outcomes[5]. On this interpretation, revolutionary acceleration reflects broad increases in reach and already-strong within-peer bundling, reinforced by higher incentives to supply precise public reasons once endorsement became politically consequential.

[5]The *cahiers de doléances* were lists of complaints and reform proposals compiled across France in early 1789 by local assemblies of the three estates to instruct their deputies to the Estates–General.

2.2.3 Welfare: amplification risk and guardrails

The positive analysis highlights a welfare trade-off. Because endorsement is bundle-based—it requires both a persuasive *why* (m) and a persuasive *how* (h)—precision investments can improve the informational basis of legitimacy. But under strong role specialization, the state-agent channel may supply mainly implementability arguments while contributing little merit discourse. In that case, implementability can become a content-agnostic amplifier: it increases the institutional realizability of whichever merit claim prevails, rather than disciplining the merit claim itself. This section separates two wedges. The first is an internalization wedge governing the overall supply of public reasons. The second is a composition wedge governing whether implementability precision completes a well-disciplined merit case or amplifies a weakly contested one.

Planner benchmark. Fix a public sphere (c,t) and suppress subscripts. In the decentralized benchmark, channel $j \in \{S,P\}$ chooses (π_j^m, π_j^h) to maximize

$$\max_{\pi_j^m,\pi_j^h \geq 0} \beta_j B u(\Lambda) \; - \; \sum_{k \in m,h} K_j^k(\pi_j^k),$$

where $\beta_j \in (0,1]$ is the share of the legitimacy payoff internalized by channel j. A utilitarian planner instead internalizes the full legitimacy payoff and chooses all precision levels:

$$\max_{\pi_j^m,\pi_{j\,j\in S,P}^h} \; Bu(\Lambda) \; - \; \sum_{j \in S,P} \sum_{k \in m,h} K_j^k(\pi_j^k).$$

The decentralized problem therefore contains a standard under-internalization wedge: when $\beta_j < 1$, the marginal private return to producing public reasons is below the marginal social return. Policies or institutions that raise β_j, or reduce the cost of producing precise and contestable reasons, increase the supply of legitimacy-relevant information.

Amplification risk. The internalization wedge is not the only welfare issue. The model also isolates a composition risk. Under the role-based implementability advantage, state agents face lower marginal costs of producing implementability precision. If this advantage generates an extreme content mix,

$$\pi_S^h \gg \pi_S^m \simeq 0,$$

then state-agent discourse supplies the institutional *how* without supplying much of the normative *why*. In the specialization extension, this creates cross-channel fit when paired with peer-side merit arguments. Proposition 2 shows that, at low exposure, such cross-channel fit can make peer and state-agent reach locally complementary. This is informative but also potentially risky: implementability arguments can amplify whatever merit discourse is available in the peer channel. When peer-side merit discourse is well contested and precise, this complementarity improves legitimacy formation. When peer-side merit discourse is weakly disciplined, implementability precision can instead scale up a one-sided or poorly justified merit claim.

Guardrails. The amplification risk is greatest when three conditions coincide: (i) state-agent discourse is strongly specialized in implementability; (ii) reach is low or fragmented, so complete bundles are scarce; and (iii) merit discourse is weakly contested. The corresponding guardrails map directly into the model's primitives. First, composition guardrails raise π_S^m or prevent extreme h-only specialization in the state-agent channel. Second, merit-discipline guardrails raise the precision and contestability of peer-side merit discourse, by improving media quality, verification capacity, civic

education, and deliberative infrastructure. Third, public-sphere integration guardrails increase cross-cutting reach and common fora, reducing the low-exposure fragmentation that makes cross-channel amplification especially powerful.

The policy implication is therefore not simply to maximize reach or precision. Higher precision is desirable when it improves both dimensions of public justification, but potentially dangerous when it strengthens implementability while leaving merit weakly disciplined. Planner-type levers that raise internalization or relax costs increase the overall supply of public reasons; composition and public-sphere guardrails ensure that additional implementability precision completes a contestable justificatory bundle rather than mechanically amplifying whichever merit claim happens to dominate.

For reference, Appendix Table A1 reports the main notation used throughout the model.

3 From Model to Empirics: Measurement and Identification

This section maps the model to empirical counterparts and lays out the identification strategy. The empirical analysis targets the model's static persuasion predictions: whether peer and state-agent effective precision—the product of exposure and content precision—increases endorsement of institutional universalism, and whether their interaction reveals within-channel bundling or cross-channel fit.

3.1 Empirical counterparts

Consider a canton-year public sphere (c, t). The dependent variable is a binary indicator of endorsement of institutional universalism. Let $Y_{i,c,t} = 1$ if individual i supports equal institutional treatment, including for foreigners, and $Y_{i,c,t} = 0$ if the individual favors preferential opportunities for Swiss citizens. Thus, the outcome measures endorsement of universalism as a rule of public action rather than a purely private distributive preference.

In the model, persuasion operates through two objects: exposure to a channel $j \in \{S, P\}$, denoted E^j, and the content carried by that channel. Empirically, I construct content-weighted exposure measures

$$X^j_{c,t} \equiv E^j_c \cdot p^j_{c,t}, \qquad j \in \{S, P\}, \tag{1}$$

where E^j_c is a predetermined reach proxy for channel j in canton c, and $p^j_{c,t} \in [0, 1]$ is the channel's universalist tilt in (c, t), measured as the share of workers in channel j who endorse institutional universalism. I proxy the state-agent channel with public-sector employees (communal, cantonal, or federal) and the peer channel with private-sector employees. I compute $p^j_{c,t}$ outside the outcome sample. The outcome sample consists of non-workers, so the channel-level tilt captures the universalist content to which respondents are exposed in their canton-year public sphere, rather than mechanically reflecting their own responses. The model distinguishes the normative *why* and the institutional *how*, but the data do not separately observe merit and implementability arguments. The empirical object $p^j_{c,t}$ should therefore be interpreted as the channel's directional universalist content, not as a direct measure of π^m_j or π^h_j. The distinction between within-channel bundling and cross-channel fit is instead recovered from the exposure interaction and from the mechanism tests.

To proxy reach, I use predetermined canton-level stock variables and infrastructure:

$$E^S_c \approx \text{Public-employment share}^{(\text{baseline})}_c, \qquad E^P_c \approx \text{Association density}^{(\text{baseline})}_c. \tag{2}$$

These proxies capture stable differences in opportunities for exposure to state agents and peers: the local public-employment share captures potential contact with state agents, while association

density captures the density of civic interaction and peer diffusion. Appendix Table A2 reports the exact construction.

In the model, exposure is individual-specific: $E^j_{i,c,t}$ scales the precision with which citizen i receives reasons from channel j in public sphere (c,t). Empirically, I observe only predetermined canton-level proxies E^j_c. A natural mapping is

$$E^j_{i,c,t} = E^j_c + \eta^j_{i,c,t}, \qquad \mathbb{E}[\eta^j_{i,c,t} \mid c,t] = 0,$$

where $\eta^j_{i,c,t}$ captures within-canton heterogeneity in contacts and participation. Estimating the exposure specification with E^j_c replaces the individual content-weighted exposure $E^j_{i,c,t}p^j_{c,t}$ by the proxy $E^j_c p^j_{c,t}$. Under a classical proxy condition—that $\eta^j_{i,c,t}$ is mean-zero and orthogonal to the structural error conditional on fixed effects and controls—this substitution introduces measurement error that attenuates the estimated exposure effects. If instead $\eta^j_{i,c,t}$ reflects endogenous contact choice, the proxy may also induce selection bias. The baseline design therefore focuses on non-workers who do not move across cantons, so that variation in content-weighted exposure comes from predetermined canton-level reach interacted with canton-year changes in channel content, rather than from individuals' own mobility or workplace contact choices.

3.2 Baseline specification

I estimate

$$Y_{i,c,t} = \theta_S X^S_{c,t} + \theta_P X^P_{c,t} + \theta_{SP}\, X^S_{c,t} X^P_{c,t} + \mathbf{Z}'_{c,t}\Gamma + \alpha_i + \lambda_t + \delta_c \cdot t + \varepsilon_{i,c,t}, \tag{3}$$

where α_i are individual fixed effects, λ_t are year fixed effects, and $\delta_c \cdot t$ are canton-specific linear trends. The vector $\mathbf{Z}_{c,t}$ includes time-varying canton-level controls that may jointly affect workers' channel-specific universalist tilt and non-workers' endorsement of institutional universalism, such as electoral timing, unemployment, immigration flows, cantonal political ideology, and other canton-year covariates. I include individual time-varying covariates in robustness checks. I demean $X^S_{c,t}$ and $X^P_{c,t}$ in the estimation sample before constructing the interaction, so that θ_P and θ_S capture the marginal effect of one channel's content-weighted exposure evaluated at the mean of the other channel.

Interpretation and link to the model. Equation (3) is a model-disciplined reduced form. The regressors are constructed to match the model's exposure-scaling logic: reach is predetermined at the canton level, while channel content varies over canton-years. The object

$$X^j_{c,t} = E^j_c p^j_{c,t}$$

therefore proxies effective exposure to universalism-relevant content in channel j. The coefficients θ_S and θ_P capture how endorsement responds to marginal increases in state-agent and peer content-weighted exposure. The interaction coefficient θ_{SP} captures whether the two channels are substitutes or complements in generating persuasive (m,h) bundles. In the benchmark with within-channel bundling, each channel can independently supply a persuasive (m,h) bundle. Concavity in endorsement then implies substitutability across channels: the cross-partial is weakly negative absent cross-channel pairing. In the specialization extension, by contrast, citizens can assemble matched (m,h) reasons *across* channels. Cross-channel fit then generates local complementarity at low reach (Proposition 2). The sign of θ_{SP} is therefore informative about the persuasion architecture:

a negative interaction is consistent with within-channel bundling, while a positive interaction is consistent with cross-channel fit.

In the model, persuasion depends on effective precision $Q_j^k = \pi_j^k E^j$ (Lemma 1). Empirically, I observe predetermined canton-level reach proxies E_c^j and channel-specific universalist tilt $p_{c,t}^j$, but not the separate merit and implementability components of precision. The specification therefore uses $X_{c,t}^j = E_c^j p_{c,t}^j$ as the empirical counterpart of effective exposure to universalism-relevant public reasons. The interaction $X_{c,t}^S X_{c,t}^P$ asks whether increases in exposure to both channels jointly raise endorsement more or less than additively, as predicted by the model's distinction between within-channel bundling and cross-channel fit.

Identification. The baseline design isolates persuasion within a fixed local public sphere. Because $X_{c,t}^j$ is defined at the canton-year level, canton-by-year fixed effects would absorb the identifying variation. I therefore use individual fixed effects, year fixed effects, canton-specific linear trends, and rich canton-year controls, and restrict the baseline sample to canton non-movers. This restriction avoids endogenous relocation across cantons and ensures that within-individual variation in content-weighted exposure comes from changes in channel-specific tilt $p_{c,t}^j$ within a predetermined canton-level reach environment, rather than from individuals moving into different exposure environments.

The identifying variation in the baseline specification comes from within-individual changes in

$$X_{c,t}^j = E_c^j p_{c,t}^j,$$

induced by canton-year changes in channel-specific tilt $p_{c,t}^j$, conditional on individual fixed effects, year fixed effects, canton trends, and canton-year covariates. The baseline exposure design is therefore best interpreted as identifying the model's public-sphere persuasion architecture: whether peer and state-agent content-weighted exposures increase endorsement, and whether their interaction is consistent with within-channel bundling or cross-channel fit.

This design does not, by itself, fully rule out residual canton-year shocks correlated with both workers' universalist tilt and non-workers' endorsement. For this reason, I complement it with a workplace-switch design, which delivers within-person shifts in direct exposure to state agents while holding canton-year discourse conditions fixed. The switch design provides the sharper causal test of the effect of state-agent exposure on endorsement.

3.3 Workplace exposure shocks

The baseline exposure design identifies the public-sphere persuasion architecture at the canton-year level. To obtain sharper causal evidence on direct exposure to state agents, I exploit switches between private and public employment. In the model, the state-agent channel is characterized by a content and precision mix (π_S^m, π_S^h), empirically proxied at the canton-year level by $p_{c,t}^S$, measured outside the outcome sample. A worker's sector switch does not mechanically change this canton-year content. What changes is the worker's access to state-agent contacts through the workplace. Entering public employment raises the intensity of day-to-day interactions with public employees, and hence the individual exposure component $E_{i,c,t}^S$; leaving public employment reduces it.

Baseline switcher specification. I estimate the following fixed-effects regression, where identification comes from individuals who switch sector:

$$Y_{i,c,t} = \eta\, State_{it} + \psi\, Fed_{it} + \alpha_i + \delta_{c,t} + \varepsilon_{i,c,t}. \tag{4}$$

The unit of observation is an individual i in canton c and year t. The variable $State_{it}$ equals one if individual i works in the public sector—communal, cantonal, or federal—and zero if the individual works in the private sector, including private nonprofit employment. The variable Fed_{it} equals one if the individual works in the federal administration and zero otherwise. Thus, η captures the effect of moving into non-federal public employment relative to private employment, while ψ captures the incremental effect of federal employment relative to other public employment. The total effect of federal employment relative to private employment is therefore $\eta + \psi$. Individual fixed effects α_i absorb time-invariant selection into sector. Canton-by-year fixed effects $\delta_{c,t}$ absorb the local public-sphere environment and all canton-year shocks common to workers in the same canton and year. Standard errors are clustered at the individual level.

The key advantage of Equation (4) relative to the baseline exposure specification is the inclusion of canton-by-year fixed effects. In the baseline design, the treatment variables $X^j_{c,t}$ vary only at the canton-year level, so canton-by-year fixed effects would absorb the identifying variation. In the workplace-switch design, treatment varies at the individual-year level through $State_{it}$. I can therefore include $\delta_{c,t}$, which holds fixed the entire canton-year public sphere: the local discourse environment, worker universalist tilt $p^j_{c,t}$, local political shocks, labor-market conditions, immigration shocks, and any other canton-year factor common to individuals in the same canton and year. Identification comes from within-person changes before and after a sector switch, benchmarked against stayers observed in the same canton-year environment. Thus, the estimate isolates the effect of direct workplace access to state-agent contacts, rather than variation in canton-year discourse conditions.

By construction, a sector switch changes the composition of daily professional interactions. It generates a discrete shift in contact intensity with public employees relative to residential exposure alone. This workplace-based exposure shock is therefore less subject to attenuation from canton-level reach proxies than the baseline content-weighted exposure specification.

Dynamic self-selection and service-type sorting. Equation (4) addresses time-invariant selection into sector, but a remaining concern is dynamic self-selection: individuals may become more favorable to institutional universalism before entering public employment. Prior work shows that civic-minded agents may sort into caring services—health, education, and social care—that are delivered by both public and private providers (Francois (2003), Gregg et al. (2011)). This implies that selection on endorsement is likely to operate primarily at the service-type level rather than at the institutional-sector level.

I probe this possibility by comparing pre-switch endorsement among future switchers and stayers while individuals are still employed in the private sector. Among private-sector workers, I estimate

$$Y^{Priv}_{i,c,t} = \phi Switch^{pub}_i + \delta_{c,t} + \mathbf{Z}'_{i,t}\gamma + u_{i,t}, \tag{5}$$

where $Y^{Priv}_{i,c,t}$ is endorsement measured while individual i is employed in the private sector. The variable $Switch^{pub}_i$ equals one if individual i switches from private to the public employment at any point in the future. Canton-by-year fixed effects $\delta_{c,t}$ compare eventual switchers to stayers within the same canton-year. The vector Z_{it} includes individual covariates: gender, number of children, age, religion, culture, college education, marital status, health status, and urban residence. I also repeat the exercise excluding federal employment, allowing selection patterns to differ across institutional tiers.

To sharpen the service-type implication, I partition jobs into HES sectors—health, education, and social care—and non-HES sectors. I then re-estimate Equation (5) separately in each private-sector subsample using destination-specific indicators: $Switch^{pub,HES}_i$, equal to one if individual i later switches from private to public HES employment, and $Switch^{pub,non\text{-}HES}_i$, equal to one if individual

i later switches from private to public non-HES employment. Appendix Figure A4 shows that pre-switch differences are confined to service type: future switchers into public HES exhibit higher baseline endorsement, whereas future switchers into public non-HES do not. Accordingly, the causal switch design focuses on private-to-public transitions within non-HES industries.

Finally, I compare baseline characteristics of private non-HES workers who subsequently move to public non-HES with those who start in public non-HES and remain there (Appendix Table A5). Switchers are slightly younger and less likely to hold a college degree. In Switzerland, comparable public jobs pay a premium—on average about 12%, and up to 29% at the lower end of the pay scale (Portmann et al. (2024))—so pecuniary considerations plausibly contribute to switching within non-HES occupations.

4 Data

Swiss Household Panel (SHP). The main individual-level data come from the Swiss Household Panel (SHP), a longitudinal survey conducted annually since 1999 that follows a random sample of approximately 5,000 households, covering over 12,000 individuals residing in Switzerland. I use waves I–V (1999–2003), the only waves that jointly contain the outcome, occupation and employment information, and the full set of covariates used in the baseline exposure specification and the workplace-switch design. The panel is geocoded at the municipality of residence, which I aggregate to the canton-year level to construct the empirical counterparts to local public-sphere environments.

For the baseline exposure design, I restrict the sample to Swiss nationals aged 18 and older who are not working. I focus on Swiss nationals because the institutional-universalism item explicitly asks respondents to trade off opportunities for Swiss citizens versus foreigners. I focus on non-workers because the canton-year tilt of workers in each channel is used to construct the exposure measures; excluding workers from the outcome sample ensures that these measures capture the discourse environment to which respondents are exposed rather than mechanically reflecting their own responses. The baseline sample contains $N = 7,413$ individuals and 20,542 person-year observations (Appendix Table A3).

Institutional universalism. The outcome measures endorsement of institutional universalism: support for equal institutional treatment of socially distant others rather than preferential treatment of the national in-group. The SHP asks respondents: *'Are you in favour of Switzerland offering foreigners the same opportunities as those offered to Swiss citizens, or in favour of Switzerland offering Swiss citizens better opportunities?"* Responses are: (i) 'in favour of equality of opportunities," (ii) 'neither," and (iii) 'in favour of better opportunities for Swiss citizens." I define $Y_{i,c,t} = 1$ for respondents choosing equality of opportunities and $Y_{i,c,t} = 0$ for respondents choosing better opportunities for Swiss citizens. Respondents answering "neither" are excluded from the baseline measure. This category accounts for less than 10% of responses and plausibly pools heterogeneous positions, including ambivalence, non-attitudes, and social-desirability responding.

The item is observed annually from 1999 to 2009 and again in selected later waves (2011, 2014, 2017). Across the full set of waves in which it is fielded, endorsement of institutional universalism is relatively stable (Appendix Figure A1). In the main estimation window (1999–2003), mean endorsement in the full Swiss-adult sample is 0.657; equivalently, 34.3% prefer preferential opportunities for Swiss citizens. In the baseline estimation sample, mean endorsement is 0.624, reflecting the additional sample restrictions.

Figure 1 documents sizeable cross-canton heterogeneity in endorsement, ranging from roughly

51% to 73% across cantons. This dispersion supports the use of the canton-year as the empirical boundary of the local public sphere. This level is substantively meaningful in Switzerland: cantons retain broad sovereignty over domains not delegated to the Confederation, and Swiss federalism allocates many public tasks according to subsidiarity. Cantons also have wide-ranging competences in policy domains such as education, health, culture, economic policy, transport, and local government. Thus, the canton is not merely a geographic aggregation unit; it is a political, administrative, and public-discourse environment in which rules are debated, implemented, and locally interpreted.

Figure 1: Endorsement of institutional universalism averaged at the state (cantonal) level

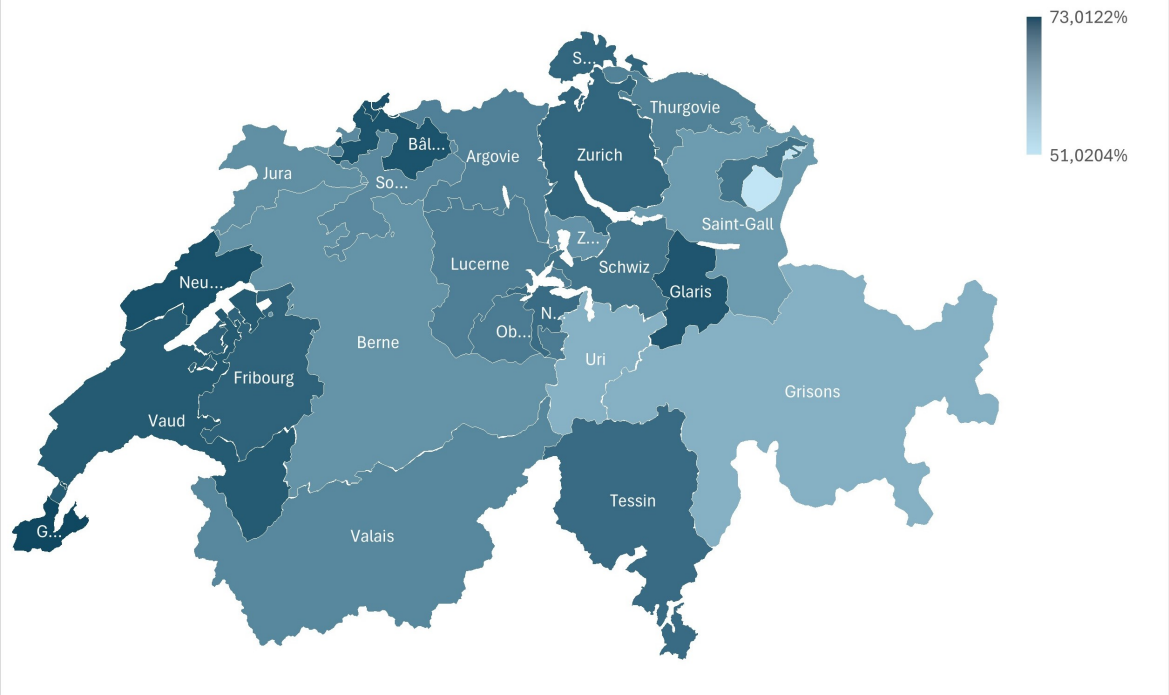

Note: The map displays the average share of Swiss individuals aged 18+ in each canton who endorse the universalism institutional norm. This corresponds to the cantonal average of the variable Y_{ict} over the period 1999–2017.

Individual-level correlates. Appendix Figure A2 reports OLS associations between $Y_{i,c,t}$ and individual covariates, including education, ideology, language, urbanicity, family structure, religiosity, and income. Endorsement is higher among college-educated, left-leaning, urban, higher-income, and French-speaking respondents, and lower among more religious respondents and those with more children. These correlations are descriptive and characterize the sample. The main specifications rely on within-individual variation. Importantly for the individual fixed-effects design, endorsement exhibits substantial within-person variation: about one quarter of respondents switch their stance at least once over 1999–2003 (Appendix Table A3).

Institutional sectors. The SHP records respondents' institutional sector through the question: *"Are you employed by a private company or a state organization?"* I define $State_{it} = 1$ if respondent i reports being employed by a state organization—communal, cantonal, or federal—and $State_{it} = 0$ if employed by a private organization, including non-profit employment. I exclude respondents

employed by international organizations, which are not part of the domestic state apparatus. I separately code federal employment using $Fed_{it} = 1$ for confederation-level employment and $Fed_{it} = 0$ otherwise. The SHP also reports industry affiliation. I partition jobs into HES industries—health, education, and social care—and non-HES industries using the economic-activity classification. This split distinguishes mission-oriented public-service industries from other activities and is used both for descriptive comparisons and to sharpen interpretation in the workplace-switch design.

Appendix Table A4 reports descriptive endorsement rates by institutional sector. Public-sector employees are more likely to endorse institutional universalism than private-sector employees, 74% versus 66%. This gap is present both within HES and within non-HES industries, suggesting that it is not confined to traditionally mission-oriented fields. Additional descriptives also reveal heterogeneity within the public sector: endorsement is higher among cantonal and communal employees than among federal employees, 78% versus 67%, respectively. Cantonal and communal employees exercise public authority within subnational institutional environments, whereas federal employees occupy roles within the national administration. I therefore allow federal employment to have a distinct effect in the workplace-switch design.

Appendix Figure A3 plots endorsement trends separately for public- and private-sector employees. The two series follow distinct trajectories and are only weakly correlated at the canton-year level ($\rho = 0.32$). This descriptive pattern supports treating the two channels separately in the construction of the canton-year content environments, p^S_{ct} and p^P_{ct}.

Switch-design sample. For the workplace-exposure design, I restrict the sample to employees aged 20–54, to focus on voluntary job-to-job mobility, and exclude the self-employed. Over 1999–2003, this worker panel contains 3,777 individuals and 9,849 person-year observations. I further partition jobs into HES and non-HES industries and emphasize switches within non-HES sectors. Among non-HES workers, 8.6% switch institutional sector during the observation window. Switches occur in both directions in roughly equal numbers, and no individual switches more than once. These mobility patterns are consistent with the relative permeability of Swiss public employment: recruitment is typically open, without competitive entrance examinations, and formal civil-service status was abolished in all but two cantons during the 1990s (Audier and Bacache-Beauvallet (2007)).

Appendix Table A2 provides definitions, construction details, and data sources for all variables used throughout the empirical analysis, while Appendix Table A3 reports their summary statistics.

5 Empirical Results

5.1 Baseline estimates

Table 1 reports estimates of the baseline content-weighted exposure specification in Equation (3). The sample consists of non-working Swiss respondents who do not move across cantons during the estimation window[6]. This design is intended to identify the architecture of public-sphere persuasion across canton-year environments: whether peer and state-agent content-weighted exposures are associated with endorsement, and whether their interaction is consistent with within-channel bundling or cross-channel fit.

[6]As described in Section 3.1, channel exposures are constructed as $X^j_{c,t} = E^j_c \, p^j_{c,t}$, using predetermined canton-level reach proxies E^j_c and canton–year variation in the channel discourse environment $p^j_{c,t}$, computed outside the outcome sample. Results are robust to including canton movers (Appendix Table A6).

Because $X^P_{c,t}$ and $X^S_{c,t}$ are centered in the estimation sample, the main-effect coefficients θ_P and θ_S can be read as marginal effects evaluated at the mean exposure of the other channel. Differentiating Equation (3) yields

$$\frac{\partial Y_{i,c,t}}{\partial X^P_{c,t}} = \theta_P + \theta_{SP} X^S_{c,t}, \qquad \frac{\partial Y_{i,c,t}}{\partial X^S_{c,t}} = \theta_S + \theta_{SP} X^P_{c,t}. \tag{6}$$

Thus, θ_{SP} captures how the marginal effect of one channel varies with exposure to the other channel. A positive interaction indicates complementarity, as predicted when citizens assemble matched ((m,h)) reasons across channels. A negative interaction indicates substitutability, as predicted when each channel supplies relatively complete bundles and marginal persuasive returns decline with exposure to the other channel.

All specifications include individual fixed effects and year fixed effects. Because the baseline sample is restricted to canton non-movers, individual fixed effects absorb time-invariant canton differences. Column (1) presents a parsimonious specification without time-varying canton controls or canton trends. Column (2) adds contemporaneous canton-year controls to absorb observable canton-year factors that could jointly affect workers' channel-specific universalist tilt and non-workers' endorsement, such as electoral timing, unemployment, immigration flows, cantonal political ideology, and other canton-year covariates. Column (3) is the preferred specification: it further adds canton-specific linear trends, absorbing gradual canton-level changes in political culture, demographic composition, or other unobservables that could correlate with trends in $p^j_{c,t}$. Column (4) adds time-varying individual controls as a robustness check. These controls are introduced to verify that the estimates are not driven by observable life-cycle changes not absorbed by individual fixed effects. Column (5) replaces contemporaneous canton controls with their one-period lags, while keeping the electoral calendar contemporaneous, to address the concern that contemporaneous canton covariates may themselves respond to shocks that also shift the public-sphere discourse environment.

Across specifications, peer and state-agent content-weighted exposures are positively associated with endorsement of institutional universalism. Since $Y_{i,c,t}$ is binary and Equation (3) is estimated as a linear probability model, coefficients can be read in percentage points. In the preferred specification, Column (3), a one-standard-deviation increase in peer content-weighted exposure, $\text{sd}(X^P_{c,t}) = 0.165$, increases endorsement by about 5.5 percentage points, evaluated at the mean state-agent exposure. A one-standard-deviation increase in state-agent content-weighted exposure, $\text{sd}(X^S_{c,t}) = 1.351$, increases endorsement by about 4.4 percentage points, evaluated at the mean peer exposure. Relative to a baseline endorsement rate of about 0.62, these correspond to increases of roughly 8.9% and 7.1%, respectively. The positive state-agent coefficient is important for the model. It suggests that state-agent discourse carries universalism-relevant persuasive content distinct from peer discourse, even conditional on individual fixed effects, year fixed effects, canton trends, and observable canton-year controls. However, because the baseline design relies on canton-level reach proxies, these estimates should be interpreted as evidence on the public-sphere persuasion architecture rather than as the sharpest causal estimate of direct state-agent exposure. The workplace-switch design below provides that complementary causal test by exploiting individual-level changes in access to state agents while holding the canton-year discourse environment fixed.

The interaction term is negative across specifications. In the preferred specification, a one-standard-deviation increase in state-agent exposure reduces the effect of a one-standard-deviation increase in peer exposure by about 1.6 percentage points. Equivalently, the peer effect falls from 5.5 percentage points at mean state-agent exposure to about 3.9 percentage points when state-agent exposure is one standard deviation above its mean. The same calculation applies symmetrically: a one-standard-deviation increase in peer exposure reduces the effect of a one-standard-deviation increase in state-agent exposure from about 4.4 to 2.8 percentage points. Thus, the interaction is

not only statistically negative but economically meaningful: the marginal effect of each channel is substantially smaller when the other channel is already strong. This substitution pattern is consistent with the benchmark bundling model, in which peer and state-agent channels carry partly overlapping bundles of normative *why* and institutional *how* arguments. The result contrasts with a purely additive diffusion process: exposure matters not only because individuals encounter more universalist agents, but because the persuasive content carried by the two channels overlaps. Once one channel is strong, the marginal persuasive value of the other declines.

The estimates are stable when adding time-varying individual controls in Column (4), supporting the view that the baseline variation is not driven by observable life-cycle changes. The content-weighted exposure coefficients also remain qualitatively similar when contemporaneous canton controls are replaced by lagged canton controls in Column (5), and the negative interaction becomes larger in magnitude. This suggests that the level effects and substitution pattern are not artifacts of the particular timing of the canton controls.

Table 1: Exposure to deliberation channels and endorsement of institutional universalism

	(1)	(2)	(3)	(4)	(5)
Y_{ict}	Endorsement of institutional universalism				
Peer c.w. exposure (X^P_{ct})	0.239***	0.231**	0.333***	0.331***	0.300**
	(0.0735)	(0.0850)	(0.0915)	(0.094)	(0.123)
State-agent c.w. exposure (X^S_{ct})	0.025*	0.027**	0.0326**	0.0325**	0.041**
	(0.012)	(0.011)	(0.0125)	(0.0123)	(0.0149)
Peer × state-agent ($X^P_{ct} \times X^S_{ct}$)	-0.0395*	-0.0352*	-0.0702**	-0.0685**	-0.108***
	(0.0224)	(0.0191)	(0.0265)	(0.0268)	(0.0345)
Observations	8,426	8,349	8,349	8,341	8,349
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Contemporaneous canton controls	No	Yes	Yes	Yes	No
Canton-specific linear trends	No	No	Yes	Yes	Yes
Time-varying individual controls	No	No	No	Yes	No
Lagged canton controls					Yes

Notes: The dependent variable Y_{ict} is a binary indicator equal to one if respondent i endorses institutional universalism in canton c and year t. Sample: Swiss respondents aged 18+; non-workers; canton non-movers. All specifications estimate equation (3) by OLS with individual, canton, and year fixed effects; standard errors are clustered at the canton level. The peer and state-agent content-weighted (c.w.) exposures, X^P_{ct} and X^S_{ct}, are mean-centered in the estimation sample so that the main-effect coefficients are evaluated at the mean of the other channel's exposure. *Contemporaneous canton controls* include: (i) the weighted-average left–right ideology of the cantonal executive (constructed by mapping parties' cabinet shares onto a 1–4 scale using party positions from Vatter et al. (2024)); (ii) foreign population share; (iii) number of referenda; (iv) unemployment rate; and (v) taxable income per capita (cantonal average). *Lagged canton controls* replace (ii)–(v) by their one-year lag (while keeping the electoral calendar contemporaneous). *Time-varying individual controls* include marital-status indicators, number of children, and an indicator for health limitations in daily activities. Column (3) is the preferred specification; Columns (4)–(5) provide robustness to adding individual controls and to using lagged canton covariates. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2 Workplace exposure shocks: Switcher estimates

Table 2 estimates Equation (4) in a linear probability model with individual fixed effects and canton-by-year fixed effects. This is the key difference with the baseline exposure design. Because the treatment varies at the individual-year level through sector switches, I can include $\delta_{c,t}$, which absorbs the entire canton-year public-sphere environment: channel-specific discourse $p^S_{c,t}$ and $p^P_{c,t}$, local political shocks, labor-market conditions, immigration shocks, campaigns, and any other canton-year factor common to individuals in the same canton and year. Identification therefore

comes from within-person changes before and after a sector switch, benchmarked against stayers observed in the same canton-year environment[7].

Column (1) uses all observed switches between private and public employment. The coefficient on $State_{it}$, $\hat{\eta}$, is positive and statistically significant: entering public employment increases the probability of endorsing institutional universalism. Since canton-by-year fixed effects hold fixed the local discourse environment, this effect is not driven by changes in the canton-year universalist tilt of public employees. It reflects a change in individual access to state-agent contacts through the workplace. Columns (2)–(4) progressively tighten the identifying comparison. Column (2) restricts attention to switches that occur within the same service type, HES versus non-HES, addressing selection into caring-service occupations that may operate in both sectors. Column (3) focuses on private-to-public switches. Column (4) further restricts to private-to-public switches within non-HES jobs, which is the preferred causal estimate because pre-switch comparisons show no differential baseline endorsement for future switchers into public non-HES employment. Across Columns (1)–(4), the point estimate increases as the comparison set becomes more tightly aligned with the identifying variation, consistent with attenuation from compositional heterogeneity in broader switch samples.

The preferred estimate in Column (4) implies that moving from private to public employment within non-HES industries raises the probability of endorsing institutional universalism by about 9.9 percentage points. Relative to a mean endorsement rate of 0.65 among private non-HES workers, this corresponds to an increase of about 15%. Appendix Figure A5 shows that the estimate is not driven by any single canton. This magnitude is informative in light of the model. The workplace switch changes the individual exposure component $E^S_{i,c,t}$ while holding the canton-year content environment fixed. It therefore provides direct evidence that increased access to state-agent contacts shifts endorsement of institutional universalism. In the model's language, the result is consistent with state agents carrying persuasive public reasons that make institutional universalism more likely to be recognized as a justified principle of public action.

The magnitude is also consistent with attenuation in the baseline exposure design. In Column (3) of Table 1, a one-standard-deviation increase in peer content-weighted exposure raises endorsement by about 5.5 percentage points, while a one-standard-deviation increase in state-agent content-weighted exposure raises endorsement by about 4.4 percentage points, both evaluated at the mean of the other channel's exposure. The switch estimate of 9.9 percentage points is roughly 2.25 times the one-standard-deviation state-agent exposure effect. This difference is expected: the baseline design relies on canton-level reach proxies that capture diffuse exposure and are likely measured with error, whereas the switch design uses a discrete individual-level change in workplace access to state agents. The larger switch estimate therefore supports the interpretation that canton-level exposure estimates are attenuated relative to direct changes in individual contact intensity.

The coefficient on Fed_{it} is negative, implying that the effect of switching into federal public employment is smaller than the effect of switching into cantonal or communal public employment. Quantitatively, the implied net effect of federal employment relative to private employment, $\eta + \psi$, is close to zero. This tier pattern is informative about the relevant public-sphere boundary. The baseline exposure design defines public-sphere environments at the canton-year level. The switch estimates show that the state-agent effect is concentrated among cantonal and communal employees, whose institutional roles operate at that same level of public authority. By contrast, entering federal employment generates little additional effect once the canton-year discourse environment is held fixed. This supports the interpretation that the relevant persuasive state-agent channel is not generic public employment status, but exposure to state agents whose institutional role is

[7]The results are robust to using a random-effects logit model (available upon request).

articulated within the relevant public sphere. In the model's terms, state agents operating at the relevant level of public authority are more likely to bundle the merit of a candidate norm with the concrete rules, procedures, and administrative means through which that norm can be publicly justified and institutionally implemented. The tier heterogeneity therefore reinforces the choice of the canton-year as the empirical boundary of the public sphere in this setting.

Column (5) considers the reverse transition, from public to private employment, within non-HES industries. The estimated change upon moving to the private sector is small and imprecise. This pattern is consistent with persistence under learning: if workplace exposure shifts beliefs, the posterior need not immediately revert when exposure falls, absent offsetting new information. Column (6) probes whether the estimated effect of state-agent exposure is instead driven by observable features of public jobs. I augment the preferred private-to-public non-HES specification with job-attribute controls: log income, perceived workplace atmosphere, and job security. The public-employment estimate remains similar in magnitude, while the added job attributes are close to zero. This suggests that the switch effect is not simply picking up contemporaneous changes in pay, job security, or reported job quality.

Taken together with the baseline content-weighted exposure regressions, the switch design provides complementary evidence on the model's persuasion mechanism. The baseline estimates identify the architecture of public-sphere persuasion across canton-year environments, while the switch estimates show that a discrete increase in individual access to state agents causally raises endorsement of institutional universalism, holding the local discourse environment fixed.

6 Mechanisms: Argumentative Content and Bundle Formation

The results so far show that content-weighted exposure predicts endorsement of institutional universalism, and that direct workplace access to state agents causally increases endorsement while holding the canton-year discourse environment fixed. I now examine the mechanism. The model implies that persuasion should not operate as generic ideological movement or simple diffusion from more universalist agents. It should instead operate through public reasons: arguments that connect the merit of a candidate norm to the institutional means through which it can be justified and implemented. I distinguish two mechanism layers. First, I use adjacent outcomes to identify the argumentative content accompanying persuasion. If exposure changes endorsement of institutional universalism together with specific policy preferences, but not with broad ideological orientation or generalized trust, this reveals which arguments make the norm persuasive as a rule of public action. Second, I leverage predetermined cantonal environments that shift whether merit and implementability travel together within channels or must be assembled across channels. These tests identify the bundle-based mechanism: whether persuasion works through within-channel bundling or cross-channel assembly of matched (m, h) reasons.

6.1 Argumentative content mechanism tests

I re-estimate the baseline exposure and workplace-switch specifications on a set of adjacent societal and political attitudes: welfare spending, redistribution through taxation of high incomes, support for a strong Swiss army, support for environmental protection over economic growth, support for European Union membership, self-reported left–right ideology, satisfaction with democracy, and trust in the federal government. Appendix Table A2 reports exact variable definitions. These outcomes are used to identify the content of persuasion: whether exposure to peer and state-agent channels changes endorsement of institutional universalism through a structured argumentative bundle rather than through generalized political movement.

Table 2: Shift in exposure to state agents and endorsement of institutional universalism

Y_{ict}	(1) All switches	(2) All switches within non-HES & HES	(3) Private to public within non-HES & HES	(4) Private to public within non-HES	(5) Public to private within non-HES	(6) Private to public within non-HES	(7) Private to public within non-HES
$State_{it}$	0.0458**	0.0526**	0.0838**	0.0987**		0.103*	0.067
	(0.0202)	(0.0266)	(0.0340)	(0.0488)		(0.056)	(0.055)
Fed_{it}	-0.0512*	-0.0477	-0.114**	-0.122*		-0.133*	-0.144**
	(0.0271)	(0.0348)	(0.0549)	(0.0687)		(0.078)	(0.069)
Private sector					-0.0274		
					(0.0427)		
Ln yearly income						0.002	
						(0.025)	
Job atmosphere						-0.003	
						(0.006)	
Job security						0.004	
						(0.011)	
$State_{it} \times French_i$							0.127**
							(0.063)
Observations	8,672	4,709	6,343	4,531	1,084	3,906	4,531
Number of individuals	3,590	2,222	2,794	1,960	441	1,805	1,960

Notes: The dependent variable is an indicator equal to one if respondent i endorses institutional universalism in year t. $State_{it}$ equals one when i is employed in the public sector (communal/cantonal or federal) and zero when employed in the private sector (for-profit or non-profit). Fed_{it} equals one when i is employed by the federal administration. The coefficient on $State_{it}$ therefore compares non-federal public employment to private employment; $State{:}Federal_{it}$ is the incremental effect of federal (relative to non-federal) public employment. All specifications include individual fixed effects and canton×year fixed effects. Standard errors are clustered at the individual level. Column (5) restricts to public-to-private switchers within non-HES and uses a private-employment indicator; the reported coefficient therefore captures the within-person change upon moving from public to private. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

Policy positions: baseline content-weighted exposure. Table 3 re-estimates Equation (3) using the above-mentioned adjacent political and societal attitudes as dependent variables. Across outcomes, the coefficients on peer exposure, state-agent exposure, and their interaction are generally small and statistically indistinguishable from zero. In particular, content-weighted exposure does not systematically predict left–right ideology, support for European Union membership, support for a strong army, environmental attitudes, satisfaction with democracy, or trust in the federal government. This pattern weighs against the interpretation that the baseline results reflect broad ideological realignment, generic political salience, or a generalized increase in institutional trust.

Two exceptions are diagnostically important. First, peer exposure significantly reduces support for welfare spending. This result is counterintuitive under a simple pro-social, egalitarian, or left-wing preference-shift interpretation: the same peer exposure that predicts stronger endorsement of institutional universalism predicts lower support for expansive public provision. The pattern is instead consistent with the model's mechanism. Peer discourse appears to make equal institutional treatment compatible with a less expansive view of public provision. In the model's terms, the peer channel does not merely increase exposure to universalist agents; it carries a particular argumentative content about how universalism should be institutionally realized.

Second, the peer–state-agent interaction is positive for satisfaction with democracy, although this result is only marginally significant. The pattern is nevertheless consistent with the bundle-based mechanism: when peer and state-agent discourse jointly circulate, democratic institutions may appear more compatible with the realization of universalistic norms. Importantly, this effect does not extend to trust in the federal government, which suggests that the mechanism is not a generalized increase in institutional trust.

Table 3: Exposure to deliberation channels and other societal preferences

	(1) Welfare	(2) Redistri.	(3) Army	(4) Env. vs growth	(5) Join EU	(6) Ideology	(7) Sat. Dem.	(8) Trust fed. gov.
X^P_{ct}	-0.275**	0.184	0.162	0.0711	0.00605	0.147	0.330	0.294
	(0.105)	(0.161)	(0.128)	(0.145)	(0.0971)	(0.297)	(0.403)	(0.382)
X^S_{ct}	-0.00504	0.00720	-0.00787	0.00838	0.00523	-0.00975	-0.00277	-0.0261
	(0.0161)	(0.0192)	(0.0190)	(0.0142)	(0.0137)	(0.0456)	(0.0442)	(0.0362)
$X^P_{ct} \times X^S_{ct}$	0.0690	-0.0156	-0.0376	0.00938	-0.00852	0.0843	0.196*	0.0184
	(0.0436)	(0.0369)	(0.0462)	(0.0502)	(0.0257)	(0.113)	(0.0983)	(0.0946)
Observations	9,494	9,604	9,732	9,707	8,497	8,528	9,781	9,819

Notes: Mechanism tests are estimated on the same outcome sample as the baseline specification: Swiss respondents aged 18+; non-workers; canton non-movers. The peer and state-agent content-weighted (c.w.) exposures X^P_{ct} and X^S_{ct} are mean-centered in the estimation sample so that main effects are evaluated at the mean of the other channel's exposure. OLS models with individual, canton, and year fixed effects, canton-specific linear trends, time-varying canton controls, and clustering at the canton level, as in column (3) of Table 1. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Policy positions: workplace switch. Table 4 re-estimates the preferred workplace-switch specification on the same adjacent political and societal outcomes as in Table 3. Across welfare spending, redistribution, environmental preferences, support for a strong army, support for European Union membership, political ideology, and trust in the federal government, the estimated effects of switching into public employment are small and statistically indistinguishable from zero. Thus, the switch effect documented above does not appear to reflect a broad ideological shift, a generalized increase in trust in the state, or a change in standard left–right policy preferences.

The only systematic movement is for satisfaction with democracy. Switching into non-federal public employment is associated with a lower level of satisfaction with democracy, although the

coefficient is only marginally significant. It is consistent with a content-specific mechanism in which exposure to state agents increases endorsement of institutional universalism while also making respondents more aware of the gap between universalistic rules and existing democratic performance. In other words, state-agent exposure does not appear to induce generalized institutional loyalty. If anything, it combines stronger endorsement of equal institutional treatment with a more critical assessment of how democratic institutions currently perform.

Table 4: Workplace exposure to state agents and adjacent political attitudes

	(1) Welfare	(2) Redistri.	(3) Env. vs growth	(4) Army	(5) Join EU	(6) Ideology	(7) Trust fed. gov.	(8) Sat. Dem.
$State_{it}$	-0.105	0.0289	-0.0357	-0.0486	0.078	-0.0905	-0.300	-0.380*
	(0.107)	(0.122)	(0.0803)	(0.115)	(0.077)	(0.209)	(0.263)	(0.226)
Fed_{it}	0.167	-0.0518	0.183	0.127	-0.003	0.117	0.234	0.191
	(0.171)	(0.153)	(0.116)	(0.139)	(0.089)	(0.337)	(0.326)	(0.311)
Observations	4,980	5,036	2,988	5,031	4,568	4,451	5,091	5,046
Number of individuals	2,030	2,040	1,559	2,036	1,951	1,898	2,052	2,039

Notes: The sample is restricted to private-to-public switchers and stayers in non-HES industries, as in the preferred switch specification. All specifications include individual fixed effects and canton-by-year fixed effects. Standard errors are clustered at the individual level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2 Bundle-based persuasion mechanism tests

The second mechanism layer tests the model's bundle-based implication. The baseline and switch estimates show that exposure to peer and state-agent channels affects endorsement of institutional universalism. The first mechanism layer showed that exposure does not operate through broad ideological movement, generic political salience, or generalized institutional trust. Instead, adjacent outcomes reveal content-specific patterns consistent with persuasion through public reasons. The second mechanism layer now tests the model's bundle-based implication. The question is whether this persuasion operates through internally bundled arguments—where a channel jointly conveys the normative *why* and institutional *how*—or through cross-channel assembly, where citizens combine merit arguments from one channel with implementability arguments from another.

Because the argumentative content circulating within each channel is not directly observed, I use predetermined cantonal environments that should shift the persuasion architecture ex ante. The empirical strategy is not to claim that these environments directly measure merit or implementability. Rather, I use theoretically and historically grounded shifters that alter whether merit and implementability are more likely to travel together within channels or to be assembled across channels. I then ask whether these shifters moderate the exposure effects in the directions predicted by the model. A key credibility check is cross-design coherence. Environments that move the persuasion architecture toward within-channel bundling should generate a stronger direct effect of state-agent exposure in the workplace-switch design, because state-agent contact carries a more complete persuasive bundle. The same environments should generate weaker cross-channel completion in the baseline content-weighted exposure design, because peer and state-agent channels become more substitutable. Conversely, environments that increase specialization should generate smaller direct state-agent effects but stronger peer–state-agent complementarity. This cross-design logic is useful because the two designs rely on different sources of variation and face distinct threats to identification.

I test the mechanism shifters using the same heterogeneity logic. For a predetermined shifter

W_c, I estimate

$$Y_{ict} = (\theta_P + \beta_{P,W} W_c) X_{ct}^P + (\theta_S + \beta_{S,W} W_c) X_{ct}^S + (\theta_{SP} + \beta_{SP,W} W_c) X_{ct}^P X_{ct}^S + \mathbf{Z}_{ct}' \Gamma + \alpha_i + \lambda_t + \delta_c \cdot t + \varepsilon_{ict}.$$

The coefficients $\beta_{P,W}$ and $\beta_{S,W}$ capture whether the shifter changes the marginal effects of peer and state-agent exposure, while $\beta_{SP,W}$ captures whether it shifts the peer–state-agent interaction toward substitution or complementarity.

I also estimate the corresponding workplace-switch specification,

$$Y_{ict} = \eta \, State_{it} + \kappa \, (State_{it} \times W_i) + \psi \, Fed_{it} + \alpha_i + \delta_{ct} + \varepsilon_{ict},$$

where W_i denotes the relevant time-invariant shifter in the workplace-switch design. The coefficient κ tests whether the same direct increase in workplace access to state agents has a different effect under the corresponding persuasion architecture.

Church–state separation and the public sphere–state link. The first shifter is cantonal church–state separation. The theoretical logic is that church–state separation changes how institutional norms are publicly justified. In a sacral or ecclesiastically grounded political order, common norms can appear as inherited features of an already-given moral order. By contrast, in a secular political order, common institutions are less directly grounded in sacred or ecclesiastical authority. Their legitimacy must instead be articulated within an immanent public order, through reasons that can be publicly exchanged and institutionally realized (Taylor (2004), Taylor (2007)). Historical accounts of moments in which monarchy, Church, and inherited authority weakened make the same point: once fixed external structures lose their taken-for-granted authority, previously unthinkable normative claims become contestable within public discourse (Morrill (1993), Hill (1972), Acemoglu and Robinson (2025)).

This implication is directly connected to the model. Church–state separation does not make the state less norm-bearing. Rather, it changes the basis on which institutional norms must be defended. Where the common institutional order is less inherited as a sacred or ecclesiastical order, the normative *why* of a rule of public action must be articulated more explicitly in the public sphere. At the same time, because the norm is at issue as a rule to be enacted through public institutions, justification also requires an account of the institutional *how*. Church–state separation should therefore increase the supply of complete justificatory bundles within channels: a given peer or state-agent discourse is more likely to connect why the norm is justified to how it can be institutionally realized.

I operationalize this variation using cantonal church-tax regimes. The relevant distinction is not simply whether a church tax exists, but whether payment is legally compulsory once the tax or contribution is levied. In most Swiss cantons, church taxes are compulsory and can be collected through ordinary enforcement procedures. Three cantons depart from this rule: Geneva, Neuchâtel, and Ticino. In these cantons, taxpayers cannot be coercively compelled to pay the church tax or contribution[8]. I code these cantons as having stronger church–state separation; they account for approximately 11% of observations in the baseline estimation sample. This measure captures whether ecclesiastical authority remains formally backed by the coercive fiscal authority of the cantonal state.

[8]Ticino is a hybrid case: taxpayers are treated as liable unless they formally request exemption from the church-tax register within the prescribed time limits. I code Ticino as a secular canton because the decisive feature for the mechanism is the absence of coercive collection once exemption is claimed. Available historical legal versions indicate that the non-compulsory-payment rule was already in force at least by 1946 in Geneva, 1960 in Neuchâtel, and 1992 in Ticino. Since earlier legal versions were not available in the cantonal legal databases consulted, these dates should be interpreted as conservative lower bounds rather than definitive dates of origin.

In the baseline content-weighted exposure design, $W_c = 1$ for these secular cantons and 0 otherwise. In the workplace-switch design, $W_i = 1$ for respondents living in a secular canton and 0 for respondents living in a non-secular canton. The prediction is that stronger church–state separation should move the persuasion architecture toward within-channel bundling. In the baseline content-weighted exposure design, this implies stronger main effects of peer and state-agent exposure and a more negative peer–state-agent interaction. In the workplace-switch design, it implies a larger direct effect of switching into public employment, because state-agent exposure carries a more complete persuasive bundle.

Figure 2 summarizes the cross-design evidence for the church–state separation shifter. In the content-weighted exposure design, secular cantons display stronger peer and state-agent exposure effects and a more negative peer–state-agent interaction, consistent with stronger within-channel bundling. In the workplace-switch design, the public-employment effect is also larger in secular cantons. The point estimates imply an effect of about 9.6 percentage points in non-secular cantons and about 30.3 percentage points in secular cantons. The difference is imprecisely estimated but sizeable, and its sign and magnitude are consistent with the same mechanism. Church–state separation makes it more likely that both channels connect the normative *why* of universalism to the institutional *how* of public realization. As channels become more internally complete, endorsement relies less on assembling matched *why* and *how* reasons across channels, reducing the marginal role of cross-channel fit.

Figure 2: State secularity and bundle-based persuasion

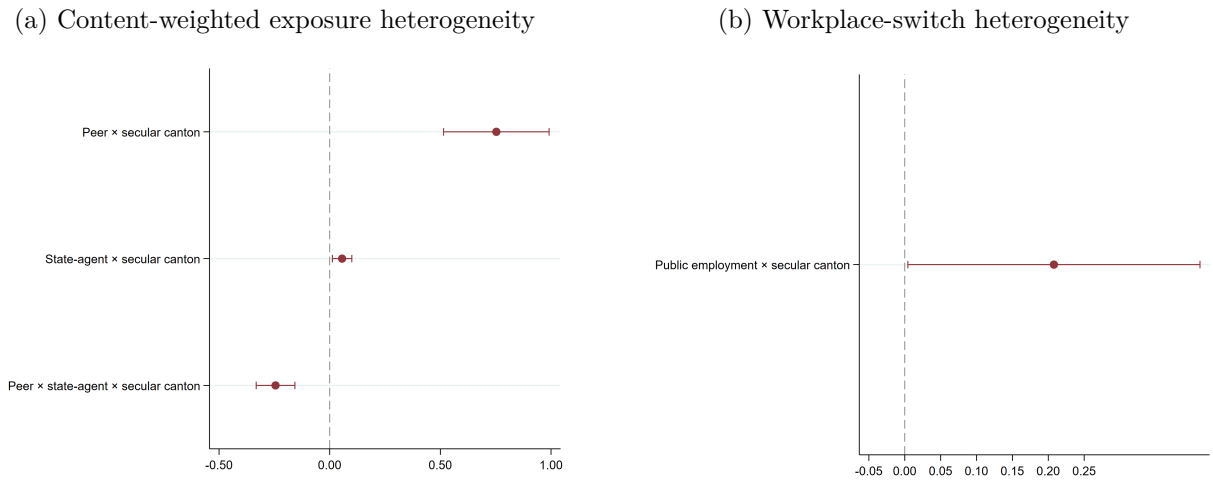

Notes: The figure reports heterogeneity estimates for the state secularity mechanism. Panel A plots coefficients from the baseline content-weighted exposure specification interacted with a secular canton indicator. Panel B plots heterogeneity in the preferred workplace-switch design, estimated among private-to-public switchers and stayers in non-HES industries by interacting public employment with residence in a secular canton. The workplace-switch specification includes individual fixed effects and canton-by-year fixed effects, so the heterogeneity is identified holding fixed the contemporaneous canton-year environment. Horizontal bars report 90% confidence intervals. In Panel A, positive coefficients on peer or state-agent exposure indicate stronger within-channel persuasion in secular states; a negative coefficient on the peer–state-agent interaction indicates greater substitution between channels, consistent with stronger within-channel bundling.

A potential concern is that institutional church–state separation may proxy for secular receiver composition rather than for the public sphere–state link. I address this concern using canton-level

vote shares in the 1980 federal popular initiative on the complete separation of Church and State. The initiative would have introduced a federal constitutional mandate requiring complete church–state separation, the abolition of cantonal church–state arrangements, and the prohibition of ecclesiastical taxes. I use cantonal support for this initiative as a predetermined measure of receiver-side support for state secularity[9]. I then allow this attitudinal measure to shift the exposure–endorsement mapping by interacting it with peer exposure, state-agent exposure, and their interaction. The institutional church–state separation shifter remains predictive of the conditional exposure effects after accounting for receiver-side support for state secularity (Table 5, Column (2)). If anything, the secular-canton interactions increase in magnitude once this receiver-side benchmark is included. This supports the interpretation that the heterogeneity estimates capture differences in how public justification is structured, rather than simply differences in the secular composition of respondents.

Table 5: Church–state separation and receiver-side support for state secularity

	(1)	(2)
Y_{ict}	Endorsement of institutional universalism	
Secular canton $\times X^P_{ct}$	0.753***	0.959***
	(0.139)	(0.207)
Secular canton $\times X^S_{ct}$	0.0560**	0.116***
	(0.0255)	(0.0359)
Secular canton $\times X^P_{ct} \times X^S_{ct}$	-0.245***	-0.402***
	(0.0508)	(0.0607)
Receiver-side secularity interactions	No	Yes
Observations	8,349	8,349

Notes: The secular-canton indicator equals one for Geneva, Neuchâtel, and Ticino, where payment of the church tax or contribution is not coercively enforced by the state. These cantons account for approximately 11% of observations in the baseline estimation sample. Column (1) reports the baseline church–state separation heterogeneity specification. Column (2) allows the exposure–endorsement mapping to vary with receiver-side support for state secularity by interacting the 1980 cantonal vote share in favor of the federal popular initiative on the complete separation of Church and State with X^P_{ct}, X^S_{ct}, and $X^P_{ct} \times X^S_{ct}$. The coefficients reported are the secular-canton interactions with peer exposure, state-agent exposure, and their interaction. All specifications include the fixed effects and controls from the baseline specification. Standard errors are clustered at the canton level.

French republican state legitimacy and merit–implementability bundling. The second shifter is exposure to French republican state-building. Whereas church–state separation shifts the regime of justification by detaching common institutions from inherited ecclesiastical authority, French republican state-building shifts the grammar through which normative claims are connected to institutional realization. In the republican tradition associated with Rousseau, legitimate law is not an external command imposed on individuals; it is binding because citizens can understand it as an expression of collective self-legislation, or the general will (Rousseau (1923)). This tradition gives a distinctive role to public authority: the state is not merely an administrative instrument for implementing independently given norms, but a political form through which collectively authored norms are given binding legal form. French republican state-building gave this grammar institutional form by tying public administration to substantive civic and moral ends (Weber (1976)).

[9]The Federal Council and Federal Assembly recommended rejection primarily on federalist grounds, arguing that church–state relations belonged to cantonal competence and reflected historically diverse cantonal arrangements. The debate therefore illustrates why the relevant public spheres in this setting are defined at the cantonal level: church–state relations were treated as matters of cantonal political authority, historical development, and democratic self-determination.

This logic is directly connected to the model. Exposure to French republican state-building should make the link between normative merit and legal-political realization more salient. Claims about what is right are more likely to be articulated together with claims about how public authority can institutionalize them. In model terms, French exposure shifts the content mix supplied by public-sphere channels toward more complete (m, h) bundles. This should strengthen within-channel persuasion and reduce the need for cross-channel completion. In the baseline content-weighted exposure design, the prediction is therefore larger peer and state-agent exposure effects, together with a more negative peer–state-agent interaction, in French-speaking cantons. Larger main effects indicate that each channel carries a more complete persuasive bundle internally; the more negative interaction indicates that the two channels become more substitutable once each can more fully connect the normative *why* to the institutional *how*. In the workplace-switch design, the corresponding prediction is a larger effect of direct state-agent exposure, because state-agent contact is more likely to convey both the institutional *how* and the normative *why*.

I operationalize exposure to a French republican logic of state legitimacy using a binary French-exposure indicator. In the baseline content-weighted exposure design, $W_c = 1$ for French-speaking cantons and 0 for German-speaking cantons. In the workplace-switch design, $W_i = 1$ for French-speaking individuals and 0 for German-speaking individuals[10]. This is not a claim that the relevant mechanism operates through contemporaneous differences in formal institutions across language regions. Rather, the claim is that French-speaking environments are more exposed to a republican logic of legitimacy in which public authority is understood as the vehicle for realizing collectively justified norms. This interpretation is consistent with evidence that state-related public-service choices differ across Swiss language regions even under closely comparable institutional conditions. Using discontinuities at the Swiss language border, Athias and Wicht (2025) show that French-speaking municipalities are substantially less likely than adjacent German-speaking municipalities to contract public services to the private sector. This pattern points to a persistent cultural difference in how public authority is understood. The workplace-switch design further sharpens this interpretation. Because it includes canton-by-year fixed effects, it absorbs formal institutions, local political conditions, and discourse environments common to all individuals in a canton-year. Moreover, the heterogeneity is estimated using individual French-speaking identity interacted with the public-employment switch. The design therefore asks whether the same increase in workplace access to state agents has a different persuasive effect for French-speaking individuals, holding fixed the contemporaneous canton-year environment. This makes the language heterogeneity difficult to attribute to formal institutional differences alone, and more consistent with a cultural grammar of state legitimacy.

Figure 3 summarizes the cross-design evidence for the French-exposure shifter. In the content-weighted exposure design, French-speaking environments strengthen the marginal effect of peer exposure and shift the peer–state-agent interaction toward greater substitution (Table 6, Column (1)). The state-agent heterogeneity in the baseline exposure design is smaller, consistent with attenuation in the aggregate canton-level reach proxy. In the workplace-switch design, where access to state agents varies directly at the individual-year level, the public-employment effect is about three times larger for French-speaking workers than for German-speaking workers (Table 2, Column (7)), roughly 20 percentage points compared with 7 percentage points. The two designs therefore suggest that exposure to French republican state-building makes both peer and state-agent discourse more

[10]The heterogeneity exercise focuses on the French–German contrast rather than a French versus non-French comparison. The mechanism concerns exposure to a French republican logic of state legitimacy; German-speaking Switzerland therefore provides the cleanest benchmark within the same federal state, as well as the largest comparison group. Pooling Italian- and Romansh-speaking regions with German-speaking Switzerland would create a heterogeneous residual category.

likely to connect the normative *why* of universalism to the institutional *how* of public realization. As channels become more internally complete, endorsement relies less on assembling matched *why* and *how* reasons across channels, reducing the marginal role of cross-channel fit.

Figure 3: French exposure and bundle-based persuasion

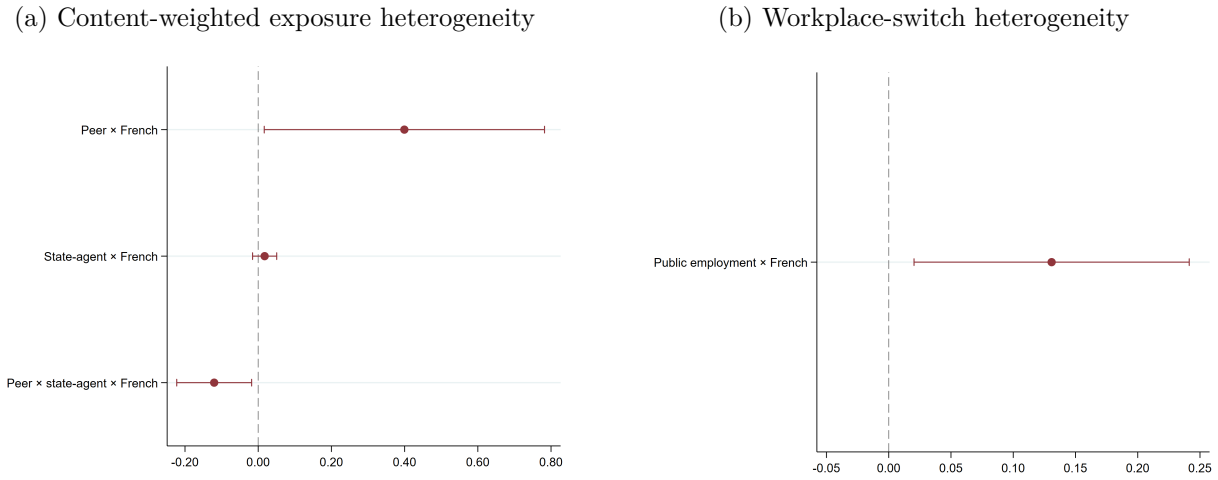

Notes: The figure reports heterogeneity estimates for the French-exposure mechanism. Panel A plots coefficients from the baseline content-weighted exposure specification interacted with a French-speaking canton indicator. Panel B plots heterogeneity in the preferred workplace-switch design, estimated among private-to-public switchers and stayers in non-HES industries by interacting public employment with individual French-speaking identity. The workplace-switch specification includes individual fixed effects and canton-by-year fixed effects, so the heterogeneity is identified holding fixed the contemporaneous canton-year environment. Horizontal bars report 90% confidence intervals. In Panel A, positive coefficients on peer or state-agent exposure indicate stronger within-channel persuasion in French-speaking environments; a negative coefficient on the peer–state-agent interaction indicates greater substitution between channels, consistent with stronger within-channel bundling.

The remaining concern is compositional. French-speaking individuals or French-speaking environments may differ in receiver-side characteristics that also shape responsiveness to public-sphere persuasion. In other words, the French-exposure pattern could reflect differential demand for certain arguments rather than a supply-side shift in the content of public discourse. I address this concern by allowing the exposure–endorsement mapping to vary with two predetermined receiver-side benchmarks: canton-level support for state secularity and predetermined cantonal ideology. Support for state secularity captures the possibility that the French pattern reflects secular receiver demand; ideology captures the possibility that it reflects a more left-leaning demand for universalistic arguments. In each case, I interact the benchmark with peer exposure, state-agent exposure, and their interaction. The French-exposure pattern is not absorbed by these adjustments: the French-speaking interaction with peer exposure remains positive, and the French-speaking interaction with the peer–state-agent term remains negative and statistically significant (Table 6, Columns (2) and (3)). This weighs against the interpretation that the results merely reflect more secular or more left-leaning receivers. Instead, the evidence is more consistent with French exposure shifting the content architecture of persuasion itself.

This interpretation also helps qualify the descriptive pattern reported in Section 4: endorsement of institutional universalism is higher among French-speaking respondents. The French-exposure results should not be read as merely reflecting a higher baseline attachment to universalistic rules or

Table 6: French exposure and receiver-side composition

Y_{ict}	(1)	(2)	(3)
	Endorsement of institutional universalism		
French $\times X^P_{ct}$	0.399*	0.519*	0.407**
	(0.221)	(0.285)	(0.188)
French $\times X^S_{ct}$	0.0176	0.0338	0.0415*
	(0.0190)	(0.0293)	(0.0217)
French $\times X^P_{ct} \times X^S_{ct}$	-0.120*	-0.183***	-0.187***
	(0.0590)	(0.0581)	(0.0543)
Receiver-composition interactions	No	Secularity support	Predetermined ideology
Observations	6,586	6,586	6,576

Notes: Column (1) reports the baseline French-exposure heterogeneity specification. Column (2) allows the exposure–endorsement mapping to vary with canton-level support for state secularity by interacting referendum support for church–state separation with peer exposure, state-agent exposure, and their interaction. Column (3) analogously allows the exposure–endorsement mapping to vary with predetermined cantonal ideology. The coefficients reported are the French-speaking canton interactions with peer exposure, state-agent exposure, and their interaction. All specifications include the fixed effects and controls from the baseline specification. Standard errors are clustered at the canton level.

greater institutional loyalty. They indicate that exposure itself maps differently into endorsement in French-speaking environments. This is more consistent with a republican grammar in which public authority is judged by its capacity to realize collectively justified norms. Such a grammar can make universalistic rules more persuasive while also making existing institutions more exposed to criticism when they fall short of those rules, a tension characteristic of modern public self-justification (Pippin (1999)).

7 Conclusion

Institutions are legitimate only if the norms they realize are recognized as publicly justifiable. This paper studies how such recognition is formed. The model formalizes a mechanism linking institutional legitimacy to a persuasion architecture: citizens must encounter public reasons about both the normative merit of a candidate norm and the institutional means through which that norm can be realized. This distinction matters because support for a norm as a private moral principle need not imply support for its institutionalization. A norm becomes a rule of public action only when citizens can see not only why it is desirable, but also how it can be sustained through public authority.This gives a distinctive role to agents with a comparative advantage in producing implementability reasons: state agents relative to peers. When state agents bundle implementability with publicly contestable merit arguments, they can strengthen legitimate endorsement. But when implementability travels without a publicly disciplined normative *why*, the same channel can amplify prevailing merit claims in the public sphere. This ambivalence is precisely why the organization of public discourse matters. When the normative *why* and institutional *how* are bundled, persuasion can generate shared endorsement. When they are split across channels or fail to meet, institutional consensus becomes harder to sustain. The model therefore links the formation of institutional legitimacy to the organization of public discourse, not simply to the amount of exposure to supportive agents.

The empirical evidence supports this mechanism. Using Swiss panel data, I show first that content-weighted exposure to both peers and state agents increases endorsement of institutional universalism. The negative interaction between the two exposure channels is consistent with within-

channel bundling: each channel is less marginally persuasive when the other already supplies a strong justificatory bundle. I then use a workplace-switch design to hold fixed the canton-year discourse environment. Moving from private to public employment substantially increases endorsement, and the effect is concentrated among state agents operating at the relevant level of public authority. These results are difficult to reconcile with simple imitation, generic ideological movement, or generalized institutional trust. They instead point to content-specific persuasion through public reasons. Adjacent outcomes help identify the argumentative content of this persuasion. They show that peer exposure makes equal treatment compatible with limits on welfare spending, while state-agent exposure increases endorsement while sharpening critical evaluations of democratic performance. Further mechanism tests show that the persuasion architecture varies with institutional and cultural environments in theoretically meaningful ways. Church–state separation strengthens the public-sphere basis of institutional justification: where ecclesiastical authority is less formally backed by the coercive fiscal authority of the state, public norms must be defended more explicitly through reasons that can be publicly exchanged and institutionally realized. French republican state legitimacy shifts a different margin: it makes the link between collectively justified norms and public authority more salient. Across both shifters, the evidence points in the same direction. Environments that make merit and implementability more likely to travel together strengthen within-channel persuasion and reduce the need for cross-channel completion.

The broader implication is that institutions shape legitimacy not only by changing incentives, constraints, or formal rights, but by structuring the reasons through which citizens evaluate institutionalized norms. Democratic authorization is therefore not a sufficient foundation of legitimacy. Authorization makes public justification consequential, but it does not itself make norms publicly justifiable or institutionally realizable. Understanding when institutions become legitimate therefore requires studying not only what institutions do, but how they make their underlying norms publicly intelligible. The paper also qualifies a common opposition between state coercion and market persuasion. The state is not per se domination, just as the market is not the only domain in which persuasion operates. Public authority can itself become an object of normative persuasion: citizens deliberate not only about how rules are implemented, but about the reasons for which institutional rules should bind them. In this sense, culture is not only an inherited stock of values or a set of stable preferences. It is also something actors deliberate about: a field in which public reasons make norms intelligible, contestable, and ultimately capable of being realized as rules of public action.

References

Acemoglu, Daron and Matthew O. Jackson (2015) "History, Expectations, and Leadership in the Evolution of Social Norms," *The Review of Economic Studies*, 82 (2), 423–456, 10.1093/restud/rdu039.

Acemoglu, Daron and James A. Robinson (2006) "De Facto Political Power and Institutional Persistence," *American Economic Review*, 96 (2), 325–330, 10.1257/000282806777212549.

——— (2008) "Persistence of Power, Elites, and Institutions," *American Economic Review*, 98 (1), 267–293, 10.1257/aer.98.1.267.

——— (2012) *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, New York: Crown Publishing Group.

——— (2025) "Culture, Institutions, and Social Equilibria: A Framework," *Journal of Economic Literature*, 63 (2), 637–92, 10.1257/jel.20241680.

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya (2015) "Radio and the Rise of The Nazis in Prewar Germany," *The Quarterly Journal of Economics*, 130 (4), 1885–1939, 10.1093/qje/qjv030.

Aneja, Abhay and Guo Xu (2024) "Strengthening State Capacity: Civil Service Reform and Public Sector Performance during the Gilded Age," *American Economic Review*, 114 (8), 2352–87, 10.1257/aer.20230019.

Arendt, Hannah (1963) *Eichmann in Jerusalem: A Report on the Banality of Evil*, New York: The Viking Press.

Arrow, K. (1951) *Social Choice and Individual Values*: Cowles Foundation.

Athias, Laure and Bruno Ventelou (2026) "The Heart Has Reasons That Reason Cannot Understand: Private vs Institutional Universalism," working paper.

Athias, Laure and Pascal Wicht (2025) "Make or buy for public services: Culture matters for efficiency considerations," *The Quarterly Review of Economics and Finance*, 99, 101899, https://doi.org/10.1016/j.qref.2024.101899.

Audier, F. and M. Bacache-Beauvallet (2007) "Emploi dans la fonction publique et fonctions "d'intérêt public": Que nous apprennent les comparaisons internationales ?" *Revue de l'OFCE*, 103, 323–350, 10.3917/reof.103.0323.

Berman, Sheri (1997) "Civil Society and the Collapse of the Weimar Republic," *World Politics*, 49 (3), 401–429, 10.1353/wp.1997.0008.

Besley, Timothy (2020) "State Capacity, Reciprocity, and the Social Contract," *Econometrica*, 88 (4), 1307–1335, https://doi.org/10.3982/ECTA16863.

Besley, Timothy, Robin Burgess, Adnan Khan, and Guo Xu (2022) "Bureaucracy and Development," *Annual Review of Economics*, 14 (Volume 14, 2022), 397–424, https://doi.org/10.1146/annurev-economics-080521-011950.

Besley, Timothy and Torsten Persson (2009) "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, 99 (4), 1218–1244, 10.1257/aer.99.4.1218.

——— (2011) *Pillars of Prosperity*, Princeton: Princeton University Press, doi:10.1515/9781400840526.

——— (2019) "Democratic Values and Institutions," *American Economic Review: Insights*, 1 (1), 59–76, 10.1257/aeri.20180248.

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro (2024) "Cross-Country Trends in Affective Polarization," *The Review of Economics and Statistics*, 106 (2), 557–565, 10.1162/rest_a_01160.

Bursztyn, Leonardo, Georgy Egorov, and Stefano Fiorin (2020) "From Extreme to Mainstream: The Erosion of Social Norms," *American Economic Review*, 110 (11), 3522–48, 10.1257/aer.20171175.

Castoriadis, Cornelius (1987) *The Imaginary Institution of Society*, Cambridge, Massachusetts: The MIT Press.

——— (1996) *La montée de l'insignifiance: Les carrefours du labyrinthe IV*, Paris: Éditions du Seuil.

Condorcet, Marie-Jean-Antoine-Nicolas de Caritat (1785) *Essay on the Application of Analysis to the Probability of Majority Decisions (Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix)*: Paris.

Darnton, Robert (1982) *The Literary Underground of the Old Regime*: Harvard University Press.

Ekeh, Peter P. (1975) "Colonialism and the Two Publics in Africa: A Theoretical Statement," *Comparative Studies in Society and History*, 17 (1), 91–112.

Enke, Benjamin, Ricardo Rodriguez-Padilla, and Florian Zimmermann (2022) "Moral Universalism: Measurement and Economic Relevance," *Management Science*, 68 (5), 3175–3973, https://doi.org/10.1287/mnsc.2021.4086.

Esping-Andersen, Gosta (1990) *The three worlds of welfare capitalism*, Princeton (NJ): Princeton University Press.

Eugster, Beatrix, Rafael Lalive, Andreas Steinhauer, and Josef Zweimüller (2011) "The Demand for Social Insurance: Does Culture Matter?" *The Economic Journal*, 121 (556), F413–F448, https://doi.org/10.1111/j.1468-0297.2011.02479.x.

Eugster, Beatrix and Raphaël Parchet (2019) "Culture and Taxes," *Journal of Political Economy*, 127 (1), 296–337, 10.1086/700760.

Francois, Patrick (2003) "Not-for-Profit Provision of Public Services," *The Economic Journal*, 113 (486), C53–C61.

Gentzkow, Matthew and Emir Kamenica (2017) "Competition in Persuasion," *The Review of Economic Studies*, 84 (1 (298)), 300–322.

Graeber, David and Marshall Sahlins (2017) *On Kings*, Chicago: Hau Books.

Gratton, Gabriele and Barton E Lee (2024) "Liberty, Security, and Accountability: The Rise and Fall of Illiberal Democracies," *The Review of Economic Studies*, 91 (1), 340–371, 10.1093/restud/rdad030.

Gregg, Paul, Paul Grout, Anita Ratcliffe, Sarah Smith, and Frank Windmeijer (2011) "How important is pro-social behaviour in the delivery of public services?" *Journal of Public Economics*, 95 (7), 758–766.

Greif, Avner (1994) "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies," *Journal of Political Economy*, 102 (5), 912–950.

Habermas, Jürgen (1990) *Moral Consciousness and Communicative Action*, Cambridge, MA: MIT Press.

Habermas, Jürgen (1989) *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*: MIT Press.

Harsanyi, John C. (1955) "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy*, 63 (4), 309–321.

Heldring, Leander (2026) "The cost of state-building: Evidence from Germany," working paper.

Henrich, Joseph (2020) *The WEIRDest People in the World: How the West Became Psychologically Peculiar and Particularly Prosperous*, New York: Farrar, Strauss, and Giroux.

Hill, Christopher (1972) *The World Turned Upside Down: Radical Ideas During the English Revolution*, London: Temple Smith.

Horkheimer, Max and Theodor W. Adorno (1944) *Dialektik der Aufklärung: Philosophische Fragmente*, Amsterdam: Querido Verlag.

Hosli, Madeleine O., Luiza Martins Santos, and Henry Boeree (2025) "A political-economic analysis of Swiss referendums 1848 to 2022: turnout, acceptance rates and the double-majority threshold," *Constitutional Political Economy*, 36 (2), 245–270, 10.1007/s10602-025-09468-1.

Kamenica, Emir and Matthew Gentzkow (2011) "Bayesian Persuasion," *American Economic Review*, 101 (6), 2590–2615, 10.1257/aer.101.6.2590.

Kohn, Hans (1950) "Romanticism and the Rise of German Nationalism," *The Review of Politics*, 12 (4), 443–472.

Kuran, Timur (1991) "The East European Revolution of 1989: Is it Surprising that We Were Surprised?" *The American Economic Review*, 81 (2), 121–125.

——— (1995) *Private Truths, Public Lies*: Harvard University Press.

Luo, Xizi and Melanie Meng Xue (2025) "Enlightenment Under Autocracy: The Origins of Liberalism in China," working paper.

Malinowski, Bronislaw (1922) *Argonauts Of The Western Pacific: An Account of Native Enterprise and Adventure in the Archipelagoes of Melanesian New Guinea*, London: Routledge & Kegan Paul.

Mokyr, Joel (2010) *The Enlightened Economy*: Yale University Press.

——— (2016) *A Culture of Growth*: Princeton University Press.

Morrill, John (1993) *The Nature of the English Revolution: Essays*, London: Longman.

Moscona, Jacob, Nathan Nunn, and James A Robinson (2026) "Searching for Fish in Trees? Economic Development when Context Matters," Working Paper 34810, National Bureau of Economic Research, 10.3386/w34810.

Mosse, George L. (1964) *The Crisis of German Ideology: Intellectual Origins of the Third Reich*, New York: Grosset and Dunlap.

North, D. C. (1990) *Institutions, Institutional Change and Economic Performance*, Cambridge: Cambridge UniversityPress.

Pippin, Robert B. (1999) *Modernism as a Philosophical Problem: On the Dissatisfactions of European High Culture*: Wiley-Blackwell.

——— (2008) *Hegel's Practical Philosophy: Rational Agency as Ethical Life*, Cambridge: Cambridge University Press.

Popkin, Jeremy D. (1990) *News and Politics in the Age of Revolution: Jean Louis Carrier and the Révolutions de Paris, 1789–1797*: Cornell University Press.

Portmann, Marco, Frederik Blümel, and Christoph A. Schaltegger (2024) "Lohnprämie in den öffentlichen Verwaltungen: Neue Analyse 2024," Policy Papers 18, IWP.

Putnam, Robert D. (2000) *Bowling Alone:The Collapse and Revival of American Community*, New York: Simon & Schuster.

Putnam, Robert D., Robert Leonardi, and Raffaella Y. Nonetti (1993) *Making Democracy Work: Civic Traditions in Modern Italy*: Princeton University Press.

Rawls, John (1971) *A Theory of Justice*, Cambridge, MA: Harvard University Press.

——— (1997) "The Idea of Public Reason Revisited," *University of Chicago Law Review*, 64, 765–807.

Rousseau, Jean-Jacques (1923) *The Social Contract and Discourses*, (G. D. H. Cole, Trans.). London and Toronto: J. M. Dent & Sons, Original work published 1762.

Sahlins, Marshall (1976) *Culture and Practical Reason*, Chicago: The University of Chicago Press.

——— (2000) *Culture in Practice: Selected Essays*, New York: Zone Books.

Sieyès, Emmanuel-Joseph (1789) *Qu'est-ce que le Tiers-État?*: Paris.

Tabellini, Guido (2008) "Institutions and Culture," *Journal of the European Economic Association*, 6 (2-3), 255–294, 10.1162/JEEA.2008.6.2-3.255.

Taylor, Charles (2004) *Modern Social Imaginaries*, Durham and London: Duke University Press.

——— (2007) *A Secular Age*, Cambridge, Massachusetts, and London, England: The Belknap Press of Harvard University Press.

Tilly, C. (1975) *Reflections on the History of European State-Making*, Princeton: Princeton University Press.

——— (1985) *War Making and State Making as Organized Crime*, Cambridge: Cambridge University.

Tyler, Tom R. (2006) *Why People Obey the Law*, Princeton, New Jersey: Princeton University Press.

Vatter, Adrian, Tobias Arnold, Alexander Arens et al. (2024) "Patterns of Democracy in the Swiss Cantons, 1979–2023," Dataset. University of Bern, Institute of Political Science.

Weber, Eugen (1976) *Peasants into Frenchmen: The Modernization of Rural France, 1870-1914*, Stanford, Calif.: Stanford University Press.

Weber, Max (1922) *Economy and Society*, 4th ed. Tübingen.

——— (1978) *Economy and Society: An Outline of Interpretive Sociology*, Berkeley, CA: University of California Press.

Zucman, Gabriel (2024) “A Blueprint for a Coordinated Minimum Effective Taxation Standard for Ultra-High-Net-Worth Individuals,”Technical report, commissioned by the Brazilian G20 Presidency.

Appendix: Figures and Tables

Table A1: Model notation

Object	Definition / interpretation
Indices and sets	
i	Citizen (agent).
(c,t)	Public-sphere context c at date t.
$j \in \{S,P\}$	Communication channel: S = state-agent channel; P = peer channel.
$k \in \{m,h\}$	Content dimension: m = merit/"why"; h = implementation/"how".
Latent states, priors, and signals	
m, h	Latent "merit" and "implementation" states relevant for persuasion.
μ_k, σ_k^2	Prior mean and variance for dimension $k \in \{m,h\}$.
$\tau_k \equiv \sigma_k^{-2}$	Prior precision for dimension k.
$\tilde{x}^k_{j,i,c,t}$	Signal received by i from channel j about dimension k in (c,t).
$\mathcal{F}_{i,c,t}$	Citizen i's information set in (c,t): $\mathcal{F}_{i,c,t} = \sigma\big(E^S_{i,c,t}, E^P_{i,c,t}, \tilde{x}^m_{S,i,c,t}, \tilde{x}^h_{S,i,c,t}, \tilde{x}^m_{P,i,c,t}, \tilde{x}^h_{P,i,c,t}\big)$.
$\sigma(\cdot)$	Sigma-algebra generated by the listed random variables (information available to the agent).
Exposure / reach	
$E^j_{i,c,t}$	Realized exposure of citizen i to channel j in (c,t) (nonnegative).
$E^j_{c,t}$	Aggregate reach (average exposure intensity) of channel j in public sphere (c,t).
$\eta^j_{i,c,t}$	Idiosyncratic reach component: $E^j_{i,c,t} = E^j_{c,t} + \eta^j_{i,c,t}$, $\mathbb{E}[\eta^j_{i,c,t} \mid c,t] = 0$.
Precision and effective precision	
e^k_j	Producer effort by channel j to generate precision on dimension k.
$\upsilon^k_j(\cdot)$	Precision technology: $\pi^k_j = \upsilon^k_j(e^k_j)$ (increasing, concave).
π^k_j	Precision supplied by channel j on dimension k.
Q^k_j	Effective precision / intensity from channel j on dimension k: $Q^k_j \equiv \pi^k_j E^j$ (suppressing indices when convenient).
$Q_{k,i,c,t}$	Total effective precision for citizen i on dimension k in (c,t): $Q_{k,i,c,t} \equiv \sum_{j\in\{S,P\}} \pi^k_j E^j_{i,c,t}$.
Bundling and endorsement	
$f(x)$	Poisson "arrival" map: $f(x) \equiv 1 - e^{-x}$.
Λ_j	Within-channel bundle intensity: $\Lambda_j \equiv f(Q^m_j)\, f(Q^h_j)$.
$\Lambda_\times$	Cross-channel fit (extension): $\Lambda_\times \equiv f(Q^m_S) f(Q^h_P) + f(Q^h_S) f(Q^m_P)$.
Λ	Total bundle intensity: $\Lambda \equiv \Lambda_S + \Lambda_P$ (benchmark) or $\Lambda \equiv \Lambda_S + \Lambda_P + \Lambda_\times$ (extension).
s	Endorsement probability / share: $s = u(\Lambda)$.
$u(\cdot)$	Endorsement link, increasing and concave ($u' > 0$, $u'' \le 0$); canonical case $u(\Lambda) = 1 - e^{-\Lambda}$.
κ	Off-diagonal "fit" index (local complementarity at the origin): $\kappa \equiv \pi^m_S \pi^h_P + \pi^h_S \pi^m_P$.
Producer payoffs and costs	
$C^k_j(\cdot)$	Cost of effort e^k_j for channel j on dimension k (convex).
B	Social benefit scale from endorsement/legitimacy ($B > 0$).
β_j	Internalization weight: channel j internalizes fraction $\beta_j \in (0,1]$ of B.
δ	Implementation comparative-advantage wedge (Assumption 1).
Dynamics (when used)	
S_t	Aggregate endorsement at date t (e.g., population average of s).
$E^j(S)$	Exposure map: reach as a function of current endorsement S (public-sphere propagation).
$\Phi(S)$	Update map: $\Phi(S) \equiv u(\Lambda(S))$ so $S_{t+1} = \Phi(S_t)$.

Table A2: Variable Descriptions and Sources

Table	Variable	Definition	Coding	Source	Time period
1,2,5,6	Institutional universalism	Are you in favour of Switzerland offering foreigners the same opportunities as those offered to Swiss citizens, or in favour of Switzerland offering Swiss citizens better opportunities?	0 = better opportunities for Swiss citizens and 1 = equal opportunities	Swiss Household Panel	1999–2003
3,4	Welfare	Are you in favour of a diminution or in favour of an increase of the Confederation social spendings?	1 = diminution of social spending, 2 = no change and 3 = more social spending	Swiss Household Panel	1999–2003
3,4	Redistri.	Are you in favour of an increase or in favour of a decrease of the tax on high incomes?	1 = decrease, 2 = no change and 3 = increase	Swiss Household Panel	1999–2003
3,4	Army	Are you in favour of Switzerland having a strong army or for Switzerland not having an army?	1 = no army, 2 = neither and 3 = strong army	Swiss Household Panel	1999–2003
3,4	Env. vs growth	Are you in favour of Switzerland being more concerned with protection of the environment than with economic growth, or in favour of Switzerland being more concerned with economic growth than with protection of the environment?	1 = more concerned with economic growth, 2 = neither and 3 = more concerned with environment protection	Swiss Household Panel	1999–2003
3,4	Join EU	Are you in favour of Switzerland joining the European Union or are you in favour of Switzerland staying outside of the European Union?	0 = staying outside of EU and 1 = joining the EU	Swiss Household Panel	1999–2003
3,4	Ideology	When they talk about politics, people mention left and right. Personally, where do you position yourself?	0 = 'left' and 10 = 'right'	Swiss Household Panel	1999–2003
3,4	Sat. Dem.	Overall, how satisfied are you with the way in which democracy works in our country?	0 = 'not at all satisfied' and 10 ='completely satisfied	Swiss Household Panel	1999–2003
3,4	Trust fed. gov.	How much confidence do you have in the Federal Government?	0 = 'no confidence' and 10 = 'full confidence	Swiss Household Panel	1999–2003
1,3,5,6	Average ideological position of the cantonal executive	Computed as the mean of seat-level party positions within the executive council. Each seat is coded on a four-point left–right scale ranging from 1 (far left) to 4 (far right); the index is obtained by summing these seat codes and dividing by the number of executive seats, so that higher values indicate a more right-leaning government.	1 = far left to 4 = far right	Vatter et al. (2024)	1999–2003
1,3,5,6	Foreign pop share	Share of the cantonal resident population holding foreign citizenship, measured as the number of non-citizens divided by the total resident population.	Percentage	State Secretariat for Migration	1999–2003

Table	Variable	Definition	Coding	Source	Time period
1,3,5,6	Referendum	Number of referenda held in canton.	Count	Vatter et al. (2024)	1999–2003
1,3,5,6	Unemployment rate	Cantonal unemployment rate in percent, measured as the share of unemployed individuals in the active population.	Percentage	State Secretariat for Economic Affairs	1999–2003
1,3,5,6	Taxable income per capita	Per-capita taxable income base, measured as the total mass of income subject to the personal income tax divided by the cantonal resident population.	Swiss francs per capita	Federal Finance Administration	1999–2003
5,6	Support for state secularity	Canton-level vote shares in the 1980 federal popular initiative “concerning the complete separation of Church and State", submitted to popular vote on 2 March 1980.	0 to 1	Swiss Federal Chancellery	1980
5	Secular canton	Cantonal-level indicator capturing whether the payment of the church tax is legally compulsory.	0 = payment of the church tax is voluntary and 1 = payment of the church tax is legally compulsory	Swiss Tax Conference	1946, 1960, 1992, predetermined
1,3,5,6	Public-employment share as a proxy for E_c^S used in X_{ct}^S	Percentage of a canton’s employed labor force working in the public sector measured as the number of full-time equivalent.	Percentage	Federal Statistical Office	2001
1,3,5,6	Association density as a proxy for E_c^P used in X_{ct}^P	Cantonal number of registered associations per 1,000 inhabitants.	Number per 1000 inhabitants	Federal Statistical Office	1995
1,3,5,6	$p_{c,t}^P$ used in X_{ct}^P	Cantonal share of workers in the private sector who endorse institutional universalism	0 to 1	Swiss Household Panel	1999–2003
1,3,5,6	$p_{c,t}^S$ used in X_{ct}^S	Cantonal share of workers in the public sector who endorse institutional universalism	0 to 1	Swiss Household Panel	1999–2003

Table A3: Summary Statistics

	Mean	Std. dev.	Min	Max	Between SD	Within SD
Panel A. Individual-level time-varying variables						
Institutional universalism	0.616	0.486	0.000	1.000	0.440	0.250
Peer content-weighted exposure (X_{ct}^{P})	0.000	0.165	-0.608	0.642	0.156	0.053
State-agent content-weighted exposure (X_{ct}^{S})	0.000	1.351	-4.633	5.850	1.317	0.339
Welfare	2.247	0.714	1.000	3.000	0.601	0.403
Redistribution	2.591	0.641	1.000	3.000	0.546	0.362
Army	2.189	0.794	1.000	3.000	0.709	0.369
Env. vs growth	2.295	0.712	1.000	3.000	0.590	0.417
Join EU	0.529	0.499	0.000	1.000	0.463	0.191
Ideology	4.959	2.115	0.000	10.000	1.903	0.925
Sat. w. Democracy	5.939	2.006	0.000	10.000	1.749	1.060
Trust in fed. gov.	5.804	2.244	0.000	10.000	1.987	1.092
Panel B. Cantonal-level time-varying variables						
Average ideology	2.710	0.347	1.800	3.571	–	–
Foreign pop share	0.194	0.054	0.074	0.334	–	–
Nb. of Referenda	3.767	3.432	0.000	24.000	–	–
Unemployment rate	2.368	1.077	0.300	6.500	–	–
Taxable inc. pc	32'555.02	6'591.835	22'634.02	75'379.49	–	–
Panel C. Cantonal-level pre-determined variables						
French-speaking canton	0.236	0.425	0.000	1.000	–	–
Secularity support	0.215	0.058	0.038	0.352	–	–
Secular canton	0.107	0.309	0.000	1.000	–	–
Public-employment share	6.161	1.532	3.500	11.100	–	–
Association density	0.899	0.190	0.444	1.316	–	–
Observations (person-years)						19,243
Individuals (idpers)						6,920
Years						1999–2003

Notes. The sample is restricted to non-working Swiss respondents aged 18+; canton non-movers. Unit of observation is the individual-year. Between SD is the standard deviation of individual means. Within SD is the standard deviation of deviations from individual means.

Figure A1: Endorsement of the universalism institutional norm over time

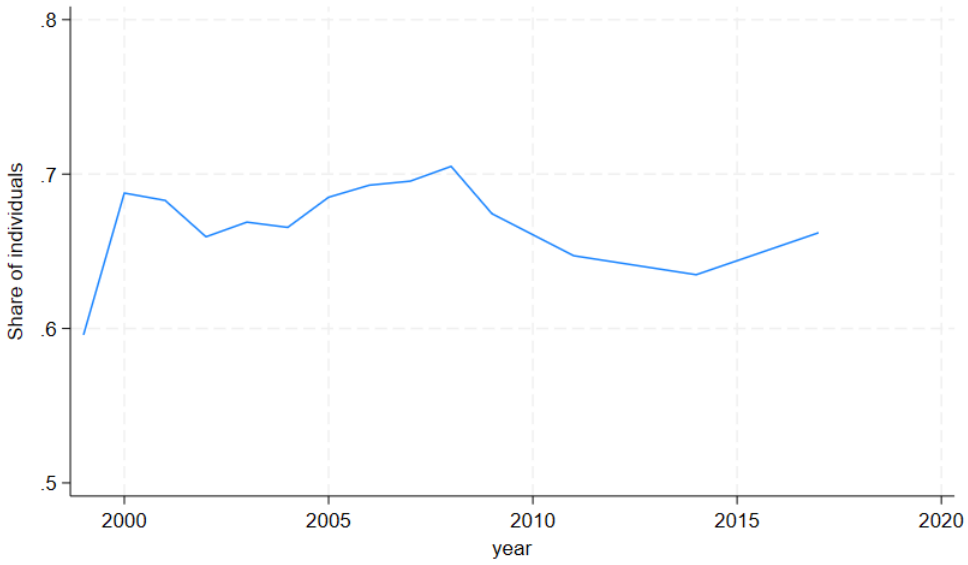

Note: The figure depicts the share of Swiss individuals aged 18 and over who endorse the universalism institutional norm over the maximal period of time with data availability, corresponding to the yearly average of the variable Y_{ict}.

Figure A2: Correlates of endorsement of the universalism institutional norm

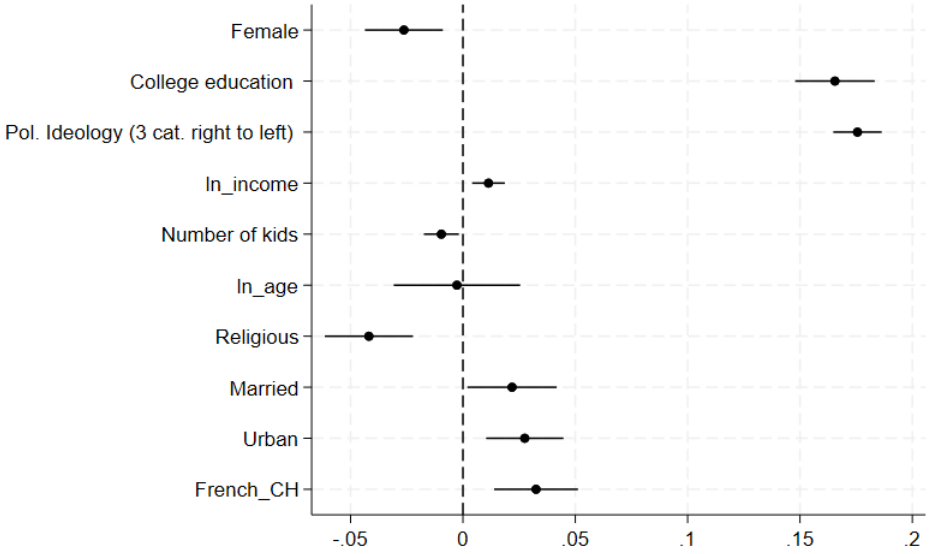

Note: This figure presents OLS Estimates from regressions of the Y_{ict} variable on all sociodemographic correlates, with SE clustered at the individual level, along with 95% confidence intervals.

Table A4: Institutional universalism endorsement by institutional sector (1999–2003)

	Public sector		Private sector		
	(1)		(2)		(3)
	Mean	N	Mean	N	Difference
All workers	0.742	4,250	0.657	7,473	0.085***
	(0.438)		(0.475)		(0.009)
Within HES industries	0.800	2,112	0.769	924	0.030*
	(0.400)		(0.421)		(0.016)
Within non-HES industries	0.666	1,569	0.641	5,265	0.026*
	(0.472)		(0.480)		(0.014)

Note: The table reports means of Y_{ict}, an indicator for endorsement of institutional universalism. The sample includes Swiss respondents aged 18+ working in the public or private sector. Industry (HES/non-HES) is observed only for a subset of employed respondents; breakdown columns therefore exclude observations with missing industry and need not sum to the overall sample. HES denotes health, education, and social care industries; non-HES denotes all other industries. Column (3) reports mean differences with robust standard errors in parentheses. Standard deviations are in parentheses below means in Columns (1) and (2). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A3: Yearly share of individuals who endorse the universalism social norm, by group

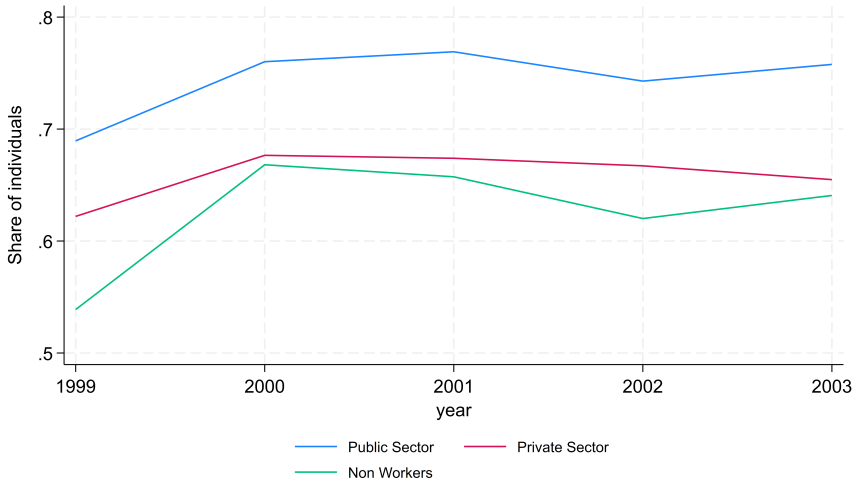

Note: The figure shows the share of Swiss individuals aged 18 and over who endorse institutional universalism over time, measured as the yearly mean of Y_{ict}, separately for non workers, state employees, and private sector workers.

Figure A4: Self-selection based on endorsement of institutional universalism across industries

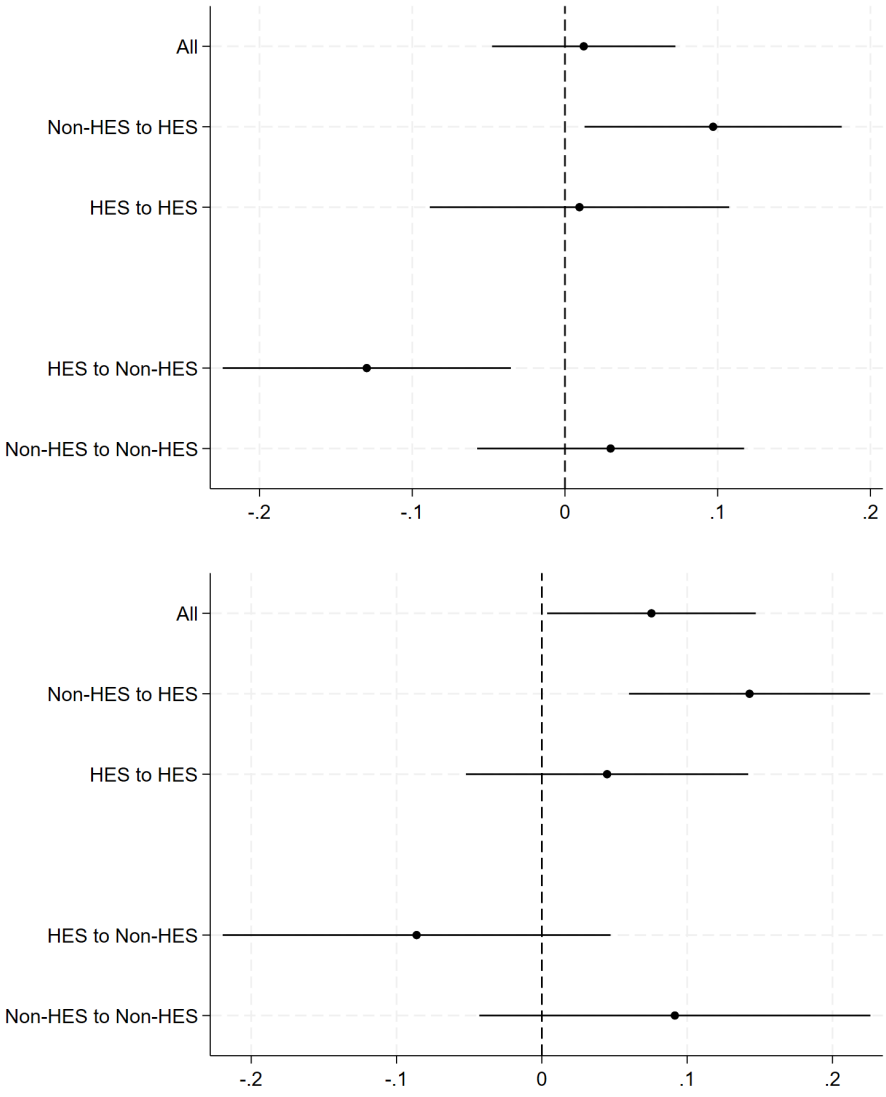

Notes: The figure summarizes regression coefficients of the switching indicators, capturing systematic differences in universalism norm endorsement between those who stay and those who switch at some future point. The above figure considers all switches while the figure below excludes switches at the federal level. See the text for further details. The 90% confidence intervals are based on standard errors clustered at the individual level.

Table A5: Switchers to *vs.* stayers in the public non-HES sector

Variable	(1) Switchers	(2) Stayers	(3) Difference
Share of female	0.429	0.351	0.077
	(0.497)	(0.478)	(0.054)
Share of married	0.623	0.676	-0.053
	(0.487)	(0.469)	(0.053)
Number of children	1.388	1.584	-0.197
	(1.289)	(1.310)	(0.141)
Age	38.714	40.512	-1.798*
	(9.467)	(8.666)	(1.021)
Share of Educ.: college level	0.143	0.242	-0.099**
	(0.352)	(0.428)	(0.040)
Share of living in urban area	0.612	0.574	0.038
	(0.490)	(0.495)	(0.053)
Ln yearly income	10.924	11.051	-0.127*
	(0.612)	(0.602)	(0.070)
State Universalism	0.679	0.729	-0.050
	(0.470)	(0.445)	(0.056)
Ideology	4.273	4.540	-0.268
	(2.263)	(2.107)	(0.258)

Notes: The table displays socio-demographic characteristics of switchers from the private into the public sector and stayers in the public sector within non-HES industries. The last column presents the difference in the mean value of each variable between the two sectors. Columns (1) and (2) standard deviation in parentheses. Columns (3) robust standard errors in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Figure A5: Endorsement of institutional universalism – robustness of public employment in the non-HES sector, dropping one cantonal bureaucracy at a time

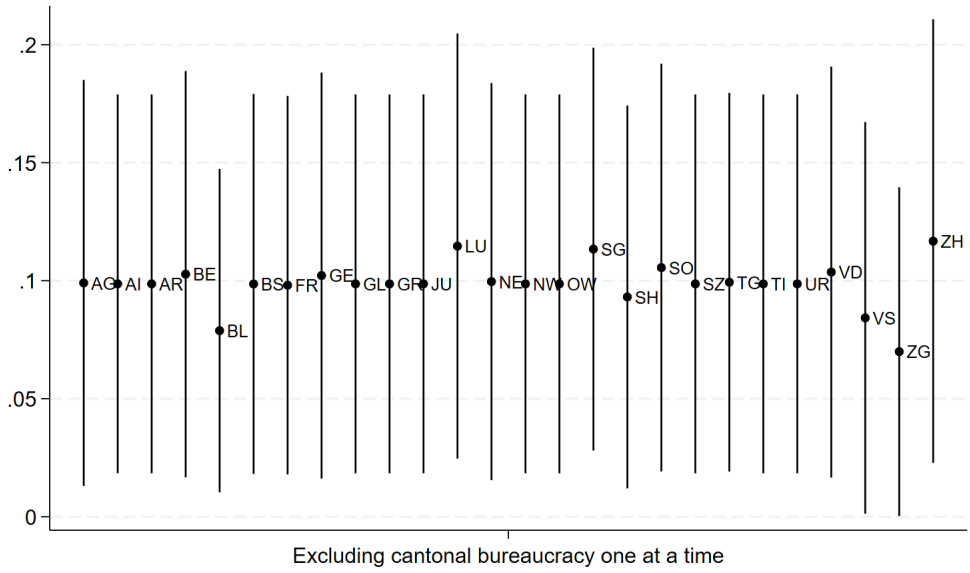

Notes: Figure reports point estimates of the state effect (specification from Column (4) in Table 2), considering switches within non-HES occupations, dropping one state at a time. Reporting 90% confidence intervals. Standard errors clustered at the individual level.

Table A6: Exposure to deliberation channels and endorsement of institutional universalism: Full sample (including movers)

	(1)	(2)	(3)	(4)	(5)
Y_{ict}	Endorsement of institutional universalism				
Peer content-weighted exposure (X^P_{ct})	0.224**	0.225**	0.311***	0.307**	0.285*
	(0.090)	(0.099)	(0.106)	(0.109)	(0.156)
State-agent content-weighted exposure (X^S_{ct})	0.021	0.023*	0.029**	0.029**	0.035**
	(0.015)	(0.012)	(0.013)	(0.013)	(0.017)
Peer × state-agent interaction ($X^P_{ct} \times X^S_{ct}$)	-0.025	-0.021	-0.052*	-0.050*	-0.087**
	(0.025)	(0.020)	(0.028)	(0.028)	(0.037)
Observations	9,041	8,957	8,957	8,949	8,957
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Canton fixed effects	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Contemporaneous canton controls	No	Yes	Yes	Yes	No
Canton-specific linear trends	No	No	Yes	Yes	Yes
Time-varying individual controls	No	No	No	Yes	No
Lagged canton controls	No	No	No	No	Yes

Notes: The dependent variable Y_{ict} is a binary indicator equal to one if respondent i endorses institutional universalism in canton c and year t. The sample is restricted to non-working Swiss respondents aged 18+. All specifications estimate equation (3) by OLS with individual, canton, and year fixed effects; standard errors are clustered at the canton level. The peer and state-agent content-weighted exposures, X^P_{ct} and X^S_{ct}, are mean-centered in the estimation sample so that the main-effect coefficients are evaluated at the mean of the other channel's content-weighted exposure. *Contemporaneous canton controls* include: (i) the weighted-average left–right ideology of the cantonal executive (constructed by mapping parties' cabinet shares onto a 1–4 scale using party positions from Vatter et al. (2024)); (ii) foreign population share; (iii) number of referenda; (iv) unemployment rate; and (v) taxable income per capita (cantonal average). *Lagged canton controls* replace (ii)–(v) by their one-year lag (while keeping the electoral calendar contemporaneous). *Time-varying individual controls* include marital-status indicators, number of children, and an indicator for health limitations in daily activities. Column (3) is the preferred specification; Columns (4)–(5) provide robustness to adding individual controls and to using lagged canton covariates. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.