Narrative Entanglement*

Adam Brzezinski
LSE

Luis Garicano
LSE and CEPR

24 May 2026

Abstract

Politicians rarely acknowledge trade-offs between a policy's economic and non-economic effects. We show that incentives to suppress such trade-offs also generate *narrative entanglement*: a shock to one policy dimension leads politicians to update their whole narrative bundle, even on unrelated dimensions, which shapes equilibrium policy. We test our theory by analyzing all speeches from the 9th European Parliament (2019–2024) with LLMs. We find entangled narratives across different policy domains, and show that following Russia's invasion of Ukraine politicians adjusted narratives on both the economic and environmental effects of climate policy, in line with our theory.

Keywords: narratives, Russia's invasion of Ukraine, climate skepticism
JEL codes: D72, D91, Q58

*We are grateful to Elliott Ash, Tim Besley, Gilat Levy, Elias Papaioannou, Stephane Wolton, and seminar participants at the Monash-Paris-Warwick-Zurich-CEPR Text-As-Data Workshop, LSE, LBS, Hong Kong University, University of Leicester, King's College London, LSE Environmental Populism Conference, IESE and the University of Oxford for helpful comments and insights. We thank Bruno Serra Lorenzo for valuable work as a research assistant.

1. Introduction

Politicians rarely acknowledge trade-offs between a policy's economic and non-economic effects. We provide a theory to show that incentives to suppress such trade-offs also result in what we call *narrative entanglement*: a shock that factually changes only one policy dimension leads politicians to re-optimize their whole narrative bundle and associated policy choice. For instance, climate policy has economic effects shaped by market conditions, and environmental effects governed by the laws of physics. Without entanglement, an increase in the cost of climate policy would lead politicians to change their statements about the economic effects of climate policy, but not about whether climate policy mitigates climate change. But with strategic incentives to hide trade-offs, rising costs would lead politicians to reconsider both their economic and environmental narratives. This matters for policymaking, as entangled narratives influence equilibrium policy and change how it reacts to economic shocks.

We provide a theory to formally study the incentives that give rise to narrative entanglement and influence policymaking. Two office-motivated parties that differ in their tax stance, "Left" and "Right", compete for votes by offering a policy alongside a narrative on the policy's economic and non-economic effects. Voters experience different private costs from the policy and form their beliefs in one of two ways: Bayesians can tell when a politician speaks the truth and update their beliefs accordingly; motivated reasoners do not weigh narratives by accuracy, but instead adopt the narrative that maximizes their prospective utility. We build our assumptions about motivated reasoning on a long tradition in psychology (Kunda, 1990) that has more recently entered economics (Bénabou, 2015; Epley and Gilovich, 2016). In the context of political economy, we find it intuitive to assume that some voters are motivated reasoners: incentives for non-pivotal voters to hold correct beliefs are limited (see also Caplan, 2007), while policy effects are very complex to estimate.

Narratives and policies differ in a fundamental way: while all voters experience the policy of the winning party, voters can disagree with each other on narratives. Disagreement can occur between Bayesian voters and motivated reasoners, but also between different motivated reasoners with diverging personal costs from the policy.

Motivated reasoners generate demand for entangled narratives that say that the policy is good or bad on every dimension. From the party's perspective, supplying entangled narratives to motivated reasoners simultaneously achieves two objectives: it maximizes the perceived payoff of the party's own platform, and it minimizes the perceived payoff of the opposing platform. Bayesian voters instead ignore entangled narratives if these are false, and only update their beliefs if they hear accurate narratives.

Truth-telling is not guaranteed even when the probability of Bayesian voters is high, and parties may instead exclusively supply entangled narratives. To see this, suppose that the Left gains electorally from implementing green policy because the median voter likes the policy. Suppose further that Bayesian voters would support the policy more strongly when learning the truth, e.g., for the sake of argument, that environmental benefits are very large but there may be some economic costs. But suppose also that motivated reasoners would be less supportive under the truth and more supportive under the no-trade-off narratives that climate policy is great for the environment and the economy. Then, the electoral benefit from telling motivated reasoners an entangled narrative may be larger even if Bayesian voters dominate the electorate—for instance, if the prior of Bayesian voters is close to the truth to start with, or if parties can come up with very extreme entangled narratives that strongly mobilize motivated reasoners.

This persistence of entanglement has fundamental policy implications, since narratives and policy are jointly determined in equilibrium. If there is no entanglement and narratives are anchored to the truth, shifts in policy costs will continuously shift the implied economic effects of the policy but not its implied non-economic effects. Under entanglement, in contrast, shifts in costs can lead politicians to discontinuously shift economic and non-economic narratives together and to update their policy accordingly. Policy change may be accelerated by entanglement, when a change in one narrative dimension transmits to the other dimension, and the new narrative bundle implies a different optimal policy choice. Policy change may also be decelerated, if one narrative dimension does not respond to a relevant shock because it is strategically tied down by the other narrative dimension. In general, policy under entangled narratives will be different from the optimal policy that would be implemented if voters learned the true narrative.

Interactions between parties' tax stance and economic narratives can give rise to polarized equilibria where parties offer different platforms and voters believe in different narratives. We provide an empirical model applied to climate policy to show why such interactions imply that, in polarized equilibria, left-wing parties support climate policy while right-wing parties oppose it. We also explicitly model the speech-making choice of individual politicians that constitute parties. Our model predicts that increases in economic costs of climate policy will tend to make left-wing politicians less likely to speak about climate policy, and right-wing politicians more likely to do so and to increasingly voice climate-skeptic narrative bundles.

We take the model's predictions to the data by studying all speeches in the 9th European Parliament (2019–2024), comprising 48,239 speeches after excluding procedural and similar comments.[1] We choose to investigate this setting because of how Russia's invasion of Ukraine

[1] One of us (Garicano) was an MEP for the first three years of the legislature.

on 24 February 2022 influenced the ongoing climate policy initiative of that parliament.[2] The invasion sharply raised EU energy prices (see Supplemental Appendix Figure C.1) and the short-run cost of green policies. Thus, the energy crisis changed the economic effects of climate policy. But it did not change whether climate policy mitigates climate change. If narratives are not entangled, politicians should only change statements about the economic costs of climate policy. If narratives are entangled, politicians should also change their environmental narratives.

We use large language models (LLMs) to identify speeches on the topic of climate policy, and to classify their environmental and economic narratives. Importantly, we classify each narrative dimension separately, so that entanglement is not an artifact of LLMs trying to be consistent in their classification. In addition to our main application on climate policy, we also classify narratives on two other prominent policy domains: immigration and digitization. We extensively verify the robustness of our approach by using a variety of different prompts as well as different models.

Figure 1 offers a first look at our narrative data. The red spikes show the daily number of parliamentary speeches expressing that climate policy is not needed to combat climate change. Climate skepticism was muted at the start of the parliamentary session but much more prevalent toward the end. The average number of climate-skeptic speeches per meeting day more than doubled after Russia's invasion of Ukraine.

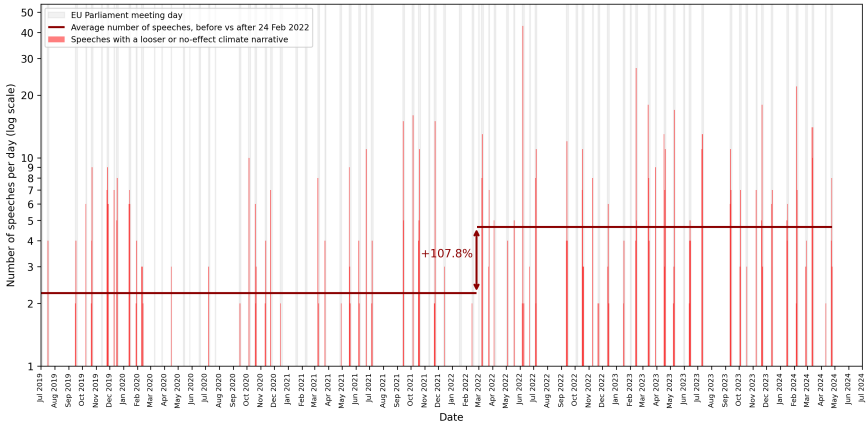

Note: Red bars show the daily number of speeches expressing a climate-skeptic environmental narrative, based on our methodology from Section 3; dark-red lines show averages before and after Russia's invasion of Ukraine (24 February 2022). Shaded bands mark parliamentary meeting days.

Figure 1 – Climate-skeptic narratives in the 9th European Parliament

[2] The European Green Deal was legislated at the European level and comprised major climate initiatives which were either directly applicable (regulations) or had to be implemented by member states in their own laws (directives). We describe this legislation in greater detail in Supplemental Appendix D.

Our empirical analysis provides three sets of main results. First, at the speech level, there are strong signs of entanglement: 91% of speeches on climate, 97% on immigration, and 83% on digitization do not discuss trade-offs between economic and non-economic effects, among speeches relevant on both dimensions. Second, after Russia's invasion of Ukraine, politicians shift the whole climate narrative bundle in line with entanglement. This holds true when only looking at the within-MEP variation in narratives. After Russia's invasion, the share of speeches with a negative economic narrative increases by 2.8 p.p., and the share questioning whether climate policy can mitigate climate change by 0.8 p.p., which are very large effects given the mean dependent variables of 5.0% and 1.9%, respectively. The shift differs by party as the model predicts: MEPs from the far-right ECR and ID party groups drove narrative shifts by being more likely to speak on the topic of climate policy and more likely to voice climate-skeptic narratives. Third, using roll-call data on legislative votes, we show that changes in climate-skeptic narratives closely relate to changes in legislative votes on climate policy.

Our paper contributes to the theoretical and empirical literature on narratives. In terms of theory, an emerging literature defines narratives as subjective causal models (Eliaz and Spiegler, 2020), while other contributions study narratives using a different terminology (Schwartzstein and Sunderam, 2021; Levy et al., 2022; Izzo et al., 2023; Aina, 2023).[3] When modeling an endogenous choice of narratives, these approaches assume that agents select narratives according to data fit. We explore the complementary assumption that some voters are motivated reasoners who in principle accept a range of narratives; equilibrium narratives are then pinned down by political incentives (see Brzezinski, 2023; Besley and Brzezinski, 2025). Our paper also connects to theoretical work that explains why climate policy is often inadequate, including commitment problems and contracting frictions (Acemoglu et al., 2012; Ulph and Ulph, 2013; Helm et al., 2003; Acemoglu and Rafey, 2023; Battaglini and Harstad, 2020), or incomplete contracting in international treaties (Harstad, 2012, 2016; Battaglini and Harstad, 2016). We add a narrative channel to this theoretical literature.

In terms of the empirical literature, a large body of research is devoted to understanding the views of citizens on climate policy, which in our framework corresponds to the demand side of climate narratives. A recent large-scale survey across 20 countries is Dechezleprêtre et al. (2025), which among other insights concludes that perceptions of climate policy are strongly determined by self-interest considerations of households—in line with the assumptions of our framework (see also Brannlund and Persson, 2012).[4] Djourelova et al. (2024)

[3] This literature is also highly related to a literature that takes mindsets, ideology and rationales seriously (Stantcheva, 2024; Chinoy et al., 2024; Bénabou, 2008; Bursztyn et al., 2023a,b).

[4] Recent literature reviews on attitudes toward climate policy include Drews and Van den Bergh (2016)

show that after natural disasters, liberals and conservatives further diverge in beliefs about climate change. We complement this analysis of the political demand side by investigating the political supply side: climate narratives voiced by politicians. Our finding of increased climate skepticism among MEPs in the EU is in line with Bombardini et al. (2025), who estimate, for the US, that both Republicans and Democrats are likely to shift toward the political right on climate policy in the future. We offer new data on speeches made in the EU, complementing databases for the US (Gentzkow and Shapiro, 2010; Quinn et al., 2010; Ash and Hansen, 2023) and national parliaments (e.g. Rauh and Schwalbach, 2020). In terms of methodology, we contribute to the emerging literature that leverages LLMs for text analysis (see, e.g., Arold et al., 2024).

The remainder of the paper is structured as follows. Section 2 presents our theoretical framework. Section 3 describes our data collection and shows descriptive evidence for narrative entanglement across different policy domains. Section 4 presents our empirical findings. Section 5 concludes by discussing the broader significance of narrative entanglement in political discourse and the implications for theory and policy. Appendix A contains our proofs, and Supplemental Appendices B–F contain additional material.

2. Theory

2.1. Set-up

Parties and the supply of narratives. Two office-motivated parties indexed by $J \in \{L, R\}$ (for *Left* and *Right*) compete for the votes of a unit mass of voters in order to win an election. They offer different exogenously given tax rates $t_L > t_R$ and choose a policy $x_J \in \{0, 1\}$ and a narrative $\nu_J = (\lambda_J, \rho_J) \in \mathcal{N} \subset \mathbb{R}^2$ that describes the economic effects via λ_J and non-economic effects via ρ_J of policy x_J. The narrative space is constrained by

$$\mathcal{N} \;=\; [\underline{\lambda}, \overline{\lambda}] \times [\underline{\rho}, \overline{\rho}], \tag{1}$$

with $\underline{\lambda} < \overline{\lambda}$ and $\underline{\rho} < \overline{\rho}$, and is assumed to contain the truth $\nu^\star = (\lambda^\star, \rho^\star)$.[5]

Voter utility and the demand for narratives. There is a continuum of voters with each voter i characterized by an idiosyncratic private economic cost from the policy $\alpha_i \sim F$, where F is atomless and symmetric with mean and median $\bar{\alpha} > 0$.

Voter utility is influenced by a party platform (x_J, t_J) together with a voter i's narrative $\nu_i \equiv (\lambda_i, \rho_i)$ and private costs α_i. This induces perceived economic outcomes $y(\alpha_i, x_J, t_J; \lambda)$

and Whitmarsh and Capstick (2018).

[5] The constraints on this narrative space can be interpreted in various ways: narratives that cannot be rejected by existing data; parameters that are only set-identified; or narratives that fall within the "Overton Window" of what is politically acceptable.

and perceived non-economic outcomes $z(x_J;\rho)$. Voter i's utility is the sum of the two:

$$u(\alpha_i, x_J, t_J; \nu) \equiv y(\alpha_i, x_J, t_J; \lambda) + z(x_J; \rho). \tag{2}$$

We assume that y, z are twice continuously differentiable, and make the following assumption on their functional form:

Assumption 1.

(i) ***Ordered narratives.*** *Higher values of λ and ρ make policy implementation weakly more attractive and non-implementation weakly less attractive:*

$$\partial_\lambda y(\alpha, 1, t; \lambda) \geq 0 \geq \partial_\lambda y(\alpha, 0, t; \lambda), \qquad \partial_\rho z(1; \rho) \geq 0 \geq \partial_\rho z(0; \rho),$$

with one inequality strict in each narrative dimension.

(ii) ***Linear, α_i-independent economic narrative.*** $\partial^2_{\lambda\lambda} y = \partial^2_{\alpha\lambda} y = 0$.

(iii) ***Costs.*** $\partial_\alpha y(\alpha, 0, t; \lambda) = 0$; $\partial_\alpha y(\alpha, 1, t; \lambda) < 0$; $\lim_{\alpha \to \mp\infty} y(\alpha, 1, t; \lambda) = \pm\infty$.

(iv) ***Progressive taxation.*** $\partial^2_{\alpha t} y(\alpha, 1, t; \lambda) > 0$.

Condition (i) says that narratives can be ordered such that higher values of λ and ρ imply more positive effects from implementing relative to not implementing the policy for the economic and non-economic dimensions, respectively. Condition (ii) implies a single-crossing property in α_i for any policy pair, and hence that the winning probability only depends on the median voter's preference (see Lemma A.1). Condition (iii) says that income only depends on private costs if the policy is implemented, in which case it is decreasing in them, and that y has full range in α, which ensures the existence of relevant cut-offs. Condition (iv) says that when the policy is implemented, lower-income citizens (i.e., citizens with higher α_i) more strongly prefer higher taxes.

Bayesian voters, motivated voters, and the truth. We assume that voters may adopt narratives in one of two ways, which is determined by nature after parties supply policies and narratives, but before voting occurs. With probability μ,[6] parties face *Bayesian voters*. They have some prior belief $\nu^{\text{PM}} \in \mathcal{N}$, do not know the truth $\nu^\star \in \mathcal{N}$ and cannot find it out on their own, but can tell whether a politician speaks the truth or lies. For the analysis to be interesting, we make the mild assumptions that $\nu^{\text{PM}} \neq \nu^\star$ as well as that $\nu^{\text{PM}}, \nu^\star$ do not take the knife-edge values where parties are exactly indifferent between voters adopting ν^{PM} or $\nu^\star$.[7]

[6] We could instead assume that there is a *proportion* μ of Bayesian voters rather than a *probability* μ. This would give exactly the same results if there was no heterogeneity in α_i. With heterogeneity in private costs, the model is analytically simpler to solve with this probabilistic set-up.

[7] The latter implies that, for any policy combination, the median voter's utility difference between Left and Right is not exactly the same under ν^{PM} as under $\nu^\star$, which only rules out knife-edge cases.

From the above it follows that Bayesian voters retain their prior belief unless at least one party supplies the truth:

$$\nu^B(\nu_L, \nu_R) = \begin{cases} \nu^\star & \text{if } \nu_L = \nu^\star \text{ or } \nu_R = \nu^\star, \\ \nu^{\text{PM}} & \text{otherwise.} \end{cases} \tag{3}$$

This is a stark way to model Bayesian voters that delivers sharp results. Other assumptions would lead to a qualitatively similar analysis but at higher analytical burden.[8]

With probability $1-\mu$ voters are not Bayesian. Then, they adopt narratives not based on accuracy, but based on whether they want to believe in them: their perceptions are driven by motivated reasoning. In particular, after observing parties' policy-narrative combinations (x_L, t_L, ν_L) and (x_R, t_R, ν_R), a *motivated voter* i adopts the narrative $\nu_i \in \{\nu_L, \nu_R\}$ that maximizes their maximally attainable utility across the two parties' policy platforms:[9]

$$\nu_i = \arg \max_{\nu \in \{\nu_L, \nu_R\}} \max \Big\{ u(\alpha_i, x_L, t_L; \nu),\ u(\alpha_i, x_R, t_R; \nu) \Big\}. \tag{4}$$

This assumption that voters maximize prospective utility has a long tradition in psychology (Kunda, 1990) and has more recently been adopted in economic models (Bénabou, 2015). In the context of political economy, our approach is similar to that of Brzezinski (2023) or Besley and Brzezinski (2025). We differ from the latter by assuming that motivated voters choose their beliefs by looking at the maximally attainable utility, rather than maximizing expected utility. In practice, this eliminates some ranges where multiple equilibria may operate through the channel that voters' beliefs are driven in part by their expected probability of which party will win.[10] We abstract from the notion that voters may trade off motivated reasoning versus accuracy of beliefs, as formalized by Little (2019) and Little et al. (2022). Instead, in Section 2.4 we consider how incentives change once politicians face both motivated reasoners and Bayesian voters such that $\mu \in (0, 1)$.

Voting. Voting occurs probabilistically. An aggregate party preference shock δ with strictly increasing cumulative distribution function G symmetric about 0 affects voter choice. Voter i votes for L if and only if, under voter i's belief ν_i, they weakly prefer L's platform after the shock:

$$u(\alpha_i, x_L, t_L; \nu_i) - \delta \ \geq \ u(\alpha_i, x_R, t_R; \nu_i). \tag{5}$$

[8] For instance, rather than assuming Bayesians only update when hearing the truth, we could assume that narratives closer to the truth are more likely to be adopted. This would capture the same trade-off for parties between truthful narratives or entangled narratives, but with a more taxonomic analysis.

[9] We assume that voters choose a narrative at random in the knife-edge cases of equal utilities, but the tie-breaking rule is inconsequential for our analysis.

[10] Such cases are explored in Besley and Brzezinski (2025) for the context of a dynamic model of narratives applied to deregulation policy narratives with only economic outcomes. They give conditions under which policy equilibria are unique.

Winning probability. Given Assumption 1.(ii), there is monotonicity of voting choices for any combination of policies, both under motivated voters and Bayesian voters (see Lemma A.1 in the Appendix for a formal statement and proof). Hence, the winning probability of party $J \in \{L, R\}$ is a μ-weighted average between winning under Bayesian voters and motivated reasoners. Define the median voter's utility difference between the two parties' platforms under narrative ν as

$$\Delta u(\bar{\alpha}, x_L, x_R, t_L, t_R; \nu) \;\equiv\; u(\bar{\alpha}, x_L, t_L; \nu) \;-\; u(\bar{\alpha}, x_R, t_R; \nu). \tag{6}$$

Let ν^B denote the common Bayesian belief from (3) and ν^M the narrative adopted by the median motivated voter under (4). Given the voting rule in (5), the winning probability of party J is then

$$\pi_J \;=\; \mu\, G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^B)\big) \;+\; (1-\mu)\, G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^M)\big). \tag{7}$$

Timeline. The timeline is as follows:

1. Parties $J \in \{L, R\}$ simultaneously choose policy–narrative pairs (x_J, ν_J).
2. Voters form beliefs according to (3) if Bayesian or to (4) if motivated reasoners.
3. The aggregate shock δ is realized and voting occurs as in (5).
4. The winning party implements the policy x_J that it announced.

Equilibrium. We characterize subgame-perfect equilibria consisting of party strategies $\{(\hat{x}_L, \hat{\nu}_L), (\hat{x}_R, \hat{\nu}_R)\}$, the Bayesian-regime equilibrium belief $\hat{\nu}^B$, and motivated voters' narrative choices $\{\hat{\nu}_i\}_i$ such that: given voters' beliefs in each regime, each party J chooses $(\hat{x}_J, \hat{\nu}_J) \in \{0, 1\} \times \mathcal{N}$ to maximize its winning probability π_J as in (7); given parties' choices, $\hat{\nu}^B$ satisfies (3) and each motivated voter i's belief $\hat{\nu}_i$ satisfies (4).

We now illustrate our framework through two specific policy applications.

Example 1: Climate policy. Take $x = 1$ as climate action (e.g. carbon-related regulation) and $x = 0$ as inaction. Let

$$u(\alpha_i, x_J, t_J; \lambda, \rho) = \underbrace{(1 - t_J)\big(w - x_J \alpha_i\big)}_{\text{post-tax income}} + \underbrace{t_J\big(w - x_J(\bar{\alpha} - \lambda)\big)}_{\text{public goods}} - \underbrace{(1 - x_J)\rho}_{\text{environmental damages}}, \tag{8}$$

where w is the common pre-tax wage; α_i is voter i's idiosyncratic private economic cost from climate policy; $\rho \geq 0$ captures the environmental damage avoided by climate policy; and λ captures an aggregate spillover of climate policy on the tax base and hence on public goods. For the latter, $\lambda = 0$ is the no-spillover benchmark, $\lambda < 0$ is a "macro-drag" narrative, and $\lambda > 0$ is a "green-growth" narrative.[11] The perceived economic effects of implementing the

[11] Formally, individual taxable income is $w - \alpha_i x_J$, so public-goods financing is $\int_{-\infty}^{\infty} t_J(w - x_J\alpha)\, dF(\alpha) = t_J(w - x_J\bar{\alpha})$ under no spillovers. With the spillover term from the economic narrative, total public goods

policy relative to not doing so are $-(1-t_J)\alpha_i - t_J(\bar{\alpha}-\lambda)$, and the perceived environmental effect is simply ρ. The former effect shows a positive interaction between taxes and economic narratives: higher taxes insure voters against private policy costs and make the aggregate-growth narrative more important.

Example 2: Immigration policy. Take $x = 1$ as tighter immigration policy and $x = 0$ as looser. Let

$$u(\alpha_i, x_J, t_J; \lambda, \rho) = \underbrace{p(x_J;\lambda)(1-t_J)w}_{\text{income if working}} + \underbrace{\big(1-p(x_J;\lambda)\big)u_J}_{\text{income if unemployed}} + \underbrace{c\,b_J - x_J\alpha_i(1-t_J)w}_{\text{healthcare}} + \underbrace{x_J\rho}_{\text{culture}},$$

where the probability of working is $p(x_J;\lambda) = \frac{1}{2} + x_J\lambda$, with $|\lambda| < \frac{1}{2}$. The unemployment transfer u_J and healthcare expenditure b_J are pinned down by the government budget constraint with equal allocation such that $(1-p)u_J = b_J = \frac{1}{2}p\,t_J w$. The constant $c \in (0,1)$ is a state efficiency wedge. Tighter immigration affects employment probability via λ, culture via ρ, and personal out-of-pocket healthcare cost via $\alpha_i(1-t_J)w$ for care that immigrants would have provided publicly under loose policy (e.g. nursing-home staffing). In contrast to the climate application, the interaction between t and λ is negative: higher taxes dilute the perceived labor-market gain from tightening, because the welfare-state efficiency wedge $c < 1$ means tax-financed expenditure delivers less than it raises.

2.2. Bayesian benchmark

We start by considering a simple benchmark where voters are Bayesian:

Proposition 1. *Assume that all voters are Bayesian with certainty, $\mu = 1$. In any equilibrium, at least one party tells the truth, and both parties implement policies that are optimal for the median voter.*

For any policy choice, one party will strictly benefit from telling the truth given the zero-sum nature of the game. In equilibrium, all Bayesian voters will hear and adopt the truth, and both parties will implement the policy that maximizes the utility of the median voter. This may not result in policy convergence, because the parties' different tax stances may imply that different policy choices are optimal for the median voter; for instance, in our climate policy application, there is a positive interaction between climate policy and high taxes, which means that for some values of $\bar{\alpha}$ the median voter may prefer the Left party to implement the policy but the Right party not to do so. But overall, parties will offer the policies that are objectively optimal for the median voter. We show below that this is not the case under motivated reasoning.

---

are instead perceived as $t_J(w - x_J(\bar{\alpha}-\lambda))$, which is the same as no spillovers under $\lambda = 0$, but smaller for $\lambda < 0$ and larger for $\lambda > 0$.

2.3. Motivated voters

We now study the situation where all voters are motivated reasoners, such that $\mu = 0$. This gives the most concise overview of our narrative entanglement result.

Narrative Entanglement. We start our analysis by defining *entangled narratives*:

$$\underline{\nu} \equiv (\underline{\lambda}, \underline{\rho}) \tag{9}$$

$$\overline{\nu} \equiv (\overline{\lambda}, \overline{\rho}) \tag{10}$$

We refer to these narratives as entangled for two reasons. First, they intertwine the different policy dimensions: the narrative $\underline{\nu}$ frames the policy as "bad" in all dimensions while the narrative $\overline{\nu}$ frames it as "good" throughout. Second, as we show below, the dimensions are entangled in the sense that, even though they are factually separate, in equilibrium a change in one is accompanied by a change in the other.

The following proposition states that parties who face only motivated voters have strategic incentives to offer entangled narratives:

Proposition 2. *Assume that voters are motivated reasoners with certainty, $\mu = 0$. Then, in any equilibrium, it is optimal for parties to offer the entangled narratives $\underline{\nu}$ or $\overline{\nu}$.*

Narrative entanglement arises because parties have two strategic aims when competing for motivated voters: (1) to offer narratives and policies that match the median voter's policy preference; and (2) to maximize their chance of winning against a party with the opposing policy. Entangled narratives simultaneously support a party's own policy *and* undermine an opponent who offers an opposing policy. Whenever parties offer differing policies, entangled narratives are strictly better than other narratives, but they remain a best response when parties offer the same policies.

The proof of Proposition 2 shows that incentives to win cause one party to supply the median voter's preferred entangled narrative in equilibrium, which makes the other party indifferent between supplying any other narratives since it will not change the median voter's choice. As a result, in principle there is always a continuum of narratives that this party could adopt in equilibrium.

To avoid considering all possible weak equilibria and to concentrate the equilibrium analysis, we make the intuitive assumption that when indifferent between narratives, parties use the one that makes their policy look best: they combine $x_J = 0$ with $\nu_J = \underline{\nu}$, and $x_J = 1$ with $\nu_J = \overline{\nu}$. We can then characterize equilibria as a function of the median voter's private costs, where we give precise cutoff values in the proof to the proposition in the Appendix:

Proposition 3. *Assume that all voters are motivated reasoners with certainty, $\mu = 0$. Then there exist thresholds $\tilde{\alpha}_1 \leq \tilde{\alpha}_2$ that characterize equilibrium party policies and narratives:*

1. *The unique equilibrium is pro-policy pooling,* $(\hat{x}_L, \hat{x}_R, \hat{\nu}_L, \hat{\nu}_R) = (1, 1, \overline{\nu}, \overline{\nu})$*, iff costs are low,* $\bar{\alpha} < \tilde{\alpha}_1$*. All voters believe in the same narrative,* $\hat{\nu}_i = \overline{\nu}$*.*

2. *The unique equilibrium is anti-policy pooling,* $(\hat{x}_L, \hat{x}_R, \hat{\nu}_L, \hat{\nu}_R) = (0, 0, \underline{\nu}, \underline{\nu})$*, iff costs are high,* $\bar{\alpha} > \tilde{\alpha}_2$*. All voters believe in the same narrative,* $\hat{\nu}_i = \underline{\nu}$*.*

3. *There is a unique pure-strategy separating equilibrium with policy and narrative divergence, or a generically unique mixed-strategy equilibrium, iff costs are intermediate,* $\bar{\alpha} \in (\tilde{\alpha}_1, \tilde{\alpha}_2)$*. When party narratives diverge, voters disagree about narratives such that* $\hat{\nu}_i = \overline{\nu}$ *for voters below an* α_i*-threshold and* $\hat{\nu}_i = \underline{\nu}$ *for voters above it.*

The first two cases laid out in the proposition are simple: when the costs to the median voter of the policy are very low, both parties support it; when the costs are very high, both parties oppose it. The third case of intermediate costs gives rise to a polarized equilibrium, in pure or mixed strategies.

Polarized equilibria. For intermediate costs, any pure-strategy equilibrium is polarized in policies—but such equilibria may not exist. Intuitively, a precondition for existence is that cross-party differences in tax stances matter sufficiently. If they do, it can be optimal for one party to implement the policy and for the other to oppose implementation. If instead tax differences do not matter, then the game is nearly symmetric and only a mixed-strategy equilibrium exists.

For the climate policy application, the necessary and sufficient condition for a polarized equilibrium in pure strategies is simply given by

$$(t_L - t_R)\overline{\lambda} > \overline{\rho} - \underline{\rho}, \tag{11}$$

which makes it explicit that tax differences need to be large enough for such an equilibrium to arise.[12]

Polarized equilibria capture the case where parties offer opposing entangled narratives $\underline{\nu}$ and $\overline{\nu}$, and voters disagree on what they believe. The proof of Proposition 3 shows that voters above a threshold value of costs will adopt the narrative $\underline{\nu}$, since that maximizes their maximally-achievable utility that occurs when the policy is not implemented. For similar reasons, voters below that cost threshold will adopt the opposing $\overline{\nu}$. Hence, heterogeneity in private costs leads to heterogeneity in the equilibrium narrative that voters adopt.

Entangled narratives and sub-optimal policy. Since the truth is not revealed in equilibrium, there is no guarantee that equilibrium policies in Proposition 3 will coincide with

[12] This is derived in the proof to Corollary 1; see also Supplemental Appendix B.2. Corollary B.1 in Supplemental Appendix B.1 gives sufficient conditions for polarized equilibria in the general theory and discusses the intuition in greater detail.

optimal policies from the benchmark Bayesian case in Proposition 1. The degree to which these two coincide will depend on the distance between the truth $\nu^\star$ and the equilibrium narrative $\overline{\nu}$ or $\underline{\nu}$. Even if $\nu^\star$ is close to one of the two entangled narratives, it is then far away from the other, and hence for some cost values the difference in equilibrium narratives will be large. By implication, equilibrium policy will in such cases generally be different.

The narrative accelerator in entangled narratives. In the Bayesian case explored in Proposition 1, voters always learn and believe in the true narrative $\nu^\star$, which is independent of costs α_i. As α_i goes up, $\nu^\star$ stays the same, but the evaluation of the policy becomes less favorable. For sufficiently high median costs $\bar{\alpha}$, the median voter will prefer parties not to implement the policy.

In the case of motivated reasoners explored in Proposition 3, an increase in costs changes not only the median voter's policy preferences, but also the narratives that they demand from parties. In equilibrium, parties will find it optimal to change the narrative, at which point the perceived policy benefit will change discontinuously. At that point, policy changes will occur abruptly. This shows a key aspect of "entanglement": even though economic and non-economic effects may be factually unrelated, in equilibrium narratives about them will move together in response to an increase in costs.

Figure 2 depicts this mechanism by contrasting party payoffs when beliefs are fixed (panel (a)) with the case when narratives adjust endogenously (panel (b)), for a particular parametrization of the model application to climate policy.[13] The figure shows the Left's payoff from different policies as a function of median voter costs, where positive values of the payoff denote an electoral advantage to the Left. In panel (a), where narratives are fixed, payoffs change smoothly in costs. In contrast, in panel (b), around the cutoff $\tilde{\alpha}_2$ there is a discontinuous jump in party payoffs. This is because for $\bar{\alpha} > \tilde{\alpha}_2$, the median voter demands a new narrative, which will in equilibrium be supplied by parties and lead to a jump in the assessed policy benefits. As a result, in panel (b) there exists an equilibrium where neither party supports the policy, which does not exist under exogenous narratives as shown in panel (a) for this range of median costs.

2.4. Motivated reasoners and Bayesians

We now study the case where voters may be motivated reasoners or Bayesians. For there to be a tension between speaking the truth and using entangled narratives, we assume that the truth is in the interior: $\lambda^\star \notin \{\underline{\lambda}, \overline{\lambda}\}$ and $\rho^\star \notin \{\underline{\rho}, \overline{\rho}\}$.

[13] We use the following parameter values for the climate policy application: $t_L = 0.75$, $t_R = 0.25$, and additionally for panel (a), we fix $(\rho, \lambda) = (1, 0.75)$; for panel (b), we use $\overline{\rho} = 1, \underline{\rho} = 0.75$, $\underline{\lambda} = -0.75$ and $\overline{\lambda} = 0.75$.

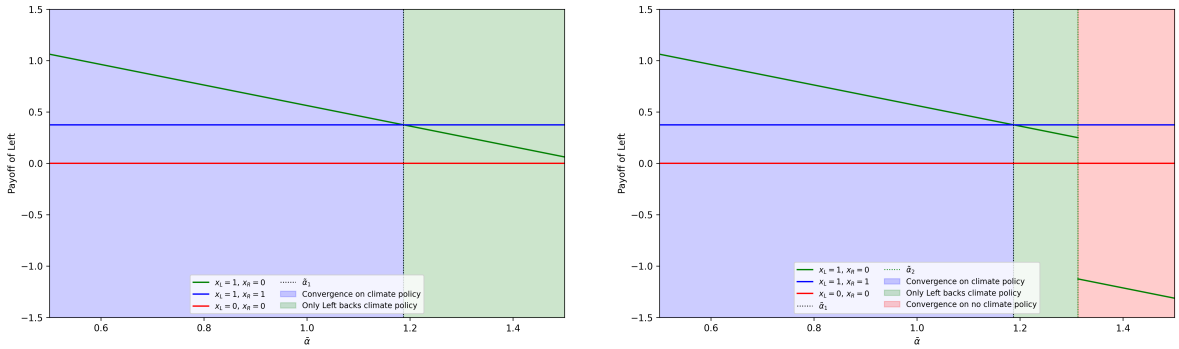

(a) Payoff of Left under exogenous narratives **(b)** The narrative accelerator

Note: Climate policy utility function (8) with the following parameter values: $t_L = 0.75$, $t_R = 0.25$. Panel (a): $(\rho, \lambda) = (1, 0.75)$. Panel (b): $\overline{\rho} = 1$, $\underline{\rho} = 0.75$, $\overline{\lambda} = 0.75$, $\underline{\lambda} = -0.75$.

Figure 2 – Party payoffs, under fixed narratives versus endogenous narratives

We start our analysis with some helpful definitions. Fixing a vector of policies (x_L, x_R), by the zero-sum nature of political competition, one party strictly prefers that Bayesian voters learn the truth and the other that they stay at their prior. Denote the former by $J^B(x_L, x_R)$. At the same time, one party strictly prefers that median-cost motivated reasoners adopt entangled narratives and the other that they adopt the truth. Denote the former by $J^M(x_L, x_R)$. Formally, using the median-voter utility difference Δu from (6), and defining $\hat{\nu}^{M,0}(x_L, x_R) \in \{\underline{\nu}, \overline{\nu}\}$, the equilibrium narrative adopted by the median voter in Proposition 3 under (x_L, x_R):

$$J^M(x_L, x_R) = \arg\max_{J \in \{L,R\}} \Big[\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \hat{\nu}^{M,0}) - \Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^\star)\Big],$$

$$J^B(x_L, x_R) = \arg\max_{J \in \{L,R\}} \Big[\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^\star) - \Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^{\text{PM}})\Big].$$

We now examine whether party narrative and policy choices that sustained an equilibrium in Proposition 3, under motivated voters only, survive the introduction of Bayesian voters. We are particularly interested in whether any party has an incentive to reveal the truth. The next proposition shows that this depends on whether, at the baseline policy profile, the same party benefits from revealing the truth to Bayesians and from providing the entangled narrative to motivated reasoners:

Proposition 4. *Consider a baseline equilibrium profile from Proposition 3 with party choices* $(\hat{x}_J^0, \hat{\nu}_J^0)$ *for* $J \in \{L, R\}$ *and median motivated-voter narrative* $\hat{\nu}^{M,0} \in \{\hat{\nu}_L^0, \hat{\nu}_R^0\}$. *Assume a positive probability that voters are Bayesian,* $\mu \in (0, 1)$, *and denote the new equilibrium strategies by* $(\hat{x}_J^\mu, \hat{\nu}_J^\mu)$ *for parties and* $\hat{\nu}^{B,\mu}, \hat{\nu}^{M,\mu}$ *for median voters.*

1. *If* $J^M(\hat{x}_L^0, \hat{x}_R^0) = J^B(\hat{x}_L^0, \hat{x}_R^0)$, *then even for arbitrarily large* $\mu < 1$, *there exist parameter*

values such that: equilibrium policies coincide with the baseline, $\hat{x}_J^\mu = \hat{x}_J^0$; the truth is never revealed and Bayesians remain at their prior, $\hat{\nu}^{B,\mu} = \nu^{PM}$; and the median motivated-reasoner narrative is unchanged relative to baseline, $\hat{\nu}^{M,\mu} = \hat{\nu}^{M,0}$.

2. *If $J^M(\hat{x}_L^0, \hat{x}_R^0) \neq J^B(\hat{x}_L^0, \hat{x}_R^0)$, then equilibrium policies coincide with the baseline if μ is small enough, $\hat{x}_J^\mu = \hat{x}_J^0$. In that case, J^B reveals the truth and Bayesians adopt it, $\hat{\nu}^{B,\mu} = \nu^\star$, while the median motivated-reasoner narrative is unchanged relative to baseline, $\hat{\nu}^{M,\mu} = \hat{\nu}^{M,0}$.*

If the same party benefits from revealing the truth to Bayesians and from providing the entangled narrative to motivated reasoners (Case 1), equilibria can remain substantively the same as with only motivated reasoners, with no truth revelation, the same equilibrium policies, and the same entangled median voter beliefs of the motivated reasoner. This can occur even for arbitrarily large $\mu \in (0, 1)$. The intuition is that, if $J^B = J^M$, then this "advantaged" party trades off net benefits of the optimal entangled narrative against net benefits of the truth. The advantaged party J will prefer to use the entangled narrative if:

$$\begin{aligned}(1-\mu)\big[G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \hat{\nu}^{M,0})\big) - G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^\star)\big)\big] \\ \geq\ \mu\big[G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^\star)\big) - G\big(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^{\mathrm{PM}})\big)\big]. \qquad (12)\end{aligned}$$

The left-hand side is the motivated-reasoner winning probability that the advantaged party forfeits by deviating from $\hat{\nu}^{M,0}$ to $\nu^\star$, weighted by the probability that voters are motivated reasoners; the right-hand side is the Bayesian winning probability it would gain from revealing the truth, weighted by the probability that voters are Bayesian.

Condition (12) is not far-fetched if one assumes the prior beliefs of Bayesian voters to be well calibrated: as $\nu^{\mathrm{PM}} \to \nu^\star$, the right-hand side goes to zero, and the inequality is always satisfied. The expression does depend on the proportion of Bayesian voters μ in the expected way, as it is more restrictive for larger μ, and impossible to satisfy for $\mu = 1$. But for any arbitrarily large $\mu < 1$ it can be satisfied.

Consider now the second case of Proposition 4, where no single party has an advantage from truth-telling and from entangled narratives, i.e., $J^M \neq J^B$. Policies remain the same as in Proposition 3 if the proportion of Bayesian voters is small enough, so the policy preferences of motivated reasoners dominate. Then, J^B will offer the truth in any equilibrium.

Equilibrium narratives may diverge even if policies converge. Indeed, the proof to Proposition 4 shows that $\hat{x}_L^\mu = \hat{x}_R^\mu = 1$ is an equilibrium only if $\hat{\nu}_L^\mu \neq \hat{\nu}_R^\mu$. In this equilibrium, both parties supply the same policy, but one voices more skeptical narratives than the other—to prevent the truth from being told and keep Bayesian voters in the dark.

3. Narrative data and descriptive evidence

3.1. Sources and overview

We collect all speeches from the 9th European Parliament, covering the period between 2 July 2019 and 25 April 2024, the first and last plenary sitting days of the legislative term. We draw the speeches directly from transcripts published by the European Parliament, giving us the full speech in the original language, the speech-ID, the MEP name and MEP-ID, the European party group of the MEP, and the country.[14] We add information on each MEP's national party from Parltrack's MEP membership histories.[15]

The total number of speeches in the 9th European Parliament was 64,250. Our focus is on substantive, orally delivered speeches by MEPs. Hence, we exclude from the total corpus speeches submitted in writing rather than orally delivered, purely procedural speeches by the presiding officer, and speeches delivered by non-MEPs (representatives of the Commission, the Council, and invited guests). This leaves us with 48,239 speeches delivered by 848 MEPs. Among these, the average number of speeches per MEP is 56.9, with a median of 46.

Each MEP was part of one of the following eight transnational party groups (with political classification in parentheses):[16] The Greens[17] (center-left to left-wing), The Left[18] (left-wing), S&D (center-left), Renew (center to center-right), EPP (center-right), ECR (right-wing with far-right factions), and ID (far-right), plus NI independents that were not part of any of the other factions. Supplemental Appendix Figure C.2 shows that the share of seats across the eight party groups (panel a) is roughly in proportion to the share of speeches (panel b). Our Parltrack bridge allows us to map observations also into national-level parties, of which there are 259 in our dataset. Our data covers 28 countries, including the United Kingdom, which was present at the start of our period of study but whose MEPs left on 31 January 2020 following Brexit.[19]

3.2. Narrative classification across three policy dimensions

We classify narratives on three policy domains that were particularly salient during the 9th European Parliament: climate policy, digitization policy, and immigration policy. Among

[14] Speeches are retrieved from https://www.europarl.europa.eu/doceo/; the list of sitting dates is obtained from the parliamentary Open Data Portal at https://data.europarl.europa.eu/.

[15] Parltrack's MEP records are available at https://parltrack.org/dumps.

[16] This is taken from the main page on each political group from Wikipedia.

[17] The official name is Greens/EFA, but for simplicity we refer to them as The Greens.

[18] This group was called GUE/NGL until October 2021, and after that The Left. For simplicity, we refer to them as "The Left" throughout.

[19] None of our results depend on the 641 speeches (1% of total) made by British MEPs, and they do not impact our analysis in Section 4 since they had left by the time of Russia's invasion.

these, climate policy was a particularly strong focus, given the European Green Deal initiative of the governing coalition. The European Climate Law, passed in June 2021, legally adopted the 2050 climate-neutrality target and the binding 55% emissions-reduction goal for 2030. Legislation to reach that target was passed in subsequent years.[20]

Digitization policy was also a major topic, in particular the data protection framework whose enforcement and extension shaped parliamentary debates over the Data Governance Act (April 2022) and the Data Act (November 2023). The very end of the parliamentary session saw the passing of the Artificial Intelligence Act (March 2024).

Immigration policy remained a highly discussed topic especially on the far-right, as we will show further below, but legislation on it was more limited compared to the other two topics. The main initiative was the New Pact on Migration and Asylum (April 2024), bundling several reforms including reforms to asylum procedures.

Narrative classification. To identify parliamentary speeches made on these three domains, we use LLMs via OpenAI's API service to classify each speech separately for each policy domain and, within domain, separately on each of two narrative dimensions. For climate policy, we ask about environmental and economic narratives; for immigration, about cultural and economic narratives; and for digitization, about security and economic narratives. Overall, each speech is independently classified six times, on two narrative dimensions for each policy domain.

Our classification strategy closely follows our model, where narratives describe causal effects of a policy. Accordingly, we structure our prompt to specify the causal parameter implied by speeches, if they are related to the topic of interest. To describe the prompt in a concise way, define the placeholder POLICY as one of {climate policy, immigration policy, digitization policy}; and OUTCOME and SECOND OUTCOME as different elements from {the economy, climate change} for climate policy, {the economy, culture} for immigration policy, and {the economy, security} for digitization policy. Each of our six prompts then follows the main structure below, replacing the placeholders with the relevant input:[21]

> Classify the speech below according to its narrative about how POLICY affects OUTCOME, and provide a 1–2 sentence justification for your classification.
>
> **Classification categories:**

[20] The Fit for 55 package constituted a reform and extension of the Emissions Trading System (18 April 2023), the introduction of the Carbon Border Adjustment Mechanism (18 April 2023), tighter CO_2 standards phasing out new internal combustion engine vehicles by 2035 (14 February 2023), and revisions to the Renewable Energy Directive (12 September 2023) and Energy Efficiency Directive (11 July 2023).

[21] The prompt additionally has a "system prompt" that briefly specifies the AI's role and a part that details the output format. There are also slight variations across prompts depending on the policy domain; the exact prompts are reproduced in Supplemental Appendix Section E.

"tighter": This speech implies that tighter POLICY would have a positive effect on OUTCOME.

"no effect": This speech implies that POLICY has no effect on OUTCOME.

"looser": This speech implies that looser POLICY would have a positive effect on OUTCOME.

"unrelated": This speech is not related to the effect of POLICY on OUTCOME.

Classification rules:

Speeches whose primary topic is not POLICY should be classified as "unrelated", even if they contain one or two sentences about POLICY.

Exclusively focus on the link between POLICY and OUTCOME, not on the link between POLICY and SECOND OUTCOME or on the link between POLICY and other outcomes. If the speech only discusses the effect of POLICY on SECOND OUTCOME or on other outcomes, but not on OUTCOME, it should be classified as "unrelated".

Capture the speech's overall desired direction for POLICY based on its effect on OUTCOME, not its stance on any single measure. For example, a speech can criticise a particular policy instrument while overall arguing that tighter POLICY benefits OUTCOME, in which case it should be classified as "tighter". Or, a speech can endorse a specific initiative while overall arguing that POLICY is too costly for OUTCOME, in which case it should be classified as "looser".

There is no memory between individual classifications across narratives. Hence, entanglement cannot be an artifact coming from the LLM aiming to consistently classify narratives for the same speech. Instead, any entanglement will come despite independent classifications.

We used GPT-5-mini (version gpt-5-mini-2025-08-07) for classifying speeches on all dimensions. This is a smaller and more cost-effective version than the larger GPT-5 model. Related literature has found that GPT-5-mini classification is highly performant relative to the larger GPT-5 model, as described by Benoit et al. (2026), and we confirm this with our own cross-validation exercise reported below.

In contrast to Benoit et al. (2026), who deal with very long party manifestos and therefore resort to first writing English-language short summaries before classifying their texts, we classify speeches based on the untranslated, unsummarized version of the speech. With a median speech length of 194 words, we deal with comparatively short text documents that do not require summarizing before analysis. The advantage of our approach of applying LLMs to the original speech in its original language is that the LLM can then consider

subtleties that would potentially be lost in translation. Some benchmark tests confirm that using LLMs on the original-language text, rather than on a translation, performs better for most languages (Intrator et al., 2024).

Topic incidence. Figure 3 plots the 10-day moving average of the share of speeches relevant to each of the three policy domains across the 9th European Parliament, where a speech is relevant if it voices a narrative in at least one narrative dimension of that policy domain. The largest policy topic was climate, which on its own accounts for 18% of speeches on average. The second was digitization at 9%, and the third immigration at 4%. Only 1% of speeches were relevant to two or more of these three topics. Together, on average 30% of speeches were relevant to at least one of the three policy topics that we consider.

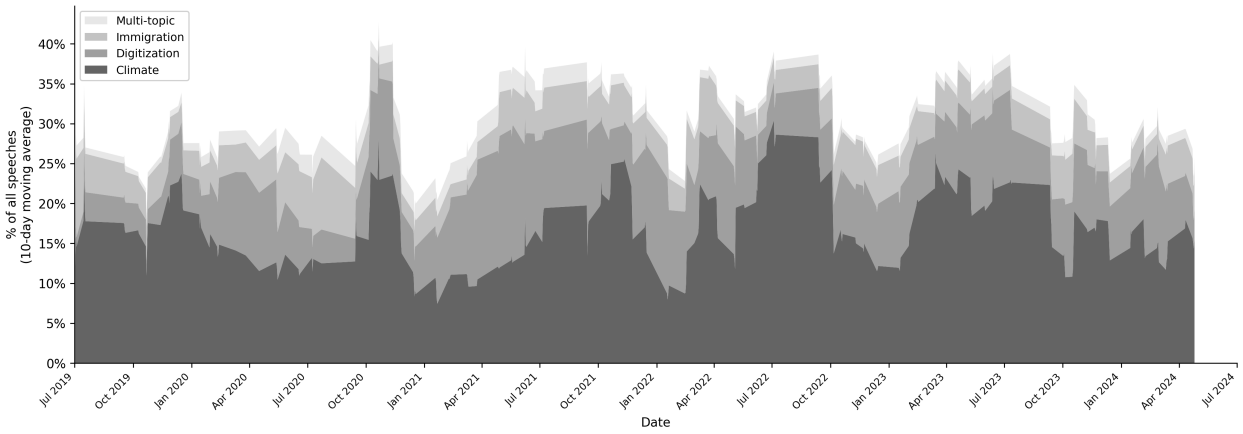

Note: 10-day moving average of the share of all parliamentary speeches with at least one narrative pertaining to climate, digitization, or immigration. "Multi-topic" captures speeches relevant to more than one domain.

Figure 3 – Climate, digitization, and immigration policy discussions

Overview of narrative classification. Table 1 shows the cross-tabulation between economic and non-economic narratives by policy domain. Among speeches that are on both economic and non-economic dimensions (i.e., not in "N/A"), there are strong signs of entanglement: most observations lie on the diagonal. Note also that the "no effect" category is the sparsest across all policy-narrative combinations. It only appears in significant numbers for environmental narratives, where intuitively "no effect" implies climate skepticism similarly to the "loose" category. To avoid a very sparse category, we group "no effect" together with "loose" in our later analysis.

Cross-validation of measures. We extensively cross-validate our classification methodology in two main ways, explained in detail in Supplemental Appendix F. First, we rerun our model classification with different LLMs, for a random subsample of 1,000 speeches

Table 1 – Cross-tabulation of economic and non-economic narratives, by policy

	Climate: environmental				Immigration: cultural				Digitization: security			
Economic	Tight	No effect	Loose	N/A	Tight	No effect	Loose	N/A	Tight	No effect	Loose	N/A
Tight	3,637	4	2	732	235	2	3	409	770	2	16	866
No effect	20	1	0	5	0	0	1	5	2	0	1	1
Loose	424	326	481	1,161	7	0	96	436	155	3	65	456
N/A	1,964	56	30	39,396	556	3	422	46,064	1,836	28	237	43,801

Note: Joint distribution of speeches across economic narratives (rows) and non-economic narratives (columns) within each policy domain; "N/A" contains speeches that use no narratives on that dimension.

for each of the six policy-narrative combinations. Second, we rerun our analysis using the baseline model but different prompts. Both approaches give very similar results, with full classification agreement of 95%–97% and substantial Cohen's κ levels between 0.72 and 0.82.

3.3. Narrative entanglement across policy domains

Figure 4 descriptively investigates entanglement by computing the share of "tight" speeches across party groups (excluding non-affiliates), separately on each dimension.[22] It then presents each party as a point in the two-dimensional %-tight space for each policy: climate in panel (a), digitization in panel (b), and immigration in panel (c). Entanglement would imply that the economic and non-economic dimensions should be very closely related, and that parties should tend to sit toward the corners.

The figure shows strong evidence for entanglement. All parties are close to the corners on climate, all except EPP close to the corners on immigration, and all except ID and ECR close to the "tight" corner on digitization. Supplemental Appendix Figure C.3 replicates this figure at the national-party level and at the MEP level, with observations again clustered around the 45° line and close to the corners.

We can also investigate narrative entanglement at the speech level instead of the party level. For this purpose, we look at speeches that discuss both the economic and non-economic effects of a policy. Speeches are classified as *entangled loose* if the narratives are "looser" or "no effect"[23] on both the non-economic and economic dimensions, *entangled tight* if they are "tighter" on both, and *not entangled* otherwise.

Following this methodology, we find that speech-level entanglement is highest in immi-

[22] The climate panel thus covers 6,945 speeches on the non-economic dimension and 6,793 on the economic dimension; digitization 3,115 and 2,337, respectively; and immigration 1,325 and 1,194, respectively.

[23] As shown in Table 1, "no effect" is by far the sparsest category, which is why we combine it with "looser". Collapsing "no effect" with "looser" is immaterial for immigration and digitization, where "no effect" appears very infrequently. For climate the collapse matters only on the environmental dimension, where 331 of 4,895 speeches argue that climate policy has no effect on climate change. Here, bundling "no effect" with "looser" is intuitive, since both imply climate skepticism.

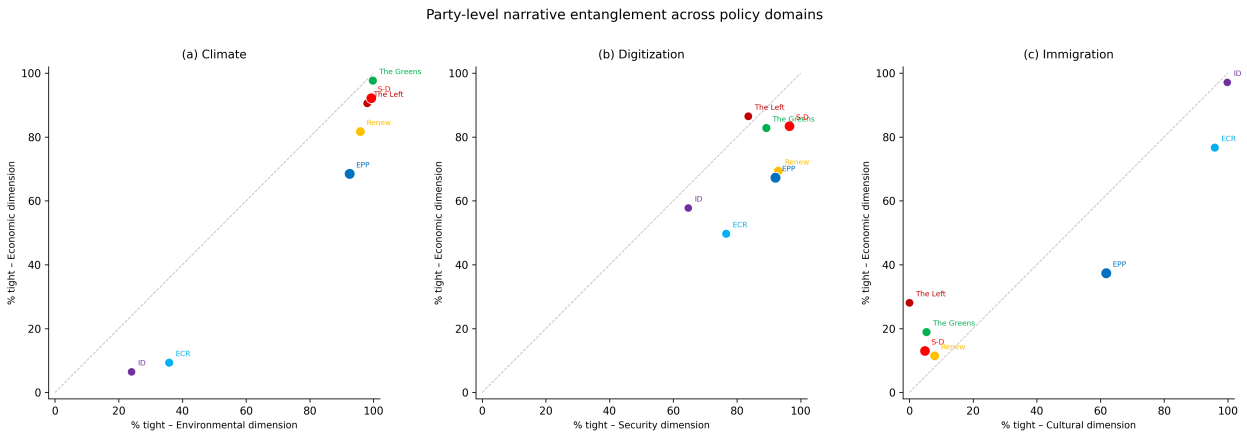

Note: Each point is a party group; point size is proportional to the square root of total speeches. For each domain, the x-axis is the share of the group's domain-relevant speeches with a "tighter" narrative on the non-economic dimension (environmental, security, or cultural), and the y-axis is the analogous share on the economic dimension.

Figure 4 – Party-level narrative entanglement across policy dimensions

gration at 96.5%, followed by climate at 90.8% and digitization at 82.7%. Figure C.4 in the Supplemental Appendix shows that entanglement varies strongly across parties and policies, with the highest proportion of entanglement among The Greens on climate policy and ID on immigration policy, and the largest non-entangled shares concentrated in digitization with ECR, ID, and Renew.

The prevalence of no-trade-off speeches is consistent with narrative entanglement. However, it does not by itself imply that a change in one narrative dimension induces a change in another. We turn to a formal test of this next.

4. Testing the theory: EU climate policy narratives

Our theory of narrative entanglement implies that a shock to one narrative dimension of a policy will induce politicians to also change their narratives on the other. More concretely, for the case of climate policy, it implies that a shock to perceived economic costs of climate policy will make politicians change their environmental narratives.

We now test these predictions by comparing speeches given before and after Russia's invasion of Ukraine on 24 February 2022. This quickly became a major topic in the European Parliament: Panel (a) of Figure 5 shows how often Ukraine was mentioned per speech per month, where Russia's invasion is marked by a red vertical line. The invasion underlined Europe's dependence on Russian fossil fuels and quickly prompted increased focus on gas prices. This is evident in Panel (b) of Figure 5, which shows the mentions per speech of the

word "gas".[24]

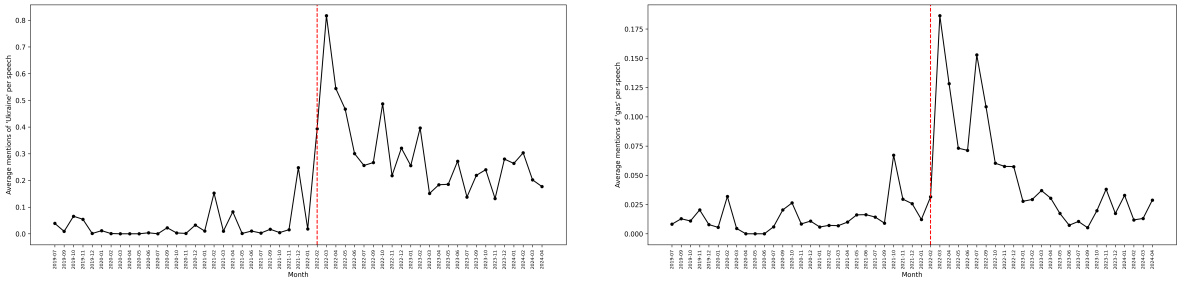

(a) Ukraine **(b)** Gas

Figure 5 – Mentions of "Ukraine" (panel a) and "gas" (panel b) per speech (monthly)

Russia's invasion sharply increased gas prices throughout the EU (see Supplemental Appendix Figure C.1), leading to higher short-run economic costs of climate policy: building up renewable energy capabilities becomes more expensive, since energy is a major input, and carbon levies are felt more strongly under higher energy prices. This increase in the median short-run costs to voters from climate policy maps into our model as an increase in $\bar{\alpha}$.[25]

While changing the economic costs of climate policy, Russia's invasion did not change the laws of physics that determine whether tighter climate policy can mitigate climate change. Without narrative entanglement, environmental narratives should therefore not change. This provides a natural testing ground for our theory of narrative entanglement.

Before moving to the empirical analysis, in Section 4.1 we apply our general framework to formulate predictions specific to climate policy narratives. In Section 4.2 we lay out our empirical strategy for analyzing shifts in climate narratives following Russia's invasion of Ukraine, and in Section 4.3 we present results. In Section 4.4 we then show how changes in narratives relate to changes in legislation based on roll-call data on climate policy.

4.1. Empirical model

To take our theory to the data, we need to make some adjustments to our general framework. First, the analysis so far has been at the party level. In our empirical application, however, the unit of observation is the individual politician, and the narratives we observe are those each politician *chooses* to voice in speeches. Since politicians do not speak on every

[24] There was already an increased frequency of the word "gas" in October 2021, when energy prices rose (see Supplemental Appendix Figure C.1) amid reports of a cut in Russian gas supplies to the EU in the last months of 2021, which was officially denied by the Kremlin. See, for instance, BBC News (2021).

[25] In the long run, the energy crisis may well decrease the relative economic costs of climate policy, as renewable energy becomes a more cost-efficient alternative. However, in the context of the model, the more fitting interpretation of $\bar{\alpha}$ is the observable cost that is felt by voters within the same election cycle. This is because the model captures the part of climate policy costs that influences politicians' electoral incentives, i.e., the short-term rather than the long-term costs.

topic, mapping the theory to the data requires modeling the speech decision. Second, our theory is deliberately broad to nest different policy topics. To generate predictions specific to climate policy, we need to impose more structure on the model.

From parties' optimal narratives to politicians' speeches. Assume that each party $J \in \{L, R\}$ has a continuum of politicians indexed by j, and each politician chooses whether to speak on the policy ($c_j = 1$) or on some other, unmodeled topic ($c_j = 0$). A politician who speaks on the policy voices their party's equilibrium narrative, as characterized in Proposition 3, and receives the party's equilibrium winning probability $\pi_J(\bar{\alpha})$, defined in (7), as a payoff. A politician who speaks on another topic receives an idiosyncratic payoff $\omega_j \sim H$, where H is a strictly increasing c.d.f. and ω_j is realized and privately observed before the speech is made. The draw ω_j captures personal preferences, expertise, the politician's role within the party, or any other factor bearing on willingness to speak on the policy.

A politician speaks on the policy if and only if the party payoff from doing so exceeds the outside payoff. The probability that a politician of party J speaks on the policy is therefore

$$\Pr\big(\pi_J(\bar{\alpha}) \geq \omega_j\big) \;=\; H\big(\pi_J(\bar{\alpha})\big). \tag{13}$$

Politicians are more likely to speak on the policy when doing so benefits their party electorally. Proposition B.1 in the Supplemental Appendix derives the implications of this assumption for the general model.

Predictions on climate policy. We now home in on Application 1 on climate policy (see equation (8)) to derive the predictions relevant to our empirical investigation. We focus on the case where $\mu = 0$, i.e., politicians face motivated reasoners with certainty. This assumption leads to the stark prediction that only entangled narratives will be voiced in equilibrium. Importantly, it allows us to formulate predictions without specifying a true narrative for climate policy, an object of controversy with no simple answer. However, the assumption is not as stark as it may seem, as Proposition 4 showed that politicians may use exclusively entangled narratives even if they face the possibility of Bayesian voters.

The next corollary gives the main predictions of the empirical model focused on climate policy under the assumption that voters are motivated reasoners. It concentrates on the empirically relevant case of polarized equilibria, where the Left and Right offer different policies and narratives. As in Section 2.3, we also concentrate on equilibria where, if indifferent, parties choose the narrative that makes their policy look best.

Corollary 1. *Assume that voters have utility function* (8)*, that they are motivated reasoners with certainty ($\mu = 0$), and that the climate tax-wedge condition* (11) *holds. Then, for the intermediate range of costs on which the pure-strategy polarized equilibrium exists, the Left implements climate policy and the Right opposes it. An increase in costs $\bar{\alpha}$ strictly decreases*

left politicians' probability of speaking on climate policy and voicing $\overline{\nu}$, and strictly increases right politicians' probability of speaking on climate policy and voicing $\underline{\nu}$.

That in a polarized equilibrium the Left implements climate policy and the Right opposes it follows from the positive interaction between taxes and climate policy: the Left's higher tax rate $t_L > t_R$ amplifies the perceived growth spillover on public goods under the "green-growth" narrative $\overline{\lambda}$. The last sentence of Corollary 1 allows us to formulate a clear prediction of how narratives change after Russia's invasion of Ukraine. As we argued before, Russia's invasion raised the short-run cost of climate policy for European voters, which corresponds to an increase in the median voter's private cost $\bar{\alpha}$. Corollary 1 then predicts that narrative shifts appear in both the economic and non-economic dimensions, and that left politicians retreat to silence on climate while right politicians speak more and voice $\underline{\nu}$.

In addition to changing costs $\bar{\alpha}$, Russia's invasion may also have shaped narrative bounds on economic effects, arguably shifting down the best-case effect $\overline{\lambda}$ and the worst-case effect $\underline{\lambda}$. A decrease in $\overline{\lambda}$ has the same directional effect as an increase in $\bar{\alpha}$: it lowers the Left's electoral advantage and shifts climate speech toward the Right, as long as the tax-wedge condition continues to hold. A decrease in $\underline{\lambda}$ does not directly affect this within-equilibrium speech margin, because the median voter adopts $\overline{\nu}$ in the polarized equilibrium; instead, it affects the surrounding cutoff regions by making implementation under the pessimistic economic narrative less attractive. Overall, a downward revision in economic narrative bounds would reinforce the predictions of Corollary 1.

4.2. Empirical strategy

We estimate the climate narrative shifts after Russia's invasion of Ukraine by using within-MEP variation. For this purpose, we use the following specification for each speech i made by MEP m at date t:

$$y_{imt} = \beta Post_t + \delta_m + X'_{imt}\gamma + \varepsilon_{imt}, \tag{14}$$

where y_{imt} is one of the three following indicator variables: (1) $y_{imt} = 1$ when speech i is on the topic of climate and 0 otherwise; (2) $y_{imt} = 1$ when speech i uses the negative economic narrative on climate change and 0 otherwise; or (3) $y_{imt} = 1$ when speech i uses the climate-skeptic narrative on the environmental dimension and 0 otherwise. On the right-hand side, $Post_t$ is equal to 1 after 24 February 2022 and 0 otherwise; δ_m are MEP fixed effects; and X_{imt} is a vector of control variables which, depending on the specification, can comprise interactions between $Post_t$ and the party of the MEP m at time t, as well as interactions between $Post_t$ and the country of the MEP.[26] We are primarily interested in the sign and

[26] While most of our observations are made by MEPs who in our sample only speak for one party, some MEPs do switch during the time period we study (see Supplemental Appendix Table C.1), for instance when

size of β and the interaction between $Post_t$ and the parties.

Our identifying assumption is that, once controlling for MEP fixed effects, the period after Russia's invasion of Ukraine constitutes a shock to the short-term economic costs of climate policy ($\bar{\alpha}$ in our model) that is exogenous to the true environmental effect of climate policies ($\rho^\star$). Conditioning on MEPs takes into account time-invariant heterogeneity across countries (since each MEP represents a single country) as well as the time-invariant personal characteristics of politicians. Moreover, we are looking at the *same* EU Parliament in the pre- and post-periods, with the same leadership under Ursula von der Leyen. In terms of heterogeneity, we allow the effect to be different across the eight party groups in the EU and additionally across the member states.

4.3. Shifts in climate policy narratives

We now turn to our main empirical results of how Russia's invasion of Ukraine shaped climate policy narratives, following the specification in (14).[27]

Probability of speaking on climate. Columns 1–3 of Table 2 show how the probability of making a speech on the topic of climate changed after Russia's invasion of Ukraine. The first column only includes $Post_t$ alongside MEP fixed effects and no interactive variables. Column 2 adds interactions between $Post_t$ and parties, while Column 3 additionally adds interactions with country. The omitted party interaction is with The Greens (for Columns 2 and 3), while the omitted country category is Spain (for Column 3). Column 1 shows that overall, there was an increase in the probability of voicing climate policy narratives. Results in Column 2 show that, after the invasion, speeches made by the Green party were 1.3 p.p. less likely to be on the topic of climate (not statistically significant), whereas the additional effect for far-right ECR and ID was +7.7 p.p. and +11.2 p.p., respectively. Compared to a mean dependent variable of 18%, these are large effects. Results in Column 3 show that these results are robust to partialing out differential effects by country.

Narrative shifts on the economic dimension. We now turn to narratives in the economic dimension. For this purpose we re-specify the dependent variable to take the value 1 if the speech implies that looser climate policy would benefit the economy or that it has no effect on the economy, and 0 otherwise. Note that the zero category contains both speeches that are not on the topic of climate and speeches that imply that a tighter climate policy

an entire country-level party switches political groups. Our results are not driven by such switchers, and remain similar when explicitly accounting for them (see the regression results of Table C.2).

[27] To allow correlation in the error terms within party groups, we cluster standard errors by country times EU party group throughout (clustering just at the level of the 8 party groups would give too few clusters). To accommodate party switchers (see Table C.1), we give each MEP who belonged to two or more party groups a separate cluster; for instance, MEPs who were in EPP and NI will be assigned to an EPP-NI cluster. This gives us 173 clusters in total.

Table 2 – Regression: Change in climate speeches and narratives after Russia's invasion of Ukraine

	Speech on climate			**Negative economic narrative**			**Climate skeptic narrative**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post	0.019***	-0.013	-0.014	0.028***	0.005***	0.003	0.008***	0.000	-0.002
	(0.006)	(0.019)	(0.017)	(0.004)	(0.001)	(0.005)	(0.002)	(0.001)	(0.004)
Post × ECR		0.077***	0.093***		0.065***	0.073***		0.019**	0.020**
		(0.022)	(0.019)		(0.009)	(0.010)		(0.009)	(0.009)
Post × EPP		0.036*	0.042**		0.023***	0.028***		0.010***	0.009**
		(0.021)	(0.018)		(0.006)	(0.006)		(0.002)	(0.004)
Post × ID		0.112***	0.122***		0.093***	0.092***		0.028***	0.024***
		(0.024)	(0.026)		(0.014)	(0.013)		(0.010)	(0.009)
Post × NI		0.063***	0.080***		0.040**	0.035**		0.009	0.003
		(0.022)	(0.023)		(0.017)	(0.017)		(0.006)	(0.006)
Post × Renew		0.002	0.008		0.011**	0.012*		0.006*	0.002
		(0.022)	(0.019)		(0.005)	(0.006)		(0.003)	(0.004)
Post × S-D		0.013	0.019		-0.001	0.004		0.000	-0.000
		(0.021)	(0.017)		(0.002)	(0.004)		(0.001)	(0.003)
Post × The Left		0.003	0.005		-0.003	0.002		-0.000	0.001
		(0.024)	(0.022)		(0.006)	(0.005)		(0.003)	(0.004)
Observations	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239
R^2	0.21	0.21	0.21	0.19	0.19	0.19	0.13	0.13	0.13
Mean Dep. Var.	0.18	0.18	0.18	0.050	0.050	0.050	0.019	0.019	0.019
MEP FE	X	X	X	X	X	X	X	X	X
Country X Post			X			X			X

Note: Standard errors clustered by country × party group. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Binary dependent variable for Columns 1–3: 1 = speech voices a narrative on either the environmental or the economic dimension of climate policy; 0 = it does not. Binary dependent variable for Columns 4–6: 1 = speech uses a "looser" or "no effect" climate policy narrative on the economy; 0 = it uses a "tighter" narrative or no narrative. Binary dependent variable for Columns 7–9: 1 = speech uses a "looser" or "no effect" climate policy narrative on climate change; 0 = it uses a "tighter" narrative or no narrative. Omitted categories: The Greens for the interactions with party (Columns 2, 3, 5, 6, 8, 9), and Spain for the interactions with country (Columns 3, 6, 9).

benefits the economy. It therefore looks jointly at the intensive and extensive margins of speech-making, covering the predictions from Corollary 1.[28]

Our theory suggests that following Russia's invasion of Ukraine, MEPs from the left should tend to fall silent instead of voicing supporting narratives, while those from the right should be more likely to speak on climate and use narratives opposing climate policy. In the specification with no interactions, we expect the coefficient on $Post_t$ to be positive as right-wing parties become more vocal and increasingly voice the narrative that climate policy is detrimental to the economy (positive effect) while left-wing parties fall more silent on the topic of climate (no effect, as the shift toward silence is captured within the 0-value of the dependent variable). When including interactions between $Post_t$ and parties, we expect a null effect for left-wing parties, and a positive effect for right-wing parties.

Columns 4–6 of Table 2 show the regression results, following the specification order of Columns 1–3 but for the economic narrative dependent variable. The results accord with the predictions of our theory. Speeches are 2.8 p.p. more likely to imply that climate policy has a looser or no-effect relationship with the economy (Column 4). Compared to a mean dependent variable of 5.0%, this is a large effect. When looking at cross-party heterogeneity (Column 5), we see that, as expected, the effects are concentrated among the far-right ECR and ID parties, and, to a lesser extent, the center-right EPP party. The coefficient on The Greens is also statistically significant but very small, in line with a precisely estimated null. The statistical significance for The Greens does not persist when additionally partialing out the interaction of $Post_t$ with the country that the MEP represents (Column 6), while the effects for other parties do not change significantly.

Narrative shifts in the environmental dimension. We now turn to results for narratives on the environmental dimension. This relates to our main theoretical insight on narrative entanglement. For this purpose, we replace our dependent variable with a binary variable which takes the value of 1 when the speech implies that looser climate policy would mitigate climate change or that it has no effect on climate change, and 0 otherwise. Again, the 0 category contains both speeches that imply that a tighter climate policy mitigates climate change and speeches that are on another topic.

If narratives are entangled across different policy dimensions, we should see similar patterns in narratives on climate as in Columns 4–6 of Table 2. If they are not entangled, then results on the environmental dimension should be unaffected.

Columns 7–9 of Table 2 show that regression results are in line with narrative entanglement. Speeches of the same MEP are 0.8 p.p. more likely to voice climate-skeptic narratives

[28] Appendix Table C.4 shows similar results for the intensive margin, when conditioning on speeches made on the topic of climate policy.

(Column 7). This is a large effect when compared to a mean dependent variable of 1.9%. Columns 8 and 9 again show that there is significant heterogeneity, with the far-right ECR and ID parties driving the results, as expected. Additionally, the EPP party shows a smaller but statistically significant response in both Columns 8 and 9.

Robustness and alternative mechanisms. Supplemental Appendix C.2 reports robustness checks and tests of alternative mechanisms. Figure C.5 shows that descriptive results that do not account for MEP fixed effects lead to the same patterns. Figure C.6 gives results with an event-study replacing $Post_t$ with year dummies centered on the invasion. Table C.2 shows that results remain very similar when separating MEPs who switch party from those who do not. Table C.3 shows that replacing MEP with national-party fixed effects leads to similar results. Table C.4 investigates the intensive margin and shows that, conditional on speaking, right-wing parties use more negative narratives on climate policy. Finally, Figure C.7 rejects a "mechanical" crowding-out effect of climate policy narratives which would occur if parties that speak more about Ukraine do so by decreasing the number of speeches on climate policy.

4.4. Narrative changes and legislative voting

Next, we examine to what extent the previously documented changes in climate narratives relate to changes in legislative behavior. For this purpose, we download roll-call data for the 9th European Parliament, identify all roll calls on climate policy, and link them to our speech data.[29]

Based on this, we investigate how changes in narratives following Russia's invasion of Ukraine correlate with changes in voting on climate policy in line with the modal vote of The Greens, where we take the latter as a proxy for regulation that is in favor of tighter climate policy. We do this at the national party rather than MEP level since most variation in voting is at the national party level (see Supplemental Appendix D).[30] Figure 6 plots changes in narratives versus changes in voting against the modal vote of The Greens on climate policy legislation, at the national-party level.[31] The size of the dot reflects the number of speeches made by the party, and the dashed line is a speech-weighted regression

[29] Parltrack provides roll-call data at https://parltrack.org/dumps; the Legislative Observatory gives OEIL classifications, including on climate policy, available at https://oeil.secure.europarl.europa.eu/; and HowTheyVote data allows us to link these two datasets, at https://github.com/HowTheyVote/data. Supplemental Appendix D discusses the roll-call data in more detail.

[30] Supplemental Appendix Figure D.5 also shows a positive relationship at the MEP level.

[31] For each national party, the x-axis is the change in the share of all speeches that use a non-tightening narrative on the environmental dimension (Panel a) or economic dimension (Panel b), and the y-axis is the change in the share of its climate-topic roll-call votes cast opposite the Greens' modal choice; roll calls on which the Greens' modal choice is abstain or missing are dropped.

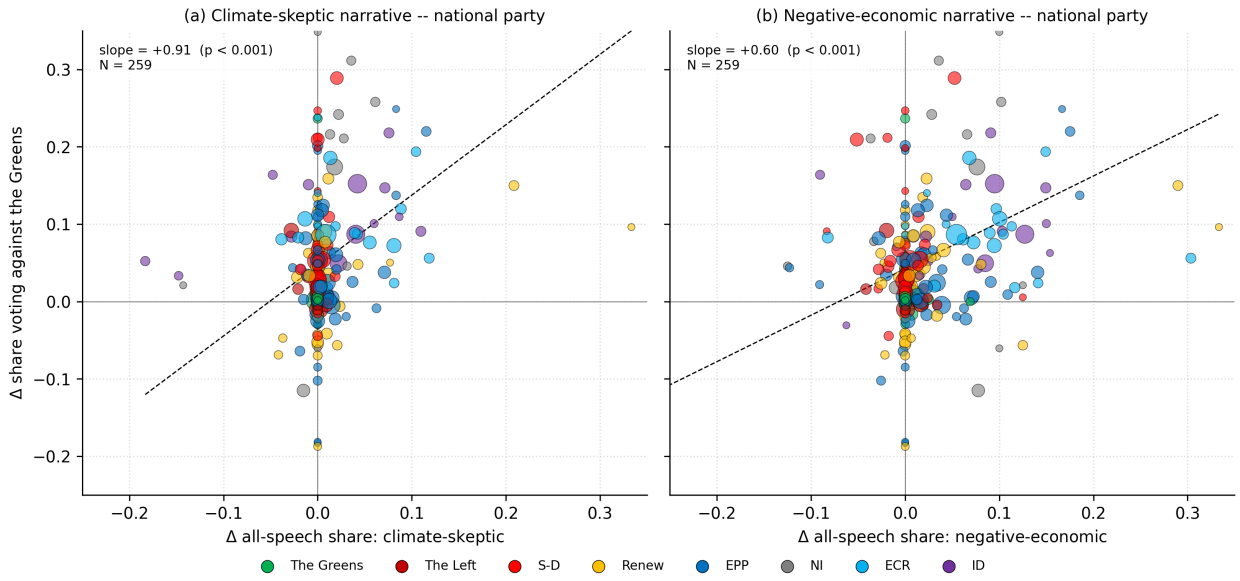

Note: Each point is a national party; changes are before versus after Russia's invasion of Ukraine. Vertical axis: change in share of the party's votes cast opposite the Greens' modal choice on climate-topic roll calls (benchmark-abstain or missing-mode roll calls dropped). Panel (a) horizontal axis: change in share of all speeches with a non-tightening climate-change narrative. Panel (b) horizontal axis: change in share of all speeches with a non-tightening economic narrative. Point size is proportional to the square root of total speeches, and color denotes the EU party group (mode for national parties with more than one party group). The dashed line is a speech-weighted WLS fit. The p-values are based on standard errors clustered by country; clustering by country times party group would give singleton clusters.

Figure 6 – Shifts in narratives versus legislative voting among national parties

line. Both panels show positive relationships between narrative change and legislative vote change, with the fitted relationship steeper for the environmental dimension.

5. Conclusion

We have developed a theory that shows why the same incentives that lead politicians to hide trade-offs also lead to entangled narratives, where changes in economic costs spill over to narratives about non-economic policy effects, and vice versa. Incentives for narrative entanglement stem from the presence of motivated voters, but can remain even if Bayesian voters are more likely. We show evidence for entangled narratives in the 9th European Parliament on the topics of climate, immigration, and digitization policy. We then test this formally by analyzing shifts in climate policy narratives before versus after Russia's invasion of Ukraine. In line with narrative entanglement, economic *and* environmental narratives on climate policy turned more negative after Russia's invasion. As our theory predicts, the results are different for right-wing versus left-wing politicians. Left-wing MEPs speak less often on the topic of climate, while far-right MEPs speak more often on it, increasingly

questioning whether climate policy is needed to address climate change.

Narrative entanglement matters because it affects the climate policy chosen in equilibrium. When economic costs increase, political messages turn skeptical, questioning the necessity of climate action. This feeds back into the beliefs of voters who take politicians' narratives into account. Hence, following increases in costs, the political discourse is likely to amplify any doubts in the electorate, and policy change is bound to follow.

In the roll-call data on legislative votes, we find a relationship between increases in climate-skeptic narratives and increases in legislative opposition to climate policy. The policy change can also be seen in partial reversals of the European Green Deal. For instance, the EU classified natural gas investments as "green" under the EU taxonomy in July 2022 (European Commission, 2022)[32] and in November 2023 dropped plans to cut fertilizer use by 50% (European Parliament, 2023). More green policy reversals followed after von der Leyen's re-election in 2024, including on supply chain assessments and the planned ban on cars with combustion engines (Politico, 2025b).

For supporters of climate policy, these results reveal the sensitivity of environmental narratives to economic disruptions. Our framework points to the importance of crafting appealing narratives to underpin policy change, especially in situations where the range of acceptable opinions is large. To be effective, policy narratives have to be able to compete with no-trade-off narratives that arise under political competition. Those designing climate policies must be mindful of narrative shifts in order to ensure long-term support for climate action, especially during economic hardships.

Narrative entanglement is not unique to climate policy. We show that narratives are particularly entangled in immigration policy, where an established literature stresses the importance of narratives (Alrabah et al., 2025). Moreover, we focus on the case of the EU because our main climate policy application can be investigated best in that context, but nothing about the mechanism of narrative entanglement makes it specific to this setting.

Our analysis focused on narrative entanglement within a policy domain. However, future extensions could examine how narrative entanglement ties into policy entanglement. Consider two policies that share an economic dimension, but have distinct non-economic dimensions, e.g., environmental narratives for climate policy and cultural narratives for immigration policy. Suppose that an economic shock makes climate policy and immigration policy more costly on the economic dimension. In the climate dimension, we investigated how this results in changes in environmental narratives. The same logic could lead cultural narratives of immigration policy to change. The result could be a spurious comovement in cultural

[32] In September 2025, the European General Court dismissed a complaint brought by Austria that this change in taxonomy breached EU law (Politico, 2025a).

narratives of immigration policy and environmental narratives of climate policy—two dimensions that are factually unrelated. In our data we see such clustering cross-sectionally across parties: as Section 3.3 shows, parties cluster into (1) entangled pro-climate narratives and entangled pro-immigration narratives, or (2) entangled anti-climate narratives and entangled anti-immigration narratives. Future work could investigate to what extent clustering across policy topics is in fact driven by narrative entanglement.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. The environment and directed technical change. *American Economic Review* 102(1): 131–166.

Acemoglu, Daron, and Will Rafey. 2023. Mirage on the horizon: Geoengineering and carbon taxation without commitment. *Journal of Public Economics* 219: 104802.

Aina, Chiara. 2023. Tailored stories. Working Paper, URL: https://chiaraaina.github.io/files/TailoredStories_Aina.pdf.

Alrabah, Ala, Andreas Beerli, Dominik Hangartner, and Dalston Ward. 2025. The Free Movement of People and the Success of Far-Right Parties: Evidence from Switzerland's Border Liberalization. *American Political Science Review* 119(3): 1426–1445.

Arold, Benjamin W, Elliott Ash, W Bentley MacLeod, and Suresh Naidu. 2024. Do Words Matter? Working Paper, URL: https://www.research-collection.ethz.ch/entities/publication/84113e60-73e3-4a70-ad1c-5fdd100b46b4.

Ash, Elliott, and Stephen Hansen. 2023. Text algorithms in economics. *Annual Review of Economics* 15(1): 659–688.

Battaglini, Marco, and Bård Harstad. 2016. Participation and duration of environmental agreements. *Journal of Political Economy* 124(1): 160–204.

Battaglini, Marco, and Bård Harstad. 2020. The political economy of weak treaties. *Journal of Political Economy* 128(2): 544–590.

BBC News. 2021. Russia denies weaponising energy amid Europe gas crisis. Published on October 13. URL: https://www.bbc.co.uk/news/world-europe-58896847.

Bénabou, Roland. 2008. Ideology. *Journal of the European Economic Association* 6(2-3): 321–352.

Bénabou, Roland. 2015. The economics of motivated beliefs. *Revue d'économie politique* (5): 665–685.

Benoit, Kenneth, Scott De Marchi, Conor Laver, Michael Laver, and Jinshuai Ma. 2026. Using Large Language Models to Analyze Political Texts through Natural Language Understanding. *American Journal of Political Science* 70(2).

Besley, Timothy, and Adam Brzezinski. 2025. A Dynamic Model of Political Narratives and the Rise of Neoliberalism. CEPR working paper DP20410, URL: https://cepr.org/publications/dp20410.

Bombardini, Matilde, Frederico Finan, Nicolas Longuet-Marx, Suresh Naidu, and Francesco Trebbi. 2025. Climate Politics in the United States. NBER working paper 34120, URL: https://www.nber.org/papers/w34120.

Brannlund, Runar, and Lars Persson. 2012. To tax, or not to tax: preferences for climate policy attributes. *Climate Policy* 12(6): 704–721.

Brzezinski, Adam. 2023. Narratives of Blame. *Oxford DPhil Chapter* 6–90.

Bursztyn, Leonardo, Georgy Egorov, Ingar Haaland, Aakaash Rao, and Christopher Roth. 2023a. Justifying dissent. *The Quarterly Journal of Economics* 138(3): 1403–1451.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. 2023b. Opinions as facts. *The Review of Economic Studies* 90(4): 1832–1864.

Caplan, Bryan. 2007. *The Myth of the Rational Voter: Why Democracies Choose Bad Policies - New Edition*. Princeton University Press, rev - revised edition.

Chinoy, Sahil, Nathan Nunn, Sandra Sequeira, and Stefanie Stantcheva. 2024. Zero-sum thinking and the roots of US political divides. CEPR working paper DP19297, URL: https://cepr.org/publications/dp19297.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2025. Fighting climate change: International attitudes toward climate policies. *American Economic Review* 115(4): 1258–1300.

Djourelova, Milena, Ruben Durante, Elliot Motte, and Eleonora Patacchini. 2024. Experience, narratives, and climate change beliefs. CEPR working paper DP18738, URL: https://cepr.org/publications/dp18738.

Drews, Stefan, and Jeroen CJM Van den Bergh. 2016. What explains public support for climate policies? A review of empirical and experimental studies. *Climate policy* 16(7): 855–876.

Eliaz, Kfir, and Ran Spiegler. 2020. A Model of Competing Narratives. *American Economic Review* 110(12): 3786–3816.

Epley, Nicholas, and Thomas Gilovich. 2016. The mechanics of motivated reasoning. *Journal of Economic Perspectives* 30(3): 133–140.

European Commission. 2022. EU Taxonomy: Complementary Climate Delegated Act to accelerate decarbonisation. URL: https://finance.ec.europa.eu/publications/eu-taxonomy-complementary-climate-delegated-act-accelerate-decarbonisation_en.

European Parliament. 2023. No Majority in Parliament for Legislation to Curb Use of Pesticides (Press release). URL: https://www.europarl.europa.eu/news/en/press-room/20231117IPR12215/no-majority-in-parliament-for-legislation-to-curb-use-of-pesticides.

Gentzkow, Matthew, and Jesse M Shapiro. 2010. What drives media slant? Evidence from US daily newspapers. *American Economic Review* 100(4): 1173–1196.

Harstad, Bård. 2012. Climate contracts: A game of emissions, investments, negotiations, and renegotiations. *Review of Economic Studies* 79(4): 1527–1557.

Harstad, Bård. 2016. The dynamics of climate agreements. *Journal of the European Economic Association* 14(3): 719–752.

Helm, Dieter, Cameron Hepburn, and Richard Mash. 2003. Credible carbon policy. *Oxford Review of Economic Policy* 19(3): 438–450.

Intrator, Yotam, Matan Halfon, Roman Goldenberg, Reut Tsarfaty, Matan Eyal, Ehud Rivlin, Yossi Matias, and Natalia Aizenberg. 2024. Breaking the Language Barrier: Can Direct Inference Outperform Pre-Translation in Multilingual LLM Applications? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, edited by Duh, Kevin, Helena Gomez, and Steven Bethard, 829–844. Mexico City, Mexico: Association for Computational Linguistics.

Izzo, Federica, Gregory J Martin, and Steven Callander. 2023. Ideological Competition. *American Journal of Political Science* 67(3): 687–700.

Kunda, Ziva. 1990. The case for motivated reasoning. *Psychological Bulletin* 108(3): 480.

Levy, Gilat, Ronny Razin, and Alwyn Young. 2022. Misspecified Politics and the Recurrence of Populism. *American Economic Review* 112(3): 928–62.

Little, Andrew T. 2019. The Distortion of Related Beliefs. *American Journal of Political Science* 63(3): 675–689.

Little, Andrew T., Keith E. Schnakenberg, and Ian R. Turner. 2022. Motivated Reasoning and Democratic Accountability. *American Political Science Review* 116(2): 751–767.

Politico. 2025a. Top European court rules nuclear power can be green. Published on September 10, URL: https://www.politico.eu/article/top-european-court-rules-nuclear-power-green/.

Politico. 2025b. Ursula von der Leyen: Green Deal — EU Politics, Economy, Policy. Published on June 25, URL: https://www.politico.eu/article/ursula-von-der-leyen-green-deal-eu-politics-economy-policy/.

Quinn, Kevin M, Burt L Monroe, Michael P Colaresi, Michael H Crespin, and Dragomir R Radev. 2010. How to analyze political attention with minimal assumptions and costs. *American Journal of Political Science* 54(1): 209–228.

Rauh, Christian, and Jan Schwalbach. 2020. The ParlSpeech V2 data set: Full-text corpora of 6.3 million parliamentary speeches in the key legislative chambers of nine representative democracies. URL: https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/L4OAKN.

Schwartzstein, Joshua, and Adi Sunderam. 2021. Using models to persuade. *American Economic Review* 111(1): 276–323.

Stantcheva, Stefanie. 2024. Perceptions, Mindsets and Beliefs Shaping Policy Views. 2024 Economica-Coase Lecture. URL: https://socialeconomicslab.org/wp-content/uploads/2025/01/Stantcheva_Economica_Coase_Lecture.pdf.

Ulph, Alistair, and David Ulph. 2013. Optimal climate change policies when governments cannot commit. *Environmental and Resource Economics* 56(2): 161–176.

Whitmarsh, Lorraine, and Stuart Capstick. 2018. Perceptions of climate change. In *Psychology and climate change*, 13–33. Elsevier.

A. Proofs

A.1. Statement and proof of Lemma A.1

Lemma A.1. *The winning probability of a party equals the probability that the median voter votes for the party.*

Proof. We start with the proof for the case of motivated reasoners. The proof for Bayesian reasoners is then a simpler version of the same.

Fix (x_L, ν_L) and (x_R, ν_R) and write ν_k for $k \in \{L, R\}$. Write $u(\alpha, x_J, t_J; \nu_k) = u_J(\alpha; \nu_k)$. By $\partial^2_{\lambda\lambda} y(\alpha, x, t; \lambda) = 0$ and $\partial^2_{\alpha\lambda} y(\alpha, x, t; \lambda) = 0$ we can write $y(\alpha, x, t; \lambda) = \tilde{y}(\alpha, x, t) + B(x, t)\lambda$.

Therefore, for $J \in \{L, R\}$ and $k \in \{L, R\}$, utility can be written as

$$u_J(\alpha; \nu_k) \equiv \tilde{y}_J(\alpha) + c_J(\nu_k),$$

where the first term depends only on policies by party J and individual cost α, while the second term depends only on policies by J and narratives ν_k.

For a given J, the difference $u_J(\alpha; \nu_L) - u_J(\alpha; \nu_R) = c_J(\nu_L) - c_J(\nu_R)$ is independent of α, so the narrative that maximizes $u_J(\alpha; \nu)$ over $\nu \in \{\nu_L, \nu_R\}$ is the same for all α. Let $\nu^L \in \arg\max_{\nu \in \{\nu_L, \nu_R\}} u_L(\alpha; \nu)$ denote the optimal narrative for $J = L$ out of ν_L, ν_R and let $\nu^R \in \arg\max_{\nu \in \{\nu_L, \nu_R\}} u_R(\alpha; \nu)$ denote the optimal narrative for $J = R$.

By motivated reasoning, $\nu_i = \nu^L$ if $u_L(\alpha_i; \nu^L) > u_R(\alpha_i; \nu^R)$, $\nu_i = \nu^R$ if $u_L(\alpha_i; \nu^L) < u_R(\alpha_i; \nu^R)$, and ν_i randomly chosen at the knife-edge case of equal utility.

Moreover, $u_L(\alpha; \nu^L) - u_R(\alpha; \nu^R) = \tilde{y}_L(\alpha) - \tilde{y}_R(\alpha) + c_L(\nu^L) - c_R(\nu^R)$, where $\tilde{y}_L(\alpha) - \tilde{y}_R(\alpha)$ is weakly monotone in α: it does not depend on α when $x_L = x_R = 0$ (since by assumption y does not depend on α when $x = 0$), it is strictly decreasing when $x_L = 1, x_R = 0$ (since by assumption $\partial_\alpha y(\alpha, 1, \cdot) < 0$), it is strictly increasing when $x_L = 0, x_R = 1$, and it is strictly increasing when $x_L = x_R = 1$ (by assumption on tax preference in that case). Hence, ν_i is characterized by a cutoff in α.

Fix any ν_k. The vote margin for a type α is

$$D_k(\alpha) \equiv u_L(\alpha;\nu_k) - u_R(\alpha;\nu_k) = \tilde{y}_L(\alpha) - \tilde{y}_R(\alpha) + C(\nu_k), \qquad C(\nu_k) \equiv c_L(\nu_k) - c_R(\nu_k),$$

so $D_k(\alpha)$ is monotone in α. Furthermore, since $c_L(\nu^L) \geq c_L(\nu^R)$ and $c_R(\nu^R) \geq c_R(\nu^L)$, we have $C(\nu^L) \geq C(\nu^R)$. Since $\nu_i = \nu^L$ if and only if $u_L(\alpha_i;\nu^L) \geq u_R(\alpha_i;\nu^R)$, narrative switching occurs at a cutoff in α. If $\tilde{y}_L(\alpha) - \tilde{y}_R(\alpha)$ is (weakly) increasing in α, then as α rises voters (weakly) switch from ν^R to ν^L, so the constant term in the vote margin jumps (weakly) upward from $C(\nu^R)$ to $C(\nu^L)$. If instead $\tilde{y}_L(\alpha) - \tilde{y}_R(\alpha)$ is (weakly) decreasing in α, then as α rises voters (weakly) switch from ν^L to ν^R, so the constant term jumps (weakly) downward from $C(\nu^L)$ to $C(\nu^R)$. In either case, the jump is in the same direction as the monotonicity of the baseline term $\tilde{y}_L(\alpha) - \tilde{y}_R(\alpha)$, and therefore $u_L(\alpha;\nu_i) - u_R(\alpha;\nu_i)$ remains monotone in α even as ν_i may change endogenously.

Therefore, when faced with motivated voters, L wins if and only if the median voter $\bar{\alpha}$ votes for L, and

$$\Pr(L \text{ wins}) = \Pr(u(\bar{\alpha}, x_L, t_L;\nu_m) - \delta \geq u(\bar{\alpha}, x_R, t_R;\nu_m)),$$

where ν_m is the narrative chosen by the median voter.

For Bayesian voters, the proof follows the same steps, except that Bayesian voters do not change their adopted narratives based on their α_i, and hence do not have any discontinuous jumps.

A.2. Proof of Proposition 1

By contradiction, assume that in equilibrium no party supplies $\nu^\star$. Then, all voters believe in ν^{PM}. By assumption, $\nu^{\text{PM}} \neq \nu^\star$, and parties are not indifferent between voters believing in ν^{PM} versus $\nu^\star$ as the median voter's cross-party utility difference between parties is, by assumption, strictly different under ν^{PM} versus $\nu^\star$ for any combination of policies. But by the zero-sum nature of the game, it must be that exactly one party would benefit from supplying $\nu^\star$, in which case all voters adopt that narrative. Hence, in any equilibrium, $\nu^\star$ gets supplied and all voters adopt that narrative. By Lemma A.1 parties maximize their probability of winning by maximizing the cross-party utility difference of the median Bayesian voter, which for a given opponent's platform and a given narrative is equivalent to maximizing the voter's utility under their narrative. Since all voters adopt $\nu^\star$, it follows that parties offer the best policy as evaluated from the true narrative $\nu^\star$.

A.3. Proof of Proposition 2

Using the definitions from the proof to Lemma A.1, we can write the median voter's utility difference between Left and Right as

$$D(\bar{\alpha}) = u_L(\bar{\alpha}; \nu_m) - u_R(\bar{\alpha}; \nu_m) = \tilde{y}_L(\bar{\alpha}) - \tilde{y}_R(\bar{\alpha}) + C(\nu_m). \tag{A.1}$$

The Left maximizes its winning chances by maximizing $D(\bar{\alpha})$, the Right does so by minimizing it.

We can now show that for any (x_J, x_{-J}), the optimal strategy is $\nu_J \in \{\underline{\nu}, \overline{\nu}\}$. We consider the four possible cases separately.

Suppose $x_L = x_R = 0$. The median voter's utility is weakly decreasing in λ_m and ρ_m, so $\underline{\nu}$ maximizes their utility and they will adopt it if it is offered. The utility difference is linear in λ_m, so it is optimal for either Left (if (A.1) is decreasing in λ_m) or Right (if (A.1) is increasing in λ_m) to offer $\underline{\nu}$, and the median voter will adopt it. Conditional on that, in equilibrium the other party is indifferent between any narrative.

Suppose $x_L = x_R = 1$. The median voter's utility is then weakly increasing in λ_m, ρ_m, and hence they will adopt $\overline{\nu}$ when it is offered. The utility difference is linear in λ, so it is optimal for either Left or Right to offer $\overline{\nu}$ and the median voter will adopt it. Conditional on that, in equilibrium the other party is indifferent between any narrative.

Suppose $x_L = 0$ and $x_R = 1$. The median voter utility difference is then strictly decreasing in λ_m, ρ_m. Across all narratives, the voter's utility is maximized either at $u_L(\bar{\alpha}; \underline{\nu})$ or at $u_R(\bar{\alpha}; \overline{\nu})$. If the former is larger, it is strictly optimal for Left to offer $\underline{\nu}$ and $\nu_m = \underline{\nu}$; if the latter is larger, then it is strictly optimal for Right to offer $\overline{\nu}$ and $\nu_m = \overline{\nu}$. In either case, the other party is indifferent between all narratives.

Suppose $x_L = 1$ and $x_R = 0$. By the same arguments as in the above paragraph, either $\nu_L = \overline{\nu}$ or $\nu_R = \underline{\nu}$ is strictly optimal, and the other party is indifferent between all narratives.

In all cases, party J has a best response in narratives belonging to $\{\underline{\nu}, \overline{\nu}\}$.

A.4. Proof of Proposition 3

By assumption, we restrict attention to strategies where implementation is paired with $\overline{\nu}$ and non-implementation with $\underline{\nu}$, which by the proof of Proposition 2 is always optimal. The policy stage is then a two-by-two zero-sum game.

For $J \in \{L, R\}$, write

$$U_J^x(\bar{\alpha}; \nu) \equiv u(\bar{\alpha}, x, t_J; \nu).$$

We now define a series of cutoffs in α. By Assumption 1, $U_J^1(\alpha; \nu)$ is continuous, strictly decreasing, and has full range in α, so each cutoff exists and has a unique solution.

Let $\hat{\alpha}^J$ be implicitly defined by

$$U_J^1(\hat{\alpha}^J; \overline{\nu}) = U_{\neg J}^0(\underline{\nu}).$$

Thus, in a separating profile where party J implements and party $\neg J$ does not, the median voter adopts $\overline{\nu}$ if and only if $\bar{\alpha} \leq \hat{\alpha}^J$.

Next define the within-narrative cutoffs $\check{\alpha}_\nu^J$ by

$$U_J^1(\check{\alpha}_\nu^J; \nu) = U_J^0(\nu), \qquad \nu \in \{\underline{\nu}, \overline{\nu}\}.$$

Finally, define the cross-narrative deviation cutoffs $\dot{\alpha}^J$ and $\ddot{\alpha}^J$ by

$$U_J^1(\dot{\alpha}^J; \overline{\nu}) - U_{\neg J}^1(\dot{\alpha}^J; \overline{\nu}) = U_J^0(\underline{\nu}) - U_{\neg J}^1(\dot{\alpha}^J; \underline{\nu}),$$

and

$$U_J^0(\underline{\nu}) - U_{\neg J}^0(\underline{\nu}) = U_J^1(\ddot{\alpha}^J; \overline{\nu}) - U_{\neg J}^0(\overline{\nu}).$$

The first is the cutoff at which party J is indifferent between staying in the $(1,1)$ pool and deviating to non-implementation when the deviation switches the median voter to $\underline{\nu}$. The second is the cutoff at which party J is indifferent between staying in the $(0,0)$ pool and deviating to implementation when the deviation switches the median voter to $\overline{\nu}$.

The primitive cutoff orderings established in Lemma A.2 give

$$\check{\alpha}_{\underline{\nu}}^J < \check{\alpha}_{\overline{\nu}}^J, \qquad \dot{\alpha}^J < \ddot{\alpha}^J, \qquad \dot{\alpha}^J < \check{\alpha}_{\overline{\nu}}^J, \qquad \check{\alpha}_{\underline{\nu}}^J < \ddot{\alpha}^J.$$

We now characterize each profile $(x_J, x_{\neg J})$. Consider first the pooling profile $(1,1)$. Party J's deviation is from implementation to non-implementation. If $\bar{\alpha} < \hat{\alpha}^{\neg J}$, the deviation does not switch the median voter away from $\overline{\nu}$, so party J does not deviate if and only if $\bar{\alpha} \leq \check{\alpha}_{\overline{\nu}}^J$. If $\bar{\alpha} > \hat{\alpha}^{\neg J}$, the deviation switches the median voter to $\underline{\nu}$, so party J does not deviate if and only if $\bar{\alpha} \leq \dot{\alpha}^J$. These two pieces stitch together into a single threshold condition $\bar{\alpha} \leq A_J$, where

$$A_J \equiv \min\left\{\check{\alpha}_{\overline{\nu}}^J, \max\{\hat{\alpha}^{\neg J}, \dot{\alpha}^J\}\right\}.$$

Thus $(1,1)$ is a pure-strategy equilibrium if and only if $\bar{\alpha} \leq A_L$ and $\bar{\alpha} \leq A_R$. Define

$$\tilde{\alpha}_1 \equiv \min\{A_L, A_R\}.$$

Then $(1,1)$ is a PSE iff $\bar{\alpha} \leq \tilde{\alpha}_1$.

Now consider the pooling profile $(0,0)$. Party J's deviation is from non-implementation to implementation. If $\bar{\alpha} < \hat{\alpha}^J$, the deviation induces the median voter to adopt $\overline{\nu}$, so party J does not deviate if and only if $\bar{\alpha} \geq \ddot{\alpha}^J$. If $\bar{\alpha} > \hat{\alpha}^J$, the deviation is evaluated under $\underline{\nu}$, so party J does not deviate if and only if $\bar{\alpha} \geq \check{\alpha}_{\underline{\nu}}^J$. These two pieces stitch together into $\bar{\alpha} \geq B_J$, where

$$B_J \equiv \max\left\{\check{\alpha}_{\underline{\nu}}^J, \min\{\hat{\alpha}^J, \ddot{\alpha}^J\}\right\}.$$

Hence $(0,0)$ is a pure-strategy equilibrium iff $\bar{\alpha} \geq B_L$ and $\bar{\alpha} \geq B_R$.

Define

$$\tilde{\alpha}_2 \equiv \max\{B_L, B_R\}.$$

Then $(0,0)$ is a PSE iff $\bar{\alpha} \geq \tilde{\alpha}_2$.

We next show that the two thresholds are ordered. Suppose, to the contrary, that $\tilde{\alpha}_1 > \tilde{\alpha}_2$. Then there exists a value of $\bar{\alpha} \in (\tilde{\alpha}_2, \tilde{\alpha}_1)$ that is not equal to any of the primitive cutoffs just defined. At this value, the characterizations above imply that both pooling profiles, $(1,1)$ and $(0,0)$, are pure-strategy equilibria. Let $D_{x_L x_R}(\bar{\alpha})$ denote Left's payoff in the induced two-by-two policy game, with the median voter's adopted narrative evaluated cell by cell. Since Left maximizes D and Right minimizes it, $(1,1)$ being a pure-strategy equilibrium implies $D_{01}(\bar{\alpha}) \leq D_{11}(\bar{\alpha}) \leq D_{10}(\bar{\alpha})$, while $(0,0)$ being a pure-strategy equilibrium implies $D_{10}(\bar{\alpha}) \leq D_{00}(\bar{\alpha}) \leq D_{01}(\bar{\alpha})$. Combining these two chains gives

$$D_{11}(\bar{\alpha}) \leq D_{10}(\bar{\alpha}) \leq D_{00}(\bar{\alpha}) \leq D_{01}(\bar{\alpha}) \leq D_{11}(\bar{\alpha}),$$

so all four payoff entries must be equal. But away from the primitive cutoff values, all active no-deviation inequalities defining the two pooling profiles are strict. Hence the four equalities cannot hold at the chosen $\bar{\alpha}$, a contradiction. Therefore it must be that $\tilde{\alpha}_1 \leq \tilde{\alpha}_2$, and the pooling regions cannot overlap on any non-degenerate interval. Thus, $(1,1)$ is the unique equilibrium for $\bar{\alpha} < \tilde{\alpha}_1$, and $(0,0)$ is the unique equilibrium for $\bar{\alpha} > \tilde{\alpha}_2$.

It remains to describe the intermediate region. If $\bar{\alpha} \in (\tilde{\alpha}_1, \tilde{\alpha}_2)$, then $\bar{\alpha} > \tilde{\alpha}_1$, so at least one party has a profitable deviation from $(1,1)$. Therefore $(1,1)$ is not an equilibrium. Similarly, $\bar{\alpha} < \tilde{\alpha}_2$, so at least one party has a profitable deviation from $(0,0)$. Therefore $(0,0)$ is not an equilibrium.

The only possible pure-strategy equilibria in this region are the two divergent profiles. From the same deviation conditions just derived, $(1,0)$ is a PSE iff $A_R \leq \bar{\alpha} \leq B_L$, because Right must not want to switch from non-implementation to implementation, and Left must not want to switch from implementation to non-implementation. Similarly, $(0,1)$ is a PSE iff $A_L \leq \bar{\alpha} \leq B_R$.

The two divergent pure-strategy equilibria cannot both occur in the intermediate region for the same cost $\bar{\alpha}$, except at knife-edge equality points: by the same zero-sum payoff-chain logic used above to rule out overlapping pooling profiles, simultaneous divergent equilibria would force all four payoff entries to equal each other. Thus at most one divergent pure-strategy equilibrium exists at (non-knife-edge) intermediate costs. If neither divergent profile is a pure-strategy equilibrium, then given that this is a finite game a mixed-strategy equilibrium exists which is unique except for knife-edge cases.

Therefore, for intermediate costs, there is no pooling equilibrium, and the unique non-

degenerate equilibrium is either a pure-strategy equilibrium with policy and narrative divergence or a mixed-strategy equilibrium.

A.5. Proof of Proposition 4

Write $\nu^e \equiv \hat{\nu}^{M,0}$ and $\Delta u(\cdot;\nu) \equiv \Delta u(\bar{\alpha}, x_{J^M}, x_{\neg J^M}, t_{J^M}, t_{\neg J^M}; \nu)$ for compactness. By the proof of Proposition 2, ν^e maximizes $\Delta u(\cdot;\nu)$ over $\mathcal{N}$, and hence $\Delta u(\cdot;\nu^e) > \Delta u(\cdot;\nu)$ for $\nu = \nu^c \equiv \{\underline{\nu}, \overline{\nu}\} \setminus \{\nu^e\}$ as well as for $\nu = \nu^\star$, since the latter is in the interior of $\mathcal{N}$. Decompose $\pi_J = (1-\mu)\,\pi_J^M + \mu\,\pi_J^B$ as in (7), where $\pi_J^M \equiv G(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^M))$ and $\pi_J^B \equiv G(\Delta u(\bar{\alpha}, x_J, x_{\neg J}, t_J, t_{\neg J}; \nu^B))$ denote winning probabilities in the all-motivated-reasoner and all-Bayesian electorates respectively, with Bayesians at $\nu^B = \nu^\star$ if some party offers $\nu^\star$, otherwise at $\nu^B = \nu^{\mathrm{PM}}$.

Part (1): $J^M = J^B$.

We start by noting that both parties cannot choose the same ν in equilibrium: (i) $\nu = \nu^e$: J^M deviates to $\nu^\star$. The median motivated reasoner, comparing $\{\nu^e, \nu^\star\}$, retains ν^e (offered by $\neg J^M$, still the voter-utility maximum), so $\pi_{J^M}^M$ is unchanged; Bayesians shift to $\nu^\star$, strictly raising $\pi_{J^M}^B$ by definition of $J^B = J^M$. (ii) $\nu = \nu^\star$: J^M deviates to ν^e. The median motivated reasoner shifts to ν^e, raising $\pi_{J^M}^M$ strictly; Bayesians remain at $\nu^\star$ (still offered by $\neg J^M$). (iii) $\nu \notin \{\nu^e, \nu^\star\}$: J^M deviates to ν^e, raising $\pi_{J^M}^M$ with the Bayesian channel unchanged.

As a candidate separating equilibrium, consider the profile $\nu_{J^M} = \nu^e$, $\nu_{\neg J^M} = \nu^c$, with policies $\hat{x}^0$ and Bayesians at ν^{PM}. The median motivated reasoner adopts ν^e (global maximum). Since $J^B = J^M$, the Bayesian channel for $\neg J^M$ is strictly maximized at $\nu^B = \nu^{\mathrm{PM}}$, so $\neg J^M$ does not offer $\nu^\star$. Among non-$\nu^\star$ narratives, $\neg J^M$ is motivated-reasoner-indifferent (median follows J^M), and ν^c is selected to deter J^M's deviation below.

For J^M, the only narrative deviation that affects the Bayesian channel is to $\nu^\star$. Define the worst-case motivated-reasoner loss and the Bayesian gain from such a deviation as

$$\underline{\Delta} \;\equiv\; G\big(\Delta u(\cdot;\nu^e)\big) - \max_{\nu \in \{\nu^c, \nu^\star\}} G\big(\Delta u(\cdot;\nu)\big),$$
$$\Gamma \;\equiv\; G\big(\Delta u(\cdot;\nu^\star)\big) - G\big(\Delta u(\cdot;\nu^{\mathrm{PM}})\big),$$

both strictly positive, the first because ν^e is the global optimum for the motivated median voter, and the second because $J^M = J^B$ so the party by assumption benefits from Bayesian voters adopting the truth. After a deviation to $\nu^\star$, $\nu_{J^M} = \nu^\star$ and $\nu_{\neg J^M} = \nu^c$: the median motivated reasoner adopts the better of $\{\nu^\star, \nu^c\}$, lowering $\pi_{J^M}^M$ by at least $\underline{\Delta}$; Bayesians shift to $\nu^\star$, raising $\pi_{J^M}^B$ by Γ. The deviation is unprofitable if

$$(1-\mu)\underline{\Delta} \;\geq\; \mu\,\Gamma \quad\Longleftrightarrow\quad \mu \;\leq\; \tilde{\mu} \;\equiv\; \frac{\underline{\Delta}}{\underline{\Delta} + \Gamma}. \tag{A.2}$$

Other deviations to any $\nu \notin \{\nu^e, \nu^\star\}$ lower $\pi_{J^M}^M$ without changing the Bayesian channel and

are therefore strictly dominated.

Holding $\underline{\Delta}$ and ν^{PM} fixed, Γ is continuous in $\nu^\star$ on the set $\{\nu : \Delta u(\cdot;\nu) > \Delta u(\cdot;\nu^{\text{PM}})\}$, on which $J^B = J^M$, and $\Gamma \to 0$ as $\nu^\star \to \nu^{\text{PM}}$. Hence for any $\mu < 1$ there exists $\nu^\star \in \mathcal{N} \setminus \{\nu^{\text{PM}}, \underline{\nu}, \overline{\nu}\}$ with Γ small enough that $\tilde{\mu} > \mu$. On this open set of parameters the truth is untold and Bayesians retain ν^{PM}, and the median voter retains their entangled belief. Finally, for a ν^{PM} that implies the same optimal policy as ν^e, no policy deviation is profitable for any $\mu < 1$, and policies coincide with the baseline from the case of only motivated reasoners.

Part (2): $J^M \neq J^B$.

To see why for low enough μ the policy will be the same as at the baseline with only motivated reasoners, note that at $\mu = 0$ the policy is by definition the same. Any payoff from deviating to another policy due to the presence of Bayesian voters is finite and continuous in μ, so for low enough μ deviation is not optimal even if the truth is revealed in equilibrium.

Given that policies are the same as at baseline, J^M strictly prefers ν^e for any $\nu_{\neg J^M}$. If J^M plays ν^e, the median motivated reasoner adopts ν^e since it is the global voter-utility maximum on $\mathcal{N}$. Any deviation $\nu_{J^M} \neq \nu^e$ delivers $\nu^M \neq \nu^e$ and strictly lowers $\pi^M_{J^M}$ since ν^e uniquely maximizes $\Delta u(\cdot;\nu)$; the Bayesian channel either is unchanged (deviation to ν^c or any $\nu \neq \nu^\star$ when $\neg J^M$ already offers $\nu^\star$) or strictly worsens for J^M (deviation to $\nu^\star$, since by $J^B = \neg J^M$, Bayesians' adoption of $\nu^\star$ rather than ν^{PM} helps $\neg J^M$). Thus $\nu^M = \nu^e$.

Given $\nu_{J^M} = \nu^e$, J^B always has an incentive to reveal the truth: the median motivated reasoner adopts ν^e regardless of $\nu_{\neg J^M}$, so $\pi^M_{\neg J^M}$ is fixed. By definition of J^B, revealing the truth increases J^B's probability of winning among Bayesian voters and is hence optimal.

A.6. Proof of Corollary 1

By Proposition 2, implementation is paired with $\overline{\nu}$ and non-implementation with $\underline{\nu}$. Under (8), the median voter's utilities are $U^1_J(\bar{\alpha};\nu) = w - \bar{\alpha} + t_J\lambda$ and $U^0_J(\nu) = w - \rho$. Hence the cutoffs from Proposition 3 are $\hat{\alpha}^J = \underline{\rho} + t_J\overline{\lambda}$, $\check{\alpha}^J_{\underline{\nu}} = \underline{\rho} + t_J\underline{\lambda}$, $\check{\alpha}^J_{\overline{\nu}} = \overline{\rho} + t_J\overline{\lambda}$, $\dot{\alpha}^J = \underline{\rho} + (t_J - t_{\neg J})\overline{\lambda} + t_{\neg J}\underline{\lambda}$, and $\ddot{\alpha}^J = \overline{\rho} + t_J\overline{\lambda} = \check{\alpha}^J_{\overline{\nu}}$. Since $t_L > t_R$ and $\overline{\lambda} > 0$, $\hat{\alpha}^L > \hat{\alpha}^R$. The interval for a polarized equilibrium with Left implementing is $A_R \leq \bar{\alpha} \leq B_L$, where $A_R = \min\{\check{\alpha}^R_{\overline{\nu}}, \max\{\hat{\alpha}^L, \dot{\alpha}^R\}\} = \min\{\check{\alpha}^R_{\overline{\nu}}, \hat{\alpha}^L\}$ and $B_L = \max\{\check{\alpha}^L_{\underline{\nu}}, \min\{\hat{\alpha}^L, \ddot{\alpha}^L\}\} = \hat{\alpha}^L$. Thus a non-degenerate polarized equilibrium exists iff $\check{\alpha}^R_{\overline{\nu}} < \hat{\alpha}^L$, which is exactly condition (11); when it holds, Proposition 3 gives the unique polarized equilibrium $(x_L, x_R) = (1, 0)$ on $\bar{\alpha} \in [\check{\alpha}^R_{\overline{\nu}}, \hat{\alpha}^L]$, with Left offering $\overline{\nu}$ and Right offering $\underline{\nu}$. The opposite polarized profile would require $A_L \leq \bar{\alpha} \leq B_R$, but $B_R = \hat{\alpha}^R$ and $A_L \geq \hat{\alpha}^R$, so there is no non-degenerate interval.

Throughout this range, the median voter adopts $\overline{\nu}$. The median voter's Left–Right utility margin is therefore $D(\bar{\alpha}) = U^1_L(\bar{\alpha};\overline{\nu}) - U^0_R(\overline{\nu}) = \overline{\rho} + t_L\overline{\lambda} - \bar{\alpha}$. Since $\bar{\alpha} \leq \hat{\alpha}^L = \underline{\rho} + t_L\overline{\lambda}$ and $\overline{\rho} > \underline{\rho}$, we have $D(\bar{\alpha}) > 0$. Hence $\pi_L = G(D(\bar{\alpha})) > 1/2 > \pi_R$. By the speech rule (13) and

strict monotonicity of H, left politicians speak with strictly higher probability than right politicians. Conditional on speaking, politicians voice their party's equilibrium narrative, so Left voices $\overline{\nu}$ and Right voices $\underline{\nu}$. Finally, $D'(\bar{\alpha}) = -1$. Since G and H are strictly increasing, an increase in $\bar{\alpha}$ strictly lowers $H(\pi_L(\bar{\alpha}))$ and strictly raises $H(\pi_R(\bar{\alpha}))$.

A.7. Statement and proof of Lemma A.2

Lemma A.2 (Primitive cutoff ordering)**.** *For each party $J \in \{L, R\}$, the cutoffs defined in the proof of Proposition 3 satisfy*

$$\check{\alpha}^J_{\underline{\nu}} < \check{\alpha}^J_{\overline{\nu}}, \qquad \dot{\alpha}^J < \ddot{\alpha}^J, \qquad \dot{\alpha}^J < \check{\alpha}^J_{\overline{\nu}}, \qquad \check{\alpha}^J_{\underline{\nu}} < \ddot{\alpha}^J.$$

Proof. Fix $J \in \{L, R\}$. For each party K, define the narrative-shift terms

$$\Delta^1_K(\bar{\alpha}) \equiv U^1_K(\bar{\alpha}; \overline{\nu}) - U^1_K(\bar{\alpha}; \underline{\nu}), \qquad \Delta^0_K \equiv U^0_K(\overline{\nu}) - U^0_K(\underline{\nu}).$$

The per-policy sign restrictions $\partial_\lambda y(\alpha, 1, t; \lambda) \geq 0$, $\partial_\lambda y(\alpha, 0, t; \lambda) \leq 0$, $\partial_\rho z(1; \rho) \geq 0$, and $\partial_\rho z(0; \rho) \leq 0$ imply $\Delta^1_K(\bar{\alpha}) \geq 0$ and $\Delta^0_K \leq 0$, and the assumption that one of these inequalities is strict in each dimension implies that every gap of the form $\Delta^1_K(\bar{\alpha}) - \Delta^0_M$ is strictly positive for $K, M \in \{L, R\}$. Moreover, by Assumption 1.(iii), both maps $\bar{\alpha} \mapsto U^1_J(\bar{\alpha}; \overline{\nu})$ and $\bar{\alpha} \mapsto U^1_J(\bar{\alpha}; \underline{\nu})$ are strictly decreasing.

First compare the two within-party cutoffs. Translating the cutoff $\check{\alpha}^J_{\underline{\nu}}$ onto the $\overline{\nu}$ curve gives $U^1_J(\check{\alpha}^J_{\underline{\nu}}; \overline{\nu}) = U^0_J(\underline{\nu}) + \Delta^1_J(\check{\alpha}^J_{\underline{\nu}})$, whereas $U^1_J(\check{\alpha}^J_{\overline{\nu}}; \overline{\nu}) = U^0_J(\overline{\nu}) = U^0_J(\underline{\nu}) + \Delta^0_J$. Since $\Delta^1_J(\check{\alpha}^J_{\underline{\nu}}) \geq 0 \geq \Delta^0_J$, with strict inequality in the height comparison, the strictly decreasing $\overline{\nu}$ curve reaches the first height at a lower value of $\bar{\alpha}$. Hence $\check{\alpha}^J_{\underline{\nu}} < \check{\alpha}^J_{\overline{\nu}}$.

Next compare the two cross-narrative cutoffs. Rearranging the definition of $\dot{\alpha}^J$ yields $U^1_J(\dot{\alpha}^J; \overline{\nu}) = U^0_J(\underline{\nu}) + \Delta^1_{\neg J}(\dot{\alpha}^J)$, while rearranging the definition of $\ddot{\alpha}^J$ yields $U^1_J(\ddot{\alpha}^J; \overline{\nu}) = U^0_J(\underline{\nu}) + \Delta^0_{\neg J}$. Again $\Delta^1_{\neg J}(\dot{\alpha}^J) \geq 0 \geq \Delta^0_{\neg J}$, with a strict height comparison. Since $U^1_J(\cdot; \overline{\nu})$ is strictly decreasing, $\dot{\alpha}^J < \ddot{\alpha}^J$.

The same level-crossing argument gives the two cross-corner inequalities. The cutoff $\dot{\alpha}^J$ lies on $U^1_J(\cdot; \overline{\nu})$ at height $U^0_J(\underline{\nu}) + \Delta^1_{\neg J}(\dot{\alpha}^J)$, whereas $\check{\alpha}^J_{\overline{\nu}}$ lies on the same curve at height $U^0_J(\underline{\nu}) + \Delta^0_J$. Since $\Delta^1_{\neg J}(\dot{\alpha}^J) \geq 0 \geq \Delta^0_J$, with strict inequality in the height comparison, strict monotonicity gives $\dot{\alpha}^J < \check{\alpha}^J_{\overline{\nu}}$. Finally, translate $\ddot{\alpha}^J$ onto the $\underline{\nu}$ curve: $U^1_J(\ddot{\alpha}^J; \underline{\nu}) = U^0_J(\underline{\nu}) + \Delta^0_{\neg J} - \Delta^1_J(\ddot{\alpha}^J)$. The cutoff $\check{\alpha}^J_{\underline{\nu}}$ lies on the same curve at height $U^0_J(\underline{\nu})$. Because $\Delta^0_{\neg J} - \Delta^1_J(\ddot{\alpha}^J) \leq 0$, with strict inequality in the height comparison, the strictly decreasing $\underline{\nu}$ curve reaches the higher height first. Therefore $\check{\alpha}^J_{\underline{\nu}} < \ddot{\alpha}^J$. □

Supplemental Appendix

for

Narrative Entanglement

Adam Brzezinski
LSE

Luis Garicano
LSE and CEPR

B. Additional theory

B.1. Additional results

B.1.1. Sufficient conditions for polarized equilibria

We give sufficient conditions for a unique polarized equilibrium in which parties offer different policies and different narratives. In particular, writing the probability of winning as $\pi_J(x_J, x_{\neg J}, \nu_J, \nu_{\neg J})$, the two following conditions are jointly sufficient for this:

$$\pi_L(0, 0, \underline{\nu}, \underline{\nu}) \approx \pi_R(0, 0, \underline{\nu}, \underline{\nu}), \tag{B.1}$$

$$\min\{\check{\alpha}^L_{\overline{\nu}}, \check{\alpha}^R_{\overline{\nu}}\} < \max\{\hat{\alpha}^L, \hat{\alpha}^R\}. \tag{B.2}$$

Condition (B.1) is a *competitiveness* requirement: in the status quo, neither party has a large electoral advantage.[33] Condition (B.2) requires a significant *tax wedge*, implying that tax differences interact significantly with the policy; the relevant cutoffs $\check{\alpha}^J_{\overline{\nu}}$ and $\hat{\alpha}^J$ are defined in the proof of Proposition 3 in Appendix A. Together, these give sufficient conditions for the existence of a unique polarized equilibrium and an intuitive condition for which party will be the policy implementer in that equilibrium:

Corollary B.1. *Assume that all voters are motivated reasoners ($\mu = 0$), the election is competitive (condition* (B.1) *holds) and the tax wedge between parties matters for policy (condition* (B.2) *holds). Then, there exists a unique polarized equilibrium if and only if costs are in an intermediate range strictly between $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$. In that range, party $J \in \{L, R\}$ is the implementer and not party $\neg J$, i.e., $(x_J, x_{\neg J}) = (1, 0)$, if and only if*

$$y(\hat{\alpha}^J, 1, t_J; \overline{\lambda}) - y(\hat{\alpha}^J, 1, t_{\neg J}; \overline{\lambda}) \;>\; -\Big(y(\cdot, 0, t_J; \underline{\lambda}) - y(\cdot, 0, t_{\neg J}; \underline{\lambda})\Big), \tag{B.3}$$

where $\hat{\alpha}^J$ is the cost at which a voter is indifferent between J implementing the policy and $\neg J$ not implementing the policy, given optimal narratives.

Proof. Let $A_J = \min\{\check{\alpha}^J_{\overline{\nu}}, \max\{\hat{\alpha}^{\neg J}, \dot{\alpha}^J\}\}$ and $B_J = \max\{\check{\alpha}^J_{\underline{\nu}}, \min\{\hat{\alpha}^J, \ddot{\alpha}^J\}\}$. By Proposition 3, a polarized equilibrium in which J implements and $\neg J$ does not exists if and only if $A_{\neg J} \leq \bar{\alpha} \leq B_J$. Under condition (B.1), $\max\{\dot{\alpha}^J, \check{\alpha}^J_{\underline{\nu}}\} \leq \hat{\alpha}^J \leq \min\{\ddot{\alpha}^J, \check{\alpha}^J_{\overline{\nu}}\}$ for each J. Let H be the party with the higher value of $\hat{\alpha}$, and let $K \neq H$. Then $B_H = \hat{\alpha}^H$, $B_K = \hat{\alpha}^K$, $A_H = \max\{\hat{\alpha}^K, \dot{\alpha}^H\}$, and $A_K = \min\{\check{\alpha}^K_{\overline{\nu}}, \hat{\alpha}^H\}$. Since $\check{\alpha}^H_{\overline{\nu}} \geq \hat{\alpha}^H = \max\{\hat{\alpha}^L, \hat{\alpha}^R\}$, condition (B.2) is equivalent to $\check{\alpha}^K_{\overline{\nu}} < \hat{\alpha}^H$. Hence $A_K = \check{\alpha}^K_{\overline{\nu}} < B_H$, so there is a non-empty interval $(\check{\alpha}^K_{\overline{\nu}}, \hat{\alpha}^H)$ on which the polarized equilibrium with H implementing exists. The opposite polarized

[33] Formally, $\max\{\Delta^0_L, \Delta^0_R, -\Delta^1_L(\hat{\alpha}^L), -\Delta^1_R(\dot{\alpha}^L)\} \leq D_{00} \leq \min\{-\Delta^0_L, -\Delta^0_R, \Delta^1_R(\hat{\alpha}^R), \Delta^1_L(\dot{\alpha}^R)\}$ is the precise condition, where the terms are defined in the proofs of Lemma A.2 and Proposition 3. A sufficient condition for this to hold is simply $D_{00} = 0$, which gives $\pi_L(0, 0, \underline{\nu}, \underline{\nu}) = \pi_R(0, 0, \underline{\nu}, \underline{\nu})$.

equilibrium cannot exist on any non-empty interval, since $A_H = \max\{\hat{\alpha}^K, \dot{\alpha}^H\} \geq \hat{\alpha}^K = B_K$. Thus, throughout this interval, the unique polarized equilibrium is $(x_H, x_K) = (1, 0)$. Finally, because $\tilde{\alpha}_1 = \min\{A_L, A_R\}$ and $\tilde{\alpha}_2 = \max\{B_L, B_R\} = \hat{\alpha}^H$, this interval is contained in the intermediate region $(\tilde{\alpha}_1, \tilde{\alpha}_2)$. If in addition $\check{\alpha}^K_{\overline{\nu}} \leq \dot{\alpha}^H$, then $A_K \leq A_H$, so $\tilde{\alpha}_1 = A_K = \check{\alpha}^K_{\overline{\nu}}$. Therefore the whole intermediate region is $(\tilde{\alpha}_1, \tilde{\alpha}_2) = (\check{\alpha}^K_{\overline{\nu}}, \hat{\alpha}^H)$, and the unique equilibrium throughout it is polarized with H implementing. □

The two conditions on which this corollary relies impose some regularity on the underlying parameter space of the model, which disciplines the possible equilibria. The competitiveness condition (B.1) ensures that no party has, *ex ante*, a higher incentive to deviate from a separating equilibrium to the status quo. The second condition ensures that tax policy matters sufficiently such that policy differentiation is possible. If that condition is not met, parties do not sufficiently differ to have different policies and narratives in equilibrium.

The final part of the corollary gives an intuitive statement of which party will be the policy implementer in a separating equilibrium. This depends on how tax policy interacts with the economic narratives. In any equilibrium with policy divergence, the implementing party is that whose tax–narrative combination is more advantageous for implementing the policy. The left-hand side of (B.3) captures J's tax advantage under implementation: the income gain at $\hat{\alpha}^J$ from having tax rate t_J rather than $t_{\neg J}$, with the optimistic economic narrative $\overline{\lambda}$ in place. The right-hand side is the mirror-image quantity for $\neg J$ under the status quo, with tax roles reversed and the pessimistic economic narrative $\underline{\lambda}$ in place; it does not depend on α, since private policy costs are paid only under policy implementation. Party J is the implementer precisely when the first tax channel dominates the second for it and the second dominates the first for party $\neg J$. The non-economic dimension plays no role in (B.3): the z-term depends on (x, ρ) but not on the tax rate or on private costs, so it is common to both parties' platforms within each (x, ρ) pair and cancels.[34]

B.1.2. Predictions for speech probabilities in the general model

Proposition 3 characterizes equilibrium winning probabilities as a function of the median voter's private cost $\bar{\alpha}$ when all voters are motivated reasoners. We focus on ranges of $\bar{\alpha}$ on which the policy game has a unique pure-strategy divergent equilibrium, and continue with our assumption that parties use the narrative that makes their policy look best when indifferent between narratives. Let $J^* \in \{L, R\}$ denote the implementing party in that equilibrium and let $\neg J^*$ denote the opposing party.

[34] The non-economic narrative bounds $\overline{\rho}$ and $\underline{\rho}$ still shape the *location* of the divergent-policy region by entering the indifference equations that define the cross-cutoffs $\hat{\alpha}^L, \hat{\alpha}^R$, but they drop out of the condition that selects which party is the implementer once divergence is the equilibrium outcome.

Proposition B.1. *Assume that all voters are motivated reasoners ($\mu = 0$), and consider an interval of median costs on which Proposition 3 delivers a unique pure-strategy divergent equilibrium with J^* implementing and $\neg J^*$ opposing.*

1. *Politicians of the electorally advantaged party are more likely to speak on the policy.*
2. *Within this divergent-equilibrium interval, an increase in $\bar{\alpha}$ weakly decreases the probability that politicians of J^* speak on the policy and weakly increases the probability that politicians of $\neg J^*$ speak on the policy. These changes are strict whenever the corresponding winning probabilities change strictly.*

Proof. By the speech setup, a politician of party J speaks on the policy if and only if $\pi_J(\bar{\alpha}) \geq \omega_j$. Since $\omega_j \sim H$, the speech probability is $\Pr(c_j = 1 \mid J) = H(\pi_J(\bar{\alpha}))$. Because H is strictly increasing, one party's politicians speak with higher probability than the other's if and only if that party has the higher equilibrium winning probability. Since $\pi_{\neg J} = 1 - \pi_J$, the implementing party is more vocal if and only if $\pi_{J^*}(\bar{\alpha}) > 1/2$. This proves part 1.

For part 2, fix the divergent equilibrium with $x_{J^*} = 1$ and $x_{\neg J^*} = 0$. By Proposition 2, J^* offers $\overline{\nu}$ and $\neg J^*$ offers $\underline{\nu}$. For either adopted narrative $\nu \in \{\underline{\nu}, \overline{\nu}\}$, the median voter's margin for J^* is $D_\nu(\bar{\alpha}) = u(\bar{\alpha}, 1, t_{J^*}; \nu) - u(\bar{\alpha}, 0, t_{\neg J^*}; \nu)$, which is strictly decreasing in $\bar{\alpha}$ by Assumption 1.(iii). By the adoption cutoff $\hat{\alpha}^{J^*}$ defined in the proof of Proposition 3, if the median voter's adopted narrative changes as $\bar{\alpha}$ rises, the change is from $\overline{\nu}$ to $\underline{\nu}$, and $D_{\overline{\nu}}(\bar{\alpha}) - D_{\underline{\nu}}(\bar{\alpha}) = [u(\bar{\alpha}, 1, t_{J^*}; \overline{\nu}) - u(\bar{\alpha}, 1, t_{J^*}; \underline{\nu})] + [u(\bar{\alpha}, 0, t_{\neg J^*}; \underline{\nu}) - u(\bar{\alpha}, 0, t_{\neg J^*}; \overline{\nu})] \geq 0$, where the first bracket is non-negative by $\partial_\lambda y(\alpha, 1, t; \lambda) \geq 0$ and $\partial_\rho z(1; \rho) \geq 0$, and the second bracket is non-negative by $\partial_\lambda y(\alpha, 0, t; \lambda) \leq 0$ and $\partial_\rho z(0; \rho) \leq 0$. Thus the realized median-voter margin for J^*, denoted $D(\bar{\alpha})$, is weakly decreasing in $\bar{\alpha}$. Since G is a c.d.f., $\pi_{J^*} = G(D(\bar{\alpha}))$ is weakly decreasing, so $\pi_{\neg J^*} = 1 - \pi_{J^*}$ is weakly increasing. Applying the increasing function H gives the stated speech-probability comparative static. □

Part 1 of the proposition separates the policy position from the speech incentive. A politician's payoff from speaking on the policy is their party's equilibrium winning probability, so the more vocal party is the party with the larger electoral return from making the policy salient. In the general model this need not always be the implementing party: the implementing party is more vocal exactly when its platform is electorally advantaged at the equilibrium profile.

Part 2 discusses how the probability of making a speech changes with $\bar{\alpha}$. An increase in $\bar{\alpha}$ raises the median voter's private cost of implementation and thereby erodes the implementing party's relative appeal. Hence, politicians of J^* become weakly less likely to speak on the policy while politicians of $\neg J^*$ become weakly more likely to do so.

Parts 1 and 2 imply a prediction for the aggregate balance of narratives voiced on the policy. By Proposition 2, each party's narrative in the divergent equilibrium is an entangled narrative that supports that party's own policy: $\overline{\nu}$ for J^* and $\underline{\nu}$ for $\neg J^*$. An increase in $\bar{\alpha}$ weakly shifts the aggregate balance of voiced narratives toward $\underline{\nu}$ as the share of speeches from $\neg J^*$ politicians rises and the share from J^* politicians falls.

B.2. Worked application to climate policy

Using the utility function for the climate policy application as given in equation (8), and using notation from the proof of Proposition 3, we write the median voter's utilities as:

$$U_J^1(\bar{\alpha};\nu) = w - \bar{\alpha} + t_J\lambda, \tag{B.4}$$

$$U_J^0(\nu) = w - \rho. \tag{B.5}$$

Two simplifying features distinguish this application: U_J^1 does not depend on ρ (environmental damages are avoided whenever the policy is implemented), and U_J^0 does not depend on J (the status-quo income is the pre-tax wage, regardless of the governing party). Throughout we maintain the sign restrictions $\underline{\lambda} < 0 < \overline{\lambda}$, $0 \leq \underline{\rho} < \overline{\rho}$, and $0 < t_R < t_L$.

We can then calculate the median-voter narrative cutoffs (see the proof of Proposition 3) defined by $\hat{\alpha}^J : U_J^1(\hat{\alpha}^J;\overline{\nu}) = U_{\neg J}^0(\underline{\nu})$. This gives us

$$\hat{\alpha}^L = \underline{\rho} + t_L\overline{\lambda} \qquad \hat{\alpha}^R = \underline{\rho} + t_R\overline{\lambda}, \tag{B.6}$$

where we can see that $\hat{\alpha}^L > \hat{\alpha}^R$.

Similarly, using the definitions of the other cutoffs from the proof of Proposition 3, we obtain

$$\check{\alpha}_{\underline{\nu}}^J = \underline{\rho} + t_J\underline{\lambda}, \qquad \check{\alpha}_{\overline{\nu}}^J = \overline{\rho} + t_J\overline{\lambda}, \tag{B.7}$$

$$\dot{\alpha}^J = \underline{\rho} + (t_J - t_{\neg J})\overline{\lambda} + t_{\neg J}\underline{\lambda}, \qquad \ddot{\alpha}^J = \overline{\rho} + t_J\overline{\lambda}. \tag{B.8}$$

Writing these out explicitly for L and R:

$$\check{\alpha}_{\underline{\nu}}^L = \underline{\rho} + t_L\underline{\lambda}, \qquad \check{\alpha}_{\underline{\nu}}^R = \underline{\rho} + t_R\underline{\lambda}, \tag{B.9}$$

$$\check{\alpha}_{\overline{\nu}}^L = \overline{\rho} + t_L\overline{\lambda}, \qquad \check{\alpha}_{\overline{\nu}}^R = \overline{\rho} + t_R\overline{\lambda}, \tag{B.10}$$

$$\dot{\alpha}^L = \underline{\rho} + t_L\overline{\lambda} - t_R(\overline{\lambda} - \underline{\lambda}), \qquad \dot{\alpha}^R = \underline{\rho} - (t_L - t_R)\overline{\lambda} + t_L\underline{\lambda}, \tag{B.11}$$

$$\ddot{\alpha}^L = \overline{\rho} + t_L\overline{\lambda} = \check{\alpha}_{\overline{\nu}}^L, \qquad \ddot{\alpha}^R = \overline{\rho} + t_R\overline{\lambda} = \check{\alpha}_{\overline{\nu}}^R. \tag{B.12}$$

Note that $\ddot{\alpha}^J = \check{\alpha}_{\overline{\nu}}^J$: this collapse arises because U_J^0 is J-independent, so the left-hand side of the defining equation for $\ddot{\alpha}^J$ vanishes.

Direct computation gives the unconditional ordering

$$\dot{\alpha}^R < \check{\alpha}_{\underline{\nu}}^L < \check{\alpha}_{\underline{\nu}}^R < \min\{\dot{\alpha}^L, \hat{\alpha}^R\} \leq \max\{\dot{\alpha}^L, \hat{\alpha}^R\} < \hat{\alpha}^L < \check{\alpha}_{\overline{\nu}}^L, \tag{B.13}$$

in line with Lemma A.2. The remaining cutoff $\check{\alpha}^R_{\overline{\nu}}$ satisfies $\hat{\alpha}^R < \check{\alpha}^R_{\overline{\nu}} < \check{\alpha}^L_{\overline{\nu}}$ unconditionally, but its rank relative to $\dot{\alpha}^L$ and to $\hat{\alpha}^L$ is not pinned down.

Equilibria follow Proposition 3 with the simplification that the bounds are given by

$$\tilde{\alpha}_1 = \min\big\{\max\{\hat{\alpha}^R, \dot{\alpha}^L\},\, \check{\alpha}^R_{\overline{\nu}}\big\} \qquad \text{and} \qquad \tilde{\alpha}_2 = \hat{\alpha}^L.$$

So we have $(x_L, x_R) = (1, 1)$ iff $\bar{\alpha} < \tilde{\alpha}_1$ and $(x_L, x_R) = (0, 0)$ iff $\bar{\alpha} > \tilde{\alpha}_2$, parallel to Cases 1 and 2 of Proposition 3.

For Case 3, i.e., $\bar{\alpha} \in (\tilde{\alpha}_1, \tilde{\alpha}_2)$, there is a mixed-strategy equilibrium on the (possibly empty) sub-range $\bar{\alpha} \in (\tilde{\alpha}_1, \min\{\check{\alpha}^R_{\overline{\nu}}, \hat{\alpha}^L\})$, and a unique polarized equilibrium $(x_L, x_R) = (1, 0)$ on the (possibly empty) sub-range $\bar{\alpha} \in (\check{\alpha}^R_{\overline{\nu}}, \hat{\alpha}^L)$. The latter is non-empty if and only if

$$\hat{\alpha}^L > \check{\alpha}^R_{\overline{\nu}} \iff (t_L - t_R)\overline{\lambda} > \overline{\rho} - \underline{\rho}, \tag{B.14}$$

i.e., if and only if the tax wedge between parties is sufficiently large relative to the spread $\overline{\rho} - \underline{\rho}$. Sufficiently different taxes mean sufficiently different parties, such that it can be optimal for them to pursue different policies in equilibrium.

C. Additional empirics

C.1. Additional descriptive tables and figures

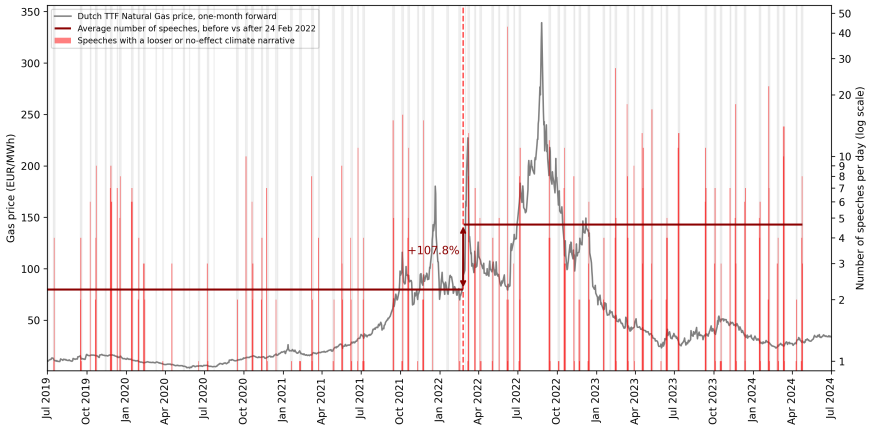

Note: This figure shows the same information as Figure 1 but with added data on Dutch TTF natural gas price as traded on the Chicago Mercantile Exchange, drawn from: *https://www.investing.com/commodities/dutch-ttf-gas-c1-futures-historical-data*.

Figure C.1 – Number of "not needed" speeches and average daily gas price

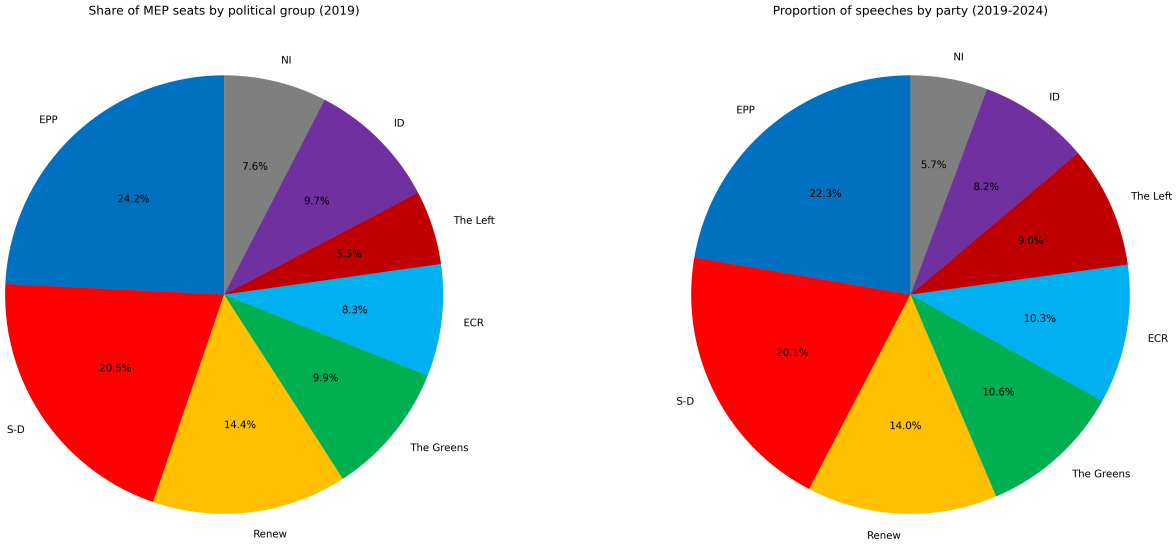

(a) MEP seat share by party (2019)

(b) Speeches by party (2019-2024)

Note: Seat shares are based on the constitutive session of the 9th European Parliament in 2019, drawn from https://results.elections.europa.eu/en/european-results/2019-2024/constitutive-session/. Speech shares are based on our speech dataset across 2019–2024.

Figure C.2 – Seat share versus speech share by party

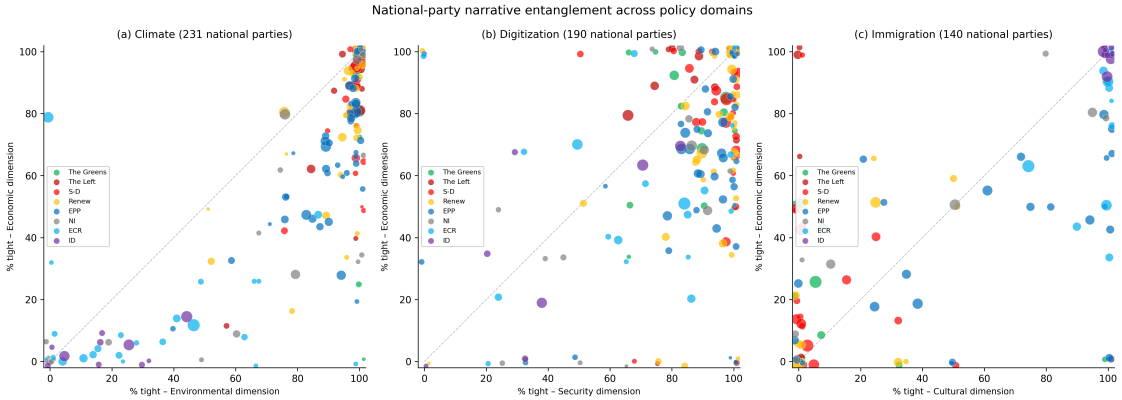

(a) National-party level

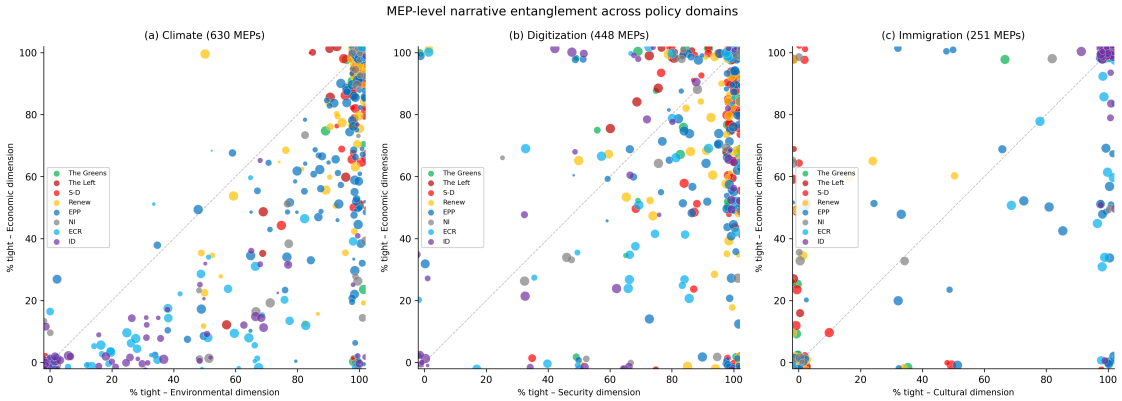

(b) MEP level

Note: Each point is a national party in panel (a) and an MEP in panel (b), where a jitter is applied to points to unstack otherwise overlapping observations. For each policy domain, parties/MEPs with at least one narrative in each dimension are shown; the x-axis is the share of the unit's domain-relevant speeches with a "tighter" narrative on the non-economic dimension (environmental, security, or cultural), and the y-axis is the analogous share on the economic dimension.

Figure C.3 – Narrative entanglement across policy dimensions: national-party and MEP levels

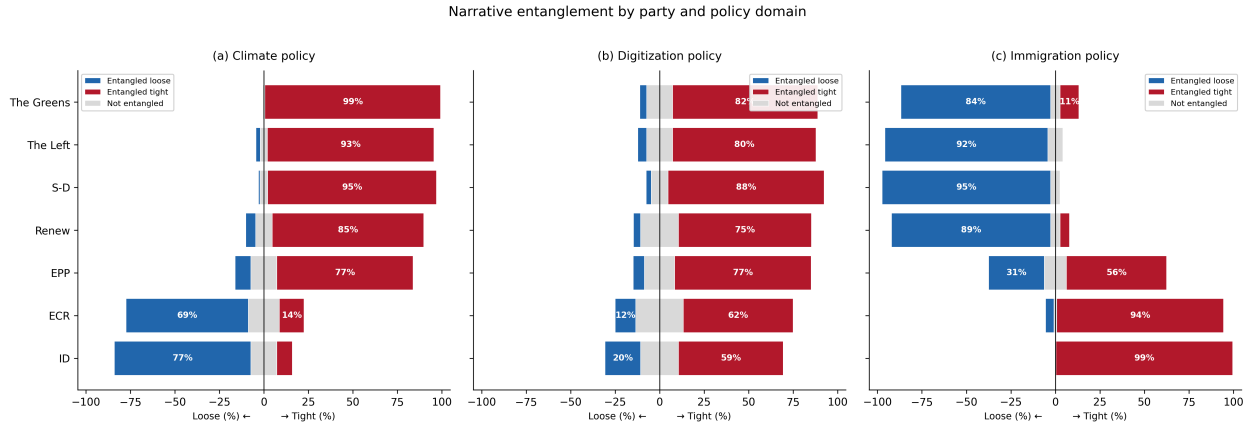

Note: Share of each party group's speeches that are relevant on both dimensions of the domain, classified as *entangled loose* (looser or no effect on both the policy and economic dimension), *entangled tight* (tighter on both), or *not entangled* (mixed). Bars are drawn diverging from zero: entangled-loose shares extend leftwards and entangled-tight shares rightwards.

Figure C.4 – Narrative entanglement by party group and policy domain

Table C.1 – EU party group combinations across MEPs and speeches

Combination	Number of MEPs	Number of Speeches
EPP only	190	10,525
S-D only	169	9,537
Renew only	120	6,341
The Greens only	85	4,818
ECR only	70	4,691
ID only	67	3,665
NI only	48	1,583
The Left only	43	4,252
EPP and NI	12	618
ID and NI	7	279
NI and The Greens	5	244
NI and S-D	4	181
EPP and ID	4	52
Renew and S-D	3	151
EPP and Renew	3	204
ECR and ID	2	84
NI and The Left	2	17
ECR and EPP	2	67
ECR and ID and NI	2	107
ECR and NI	2	128
Renew and The Greens	2	185
ECR and S-D	1	87
The Greens and The Left	1	69
EPP and S-D	1	36
S-D and The Greens	1	60
NI and Renew	1	217
ECR and Renew	1	41
Total	848	48,239

C.2. Robustness and alternative mechanisms

Descriptive overview of narrative changes. Figure C.5 descriptively summarizes the narrative shifts by party group, plotting unconditional pre- and post-invasion shares of climate-skeptic and negative economic speeches as a fraction of all speeches. Comparing with our model predictions set out in Section 4.1, the data align with our framework. On both the climate-skeptic dimension (panel a) and the negative economic dimension (panel b), ECR and ID show the largest shifts toward climate skepticism and negative economic narratives.

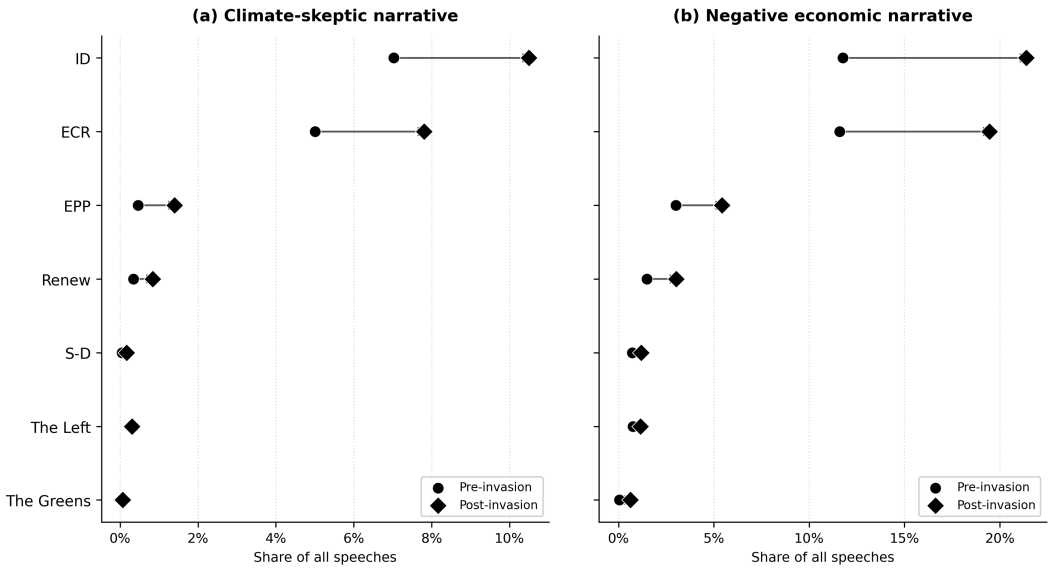

Note: Each arrow connects the pre-invasion share (circle) to the post-invasion share (diamond) of speeches with the indicated narrative, as a fraction of all speeches. Panel (a): share of speeches with the "looser" or "no effect" non-economic narrative. Panel (b): analogous share for the economic narrative.

Figure C.5 – Narrative shifts by party: before vs. after Russia's invasion

Event-study design. The main analysis uses a binary $Post_t$ indicator to compare speeches before and after Russia's invasion. We now estimate the model more flexibly by replacing $Post_t$ with a set of year dummies centered on the invasion date. Specifically, we define Year 0 as the 365-day window centered on 24 February 2022, with preceding years Y_{-1}, Y_{-2} and subsequent years Y_{+1}, Y_{+2} each spanning 365 days. The omitted category is Y_{-1}, the last year entirely before the invasion window.

Figure C.6 plots the year-by-year coefficients for the negative economic and climate-skeptic narrative outcomes, estimated with MEP fixed effects, separately for MEPs from the three parties that saw significant shifts in the main regression from Table 2 (panels (a) and (b)) and the remaining MEPs (panels (c) and (d)). The coefficients on Y_{-2} are statistically

insignificant for 3 out of 4 panels and negative but close to zero for panel (a). The invasion year (Y_0) and subsequent years show a clear upward shift that is individually significant; however, the changes in panels (a) and (b) are far larger, showing a sharp post-invasion break in both dimensions. Shifts in panels (c) and (d) are much smaller, and partly driven by NI non-affiliates.

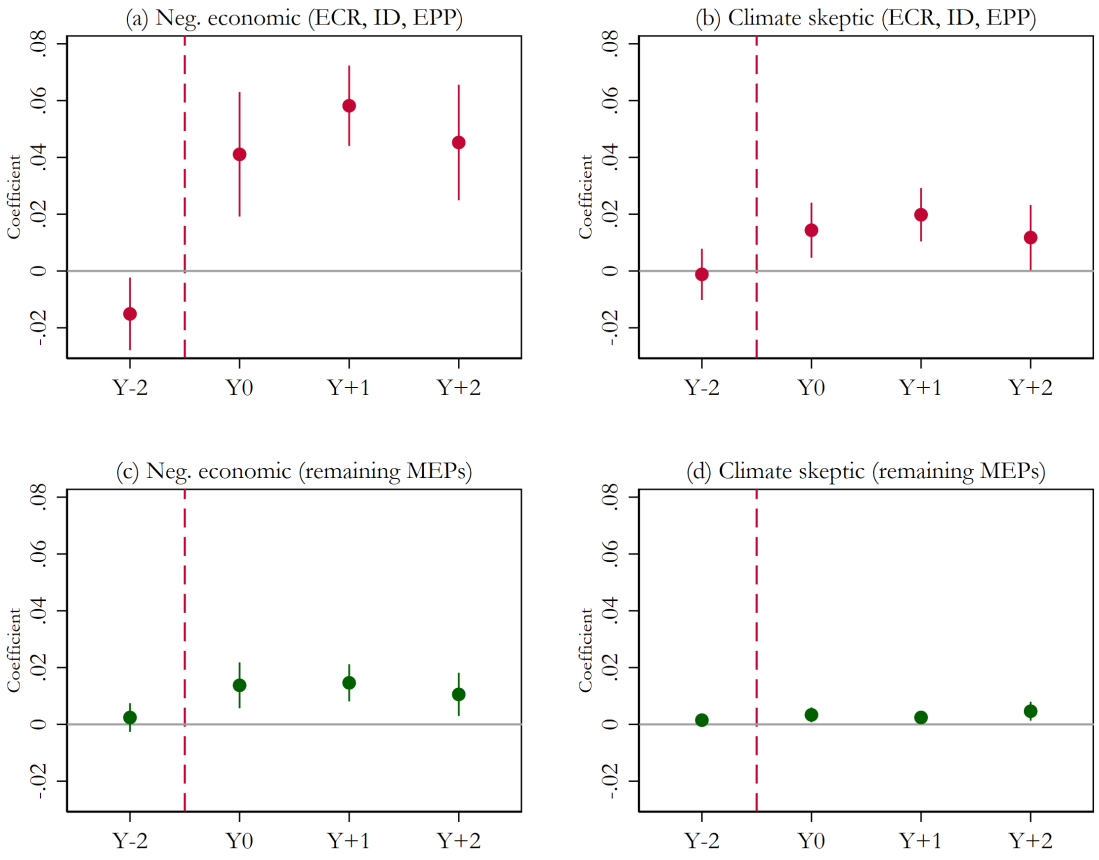

Note: Event-study version of the main regression from Table 2 with MEP fixed effects. Panels (a) and (c) investigate negative economic narratives, panels (b) and (d) the climate-skeptic environmental narratives; top panels restricted to sample of MEPs from ECR, ID, and EPP; bottom panels restricted to remaining MEPs; years are 365-day windows centered on the invasion date (24 February 2022); Y_0 spans the invasion $\pm$ 182 days. All four panels share a common y-axis scale. Base (omitted): Y_{-1}. Horizontal bars show 95% confidence intervals, based on standard errors clustered by country times EU party group.

Figure C.6 – Event-study by party subsample

Partialing out MEPs who switch parties. The main results in Table 2 always assign an MEP's speech to their current party. Table C.2 instead partials out MEPs who switch parties by pooling everyone affiliated with two or more party groups during the 9th Parliament into a single "2 or more parties" interaction. ID's post-invasion shift toward negative economic

narratives remains large and strongly significant, and ECR's shift attenuates only mildly while staying positive and significant. The pooled switcher coefficient is small and only marginally significant, confirming that switchers do not drive the within-group far-right response.

National party fixed effects. The main Table 2 uses MEP fixed effects. Table C.3 shows that results are similar when instead accounting for national-party fixed effects, so that the identifying variation is instead within the same national party. National party fixed effects assign MEPs with two or more national parties a separate fixed effect for their combination. The post-invasion coefficients and their party-group interactions are quantitatively very similar across the two specifications, consistent with most of the cross-party variation in climate narratives being aligned with national-party rather than individual-MEP characteristics.

Extensive versus intensive margin. The post-invasion shifts documented in Section 4.3 combine an extensive and intensive margin: MEPs may speak about climate or the economy more often, and, conditional on doing so, may adopt a different narrative. Columns 1–3 of Table 2 capture the extensive margin of making a speech. Table C.4 isolates the intensive margin by restricting the sample to domain-relevant speeches: on climate, the narrative shift toward looser/no-effect narratives is driven by the ECR and the EPP (the ID, which was already overwhelmingly skeptic before the invasion, has little room to move further), and on the economic dimension the shift is carried by the EPP and Renew. Overall, the intensive margin moves as predicted by our model toward more negative economic narratives and more climate-skeptic narratives.

Mechanical crowding out of climate policy because of Russia's invasion of Ukraine. We suggest that Russia's invasion of Ukraine changed economic narratives on climate policies and, via entanglement, environmental narratives. A natural alternative explanation is that the changes in climate policy narratives are more mechanical: parties may speak less frequently and less favorably about climate policy simply because Ukraine now absorbs some of their limited speaking time. This would predict that all parties reduce their speeches on climate policy, and those parties that reduce their share of speeches on climate policy more are also those that talk more about Ukraine. Figure C.7 argues against this by showing the change in speeches mentioning Ukraine versus the change in climate-relevant speeches at the national party level; Table C.5 shows similar results for the EU party group level. First, almost all national parties deliver more speeches on the topic of Ukraine after the invasion, yet many parties also deliver more speeches on the topic of climate policy. Second, the parties driving the climate-skeptic shift — ECR and ID — *raise* their share of speeches on climate after the invasion even as Ukraine mentions become more common, which is the

Table C.2 – Regression results separating MEPs with 2+ parties

	Speech on climate			Negative economic narrative			Climate skeptic narrative		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post	0.019***	-0.008	-0.010	0.028***	0.005***	0.002	0.008***	0.000	-0.002
	(0.006)	(0.019)	(0.017)	(0.004)	(0.001)	(0.005)	(0.002)	(0.001)	(0.004)
Post × ECR		0.073***	0.092***		0.066***	0.078***		0.019**	0.021**
		(0.022)	(0.020)		(0.010)	(0.010)		(0.009)	(0.009)
Post × EPP		0.030	0.037**		0.024***	0.029***		0.010***	0.010***
		(0.021)	(0.019)		(0.006)	(0.006)		(0.002)	(0.004)
Post × ID		0.105***	0.115***		0.092***	0.091***		0.026**	0.021**
		(0.025)	(0.026)		(0.014)	(0.013)		(0.010)	(0.008)
Post × NI		0.054**	0.068***		0.038*	0.039*		0.008	0.005
		(0.024)	(0.025)		(0.020)	(0.022)		(0.008)	(0.007)
Post × Renew		-0.002	0.002		0.013**	0.014*		0.006*	0.003
		(0.022)	(0.020)		(0.006)	(0.007)		(0.003)	(0.004)
Post × S-D		0.009	0.017		-0.002	0.004		0.001	0.000
		(0.021)	(0.017)		(0.002)	(0.005)		(0.001)	(0.003)
Post × The Left		-0.001	0.003		-0.003	0.003		-0.000	0.001
		(0.024)	(0.022)		(0.006)	(0.006)		(0.003)	(0.004)
Post × 2 or more parties		0.021	0.032		0.029	0.023*		0.010**	0.005
		(0.029)	(0.027)		(0.018)	(0.014)		(0.005)	(0.006)
Observations	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239
R^2	0.21	0.21	0.21	0.19	0.19	0.19	0.13	0.13	0.13
Mean Dep. Var.	0.18	0.18	0.18	0.050	0.050	0.050	0.019	0.019	0.019
MEP FE	X	X	X	X	X	X	X	X	X
Country X Post			X			X			X

Note: The specifications are the same as in Table 2, except that the interactions with party additionally distinguish MEPs that were part of two or more parties.

Table C.3 – Regression results with national party fixed effects

	Speech on climate			**Negative economic narrative**			**Climate skeptic narrative**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post	0.017***	-0.008	-0.021	0.027***	0.005***	0.006	0.008***	0.000	-0.001
	(0.006)	(0.016)	(0.014)	(0.005)	(0.001)	(0.005)	(0.002)	(0.001)	(0.004)
Post × ECR		0.068***	0.081***		0.061***	0.066***		0.018**	0.017**
		(0.017)	(0.017)		(0.006)	(0.008)		(0.008)	(0.008)
Post × EPP		0.029	0.034**		0.020***	0.024***		0.009***	0.008**
		(0.018)	(0.014)		(0.006)	(0.007)		(0.003)	(0.004)
Post × ID		0.112***	0.118***		0.095***	0.095***		0.033***	0.028***
		(0.020)	(0.020)		(0.011)	(0.011)		(0.009)	(0.008)
Post × NI		0.055***	0.075***		0.035**	0.032*		0.009	0.004
		(0.020)	(0.021)		(0.017)	(0.016)		(0.006)	(0.006)
Post × Renew		-0.003	0.005		0.010**	0.011**		0.005*	0.002
		(0.020)	(0.017)		(0.005)	(0.005)		(0.003)	(0.004)
Post × S-D		0.005	0.016		-0.003	0.002		0.000	0.000
		(0.018)	(0.015)		(0.002)	(0.004)		(0.001)	(0.003)
Post × The Left		-0.005	-0.003		-0.001	0.004		-0.001	0.000
		(0.020)	(0.019)		(0.005)	(0.005)		(0.003)	(0.003)
Observations	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239	48,239
R^2	0.069	0.070	0.071	0.11	0.11	0.11	0.081	0.082	0.083
Mean Dep. Var.	0.18	0.18	0.18	0.050	0.050	0.050	0.019	0.019	0.019
National Party FE	X	X	X	X	X	X	X	X	X
Country X Post			X			X			X

Note: The specifications are the same as in Table 2, but with national party fixed effects in place of MEP fixed effects.

Table C.4 – Intensive margin: narrative shifts conditional on climate speeches

	Negative economic narrative			Climate skeptic narrative		
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.064***	0.026***	-0.025	0.034***	0.002	-0.001
	(0.011)	(0.006)	(0.023)	(0.008)	(0.003)	(0.025)
Post × ECR		0.035	0.041		0.103**	0.100*
		(0.022)	(0.035)		(0.041)	(0.054)
Post × EPP		0.111***	0.130***		0.074***	0.086***
		(0.027)	(0.026)		(0.016)	(0.019)
Post × ID		-0.018	-0.022		0.027	0.034
		(0.012)	(0.018)		(0.059)	(0.050)
Post × NI		0.054	0.001		0.111*	0.070
		(0.069)	(0.062)		(0.059)	(0.046)
Post × Renew		0.070***	0.086***		0.031**	0.027
		(0.023)	(0.026)		(0.016)	(0.020)
Post × S-D		0.002	0.035		0.004	0.014
		(0.015)	(0.022)		(0.005)	(0.013)
Post × The Left		-0.028	0.005		0.000	0.016
		(0.041)	(0.042)		(0.021)	(0.021)
Observations	6,793	6,793	6,793	6,945	6,945	6,945
R^2	0.67	0.67	0.67	0.65	0.65	0.65
Mean Dep. Var.	0.36	0.36	0.36	0.13	0.13	0.13
MEP FE	X	X	X	X	X	X
Country x Post			X			X

Note: The specifications are the same as in Table 2, but the sample is restricted to speeches that voice a narrative on the relevant dimension rather than covering all speeches.

opposite sign of the mechanical crowd-out prediction. Third, the climate-speaking shares of some left-leaning parties do fall somewhat, but too little to mechanically generate the observed narrative shifts. These patterns are not consistent with a mechanical crowd-out effect, but are consistent with the strategic silence our model predicts for left-wing parties and the strategic shift toward more vocal climate skepticism among right-wing parties.

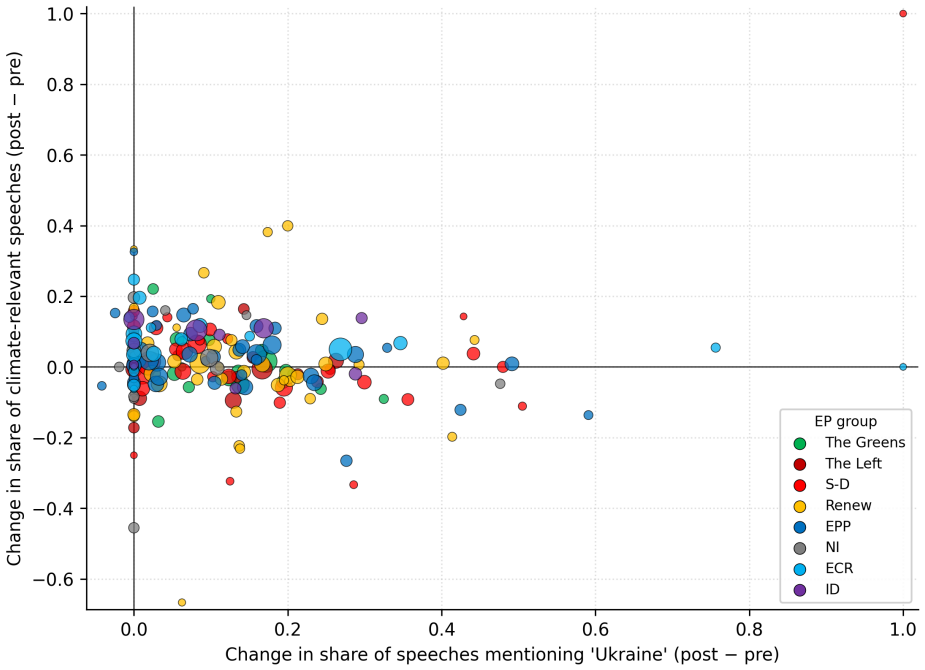

Note: Each point is a national party; colors denote EU party group (the mode for those national parties that have more than one). Point size scales with the square root of total speeches. The horizontal axis gives the change in the share of speeches mentioning "Ukraine" (post–pre invasion); the vertical axis gives the change in the share of speeches classified as climate-relevant.

Figure C.7 – Mechanical crowd-out hypothesis

Table C.5 – Speech totals, Ukraine mentions, and climate speaking by party

	Speeches (N)		Share on Ukraine		Share on climate		Share climate-skeptic	
Party	Pre	Post	Pre	Post	Pre	Post	Pre	Post
The Greens	2,229	2,861	0.018	0.150	0.267	0.253	0.000	0.001
The Left	1,709	2,622	0.008	0.123	0.185	0.163	0.004	0.003
S-D	4,412	5,263	0.010	0.101	0.175	0.173	0.000	0.002
Renew	2,939	3,831	0.020	0.147	0.180	0.174	0.003	0.008
EPP	4,747	6,008	0.021	0.139	0.156	0.176	0.005	0.014
NI	1,225	1,511	0.002	0.051	0.126	0.175	0.016	0.024
ECR	2,217	2,729	0.025	0.141	0.149	0.221	0.050	0.078
ID	1,963	1,973	0.016	0.107	0.141	0.241	0.070	0.105

Note: For each of the eight EU party groups, the table reports the number of speeches, the share of speeches mentioning "Ukraine", the share of speeches classified as climate-relevant, and the share classified as climate-skeptic, separately for the pre-invasion period (before 24 February 2022) and the post-invasion period.

D. Roll-call data and analysis

Our model implies that narratives and policy are jointly determined in equilibrium. In the main text, we focused on the narrative aspect in order to investigate our key prediction of narrative entanglement. We now examine the link between narratives and policy empirically by linking our narrative patterns from the previous sections to legislative votes in the European Parliament, using roll-call data for each of the three policy domains we have investigated, i.e., climate, immigration, and digitization. We outline our roll-call data in Section D.1 and describe in Section D.2 how we classify them into the three policy domains we focus on: climate, immigration, and digitization. In Section D.3, we then show the relationship between these legislation measures and our narrative measures from Section 3 of the main text.

D.1. Roll-call data

Source. We download full roll-call data for the 9th European Parliament from Parltrack, an independent mirror of the European Parliament's legislative records.[35] This gives us, for each roll call, the date of the vote, a unique vote identifier, a short textual description of the item being voted on, and the individual vote choice (for, against, or abstain) of each MEP. We then enrich this with legislative metadata from the HowTheyVote project, which links each roll call to the dossier being voted on and to its subject codes in the European Parliament's *Legislative Observatory* (OEIL) taxonomy, which details the subject of the vote. This gives us, for every roll call, the set of hierarchical OEIL subject codes and a flag marking whether the roll call is a main vote (i.e., the final vote on a legislative text) rather than a procedural or amendment vote.[36]

Overview. The 9th European Parliament held 18,827 roll-call votes, cast by 867 MEPs. Of these, 9.5% are classified as *main* votes by HowTheyVote, which comprise final plenary votes on the text of a legislative act, amended Commission proposals, or stand-alone resolutions. The remaining votes are amendment-level ballots, procedural motions, rejection motions, and other intermediate votes. The median MEP cast 17,242 votes, of which 1,584 were main votes. Individual vote choices split 56.9% *for*, 37.7% *against*, and 5.4% *abstain*.

D.2. Roll-call classification of policy-relevant votes

To select the relevant roll calls for each of our policies we use the official OEIL subject-code taxonomy. OEIL is the European Parliament's registry of legislative dossiers: each

[35] Parltrack's data can be accessed at https://parltrack.org/dumps.

[36] HowTheyVote's data releases are available at https://github.com/HowTheyVote/data. OEIL is the EU parliament's authoritative registry of legislative dossiers; its subject-code hierarchy is maintained at https://oeil.secure.europarl.europa.eu/.

plenary roll call is linked to the dossier whose text is being voted on, and every dossier is tagged with one or more hierarchical OEIL subject codes.

OEIL codes are split into 8 main topic areas (e.g. "European citizenship", "Area of freedom, security and justice"), 44 sub-areas (e.g. 3.70 "Environmental policy"), 255 sub-sub-areas (e.g. 3.70.03 "Climate policy, climate change, ozone layer") and 64 additional, still more fine-grained areas (e.g. 3.10.01.02 "Rural development, European Agricultural Fund for Rural Development"). We identify environmental policy directly with sub-area 3.70 ("Environmental policy") and immigration policy with sub-area 7.10 ("Free movement and integration of third-country nationals"). Digitization does not have a sub-area of its own; instead, within the communications sub-area 3.30 we retain the four sub-sub-areas most directly concerned with digital technology: 3.30.05 ("Electronic and mobile communications"), 3.30.06 ("Information and communication technologies, digital technologies"), 3.30.07 ("Cybersecurity, cyberspace policy"), and 3.30.25 ("International information networks and society, internet").

Figure D.1 plots the 10-day moving average of the share of all roll-call votes that are relevant to each of the three policy domains across the 9th European Parliament, where "relevance" denotes that the vote's dossier carries at least one OEIL subject code in the domain's branch of the hierarchy. Panel (a) uses the full set of roll-call votes as the denominator; Panel (b) restricts to main votes only as identified by the HowTheyVote project. Across all votes, climate accounts on average for 10% of roll calls on a sitting day, followed by digitization at 3% and immigration at 2%. Restricting to main votes, the shares rise to 11%, 4%, and 3% respectively. Together, the three domains account for 15% of all roll calls and 18% of main votes per sitting day.

The time series has spikes corresponding to key legislations across these topics. For climate, our main application, these correspond to an early concentration in January–February 2020 coinciding with the launch of the European Green Deal, the European Climate Law at first reading in October 2020 and again at its post-trilogue final adoption in June 2021, a very pronounced peak in summer 2022 when the Fit-for-55 package (Emissions Trading System revision, Carbon Border Adjustment Mechanism, Social Climate Fund, ReFuelEU Aviation, and the Taxonomy Delegated Act on gas and nuclear) reached first reading, and two further clusters in autumn 2023 and early 2024 as the same package was finalized alongside the Critical Raw Materials Act, the Packaging Regulation, and the CO_2 standards for heavy-duty vehicles.

Table D.1 summarizes MEP×vote observations on all domain-relevant roll calls. There is strong variation in voting for, against, and abstaining across the policy topics.

Table D.2 shows the same content but restricted to the smaller sample of main votes: the

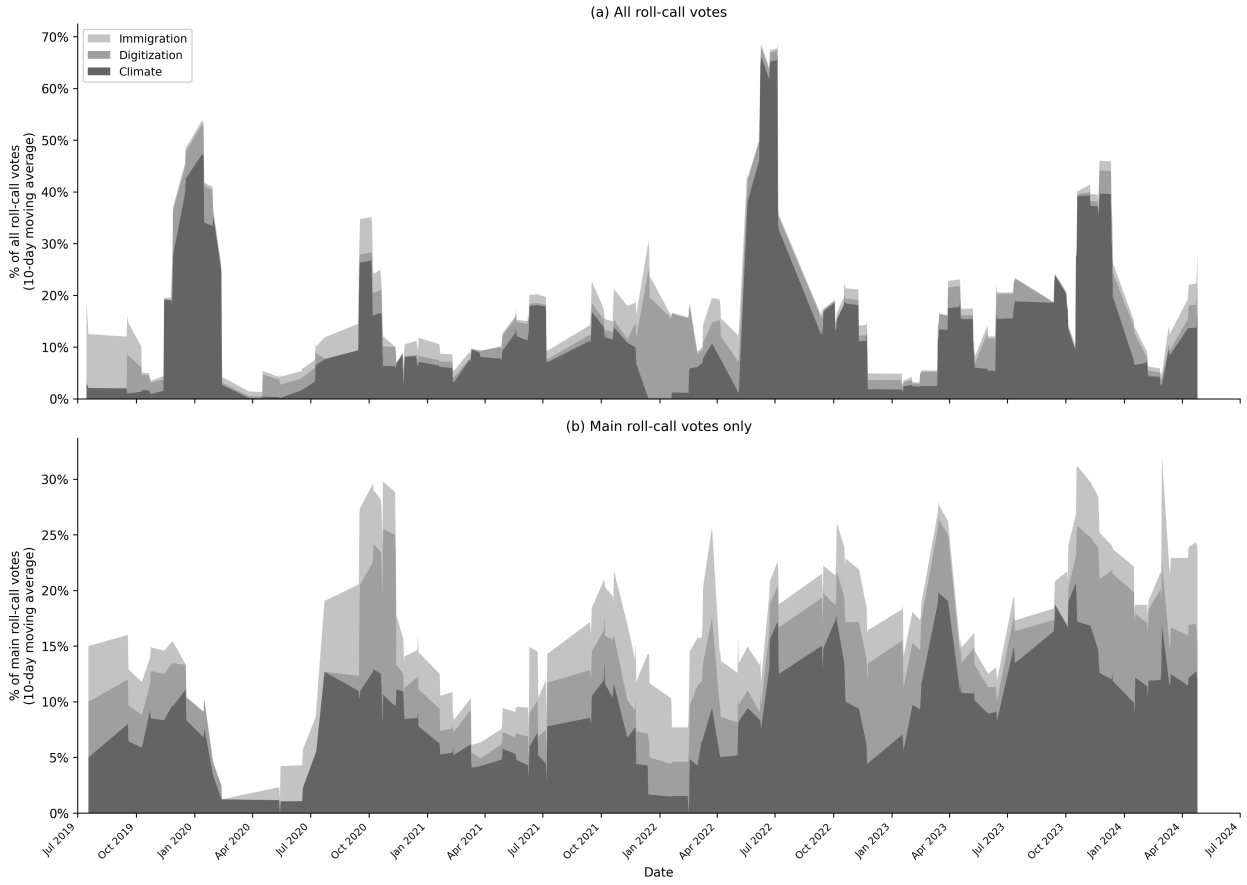

Note: 10-day moving average of the share of roll-call votes whose underlying dossier carries at least one OEIL subject code associated with climate (any code beginning with 3.70), digitization (codes 3.30.05, 3.30.06, 3.30.07, 3.30.25), or immigration (any code beginning with 7.10). Panel (a) uses all roll-call votes; Panel (b) restricts to main votes.

Figure D.1 – Climate, digitization, and immigration policy roll-call votes

final plenary vote on the text of a legislative act, an amended Commission proposal, or a stand-alone resolution, as flagged by HowTheyVote. The three samples reduce to 173 climate, 64 immigration, and 74 digitization main roll calls. The main-vote restriction discards amendment-level ballots and procedural motions, sharpening the substantive content of each vote at the cost of sample size.

Directionality. So far we have identified votes on the topics of climate, immigration, and digitization policy. However, we have not determined the directionality of the vote: whether each vote is a tightening or loosening of the underlying legislation.

Whether a given regulation should be classified as tightening or loosening policy is an inherently subjective judgment. Instead of trying to judge each of the relevant votes one by one, we benchmark votes against the modal vote of the party group that most directly

Table D.1 – Voting on all domain-relevant roll calls by party group

	Climate			Immigration			Digitization		
	For	Against	Abstain	For	Against	Abstain	For	Against	Abstain
The Greens	0.679	0.309	0.013	0.641	0.330	0.029	0.672	0.308	0.020
The Left	0.657	0.312	0.031	0.578	0.359	0.063	0.620	0.284	0.096
S&D	0.652	0.329	0.018	0.685	0.291	0.024	0.586	0.396	0.018
Renew	0.602	0.378	0.019	0.670	0.309	0.021	0.539	0.444	0.017
EPP	0.578	0.406	0.017	0.625	0.355	0.020	0.524	0.467	0.010
NI	0.524	0.352	0.123	0.562	0.323	0.115	0.558	0.305	0.137
ECR	0.427	0.486	0.087	0.527	0.369	0.104	0.498	0.398	0.104
ID	0.384	0.500	0.116	0.477	0.376	0.148	0.526	0.307	0.167
Obs.		1,615,250			215,031			387,125	

Note: Share of MEP×vote observations casting for, against, or abstain on all roll calls whose dossier carries at least one OEIL subject code relevant to the domain: 2,503 climate roll calls (OEIL 3.70 codes); 333 immigration roll calls (7.10 codes); 591 digitization roll calls (3.30.05, 3.30.06, 3.30.07, 3.30.25).

Table D.2 – Voting on main roll calls by party group, across the three policy domains

	Climate			Immigration			Digitization		
	For	Against	Abstain	For	Against	Abstain	For	Against	Abstain
The Greens	0.858	0.106	0.036	0.648	0.258	0.094	0.811	0.136	0.054
The Left	0.717	0.167	0.116	0.529	0.326	0.145	0.532	0.235	0.234
S&D	0.925	0.058	0.017	0.893	0.067	0.040	0.966	0.025	0.009
Renew	0.894	0.070	0.037	0.919	0.056	0.025	0.951	0.028	0.021
EPP	0.838	0.106	0.057	0.824	0.134	0.043	0.951	0.031	0.017
NI	0.619	0.246	0.135	0.548	0.354	0.098	0.688	0.210	0.101
ECR	0.381	0.414	0.205	0.447	0.438	0.116	0.644	0.161	0.195
ID	0.334	0.475	0.191	0.289	0.581	0.130	0.516	0.232	0.252
Obs.		108,530			39,601			46,136	

Note: Main-roll-call counterpart of Table D.1.

represents a tightening narrative on the policy topic. For climate policy, we choose the modal vote of The Greens; for immigration policy, of ID; and for digitization policy, of The Left. In each case, we drop votes where the benchmark party abstained, and hence only look at agreement with the benchmark party if the latter had an overall *for* or *against* modal vote. Within each benchmarked vote, an MEP is coded as agreeing if her vote coincides with the benchmark's modal choice and as disagreeing if her vote is the opposite of the benchmark's modal choice; MEP abstentions are thus treated as missing in the headline disagree-share.[37]

[37] An alternative coding that counts MEP abstentions as disagreement leaves all conclusions unchanged.

While our choice is relatively clear-cut for climate and immigration policy (also confirmed by the narrative evidence in Section 3.3), the choice of the benchmark party for digitization policy is less obvious. We use The Left as our benchmark for two reasons. First, The Left is among the parties with the largest proportion of pro-tightening narratives in digitization, alongside The Greens and S&D, as seen in the narrative evidence in Section 3.3. Second, among these pro-tightening parties, The Left is the only group whose modal vote varies meaningfully across the digitization main votes (see Appendix Table D.2).

Table D.3 reports voting shares across parties that agree or disagree with the benchmark party on each of the policy dimensions (Table D.4 reports the same for main roll calls only). The drop in observations relative to Table D.1 is due to the dropping of votes where the benchmark party abstained: 1.0% of MEP×vote observations for climate, 12.3% for immigration, and 6.7% for digitization. The agreement rates on climate and immigration are near mirror images of one another, confirming the left-right cleavage on these issues. Digitization policy shows a less clear-cut left-right split. Note that agreement rates of benchmark parties with themselves are not perfect, as not all MEPs vote with the mode of their party. Much of this variation happens at the national party level (see Table D.5).

Table D.3 – Agreement with the benchmark party's modal vote, by party group

	Climate (vs The Greens)			Immigration (vs ID)			Digitization (vs The Left)		
	Agree	Disagree	Abstain	Agree	Disagree	Abstain	Agree	Disagree	Abstain
The Greens	0.990	0.005	0.004	0.234	0.733	0.032	0.780	0.201	0.020
The Left	0.901	0.070	0.030	0.243	0.694	0.063	0.907	0.044	0.049
S&D	0.841	0.141	0.017	0.281	0.692	0.027	0.638	0.343	0.019
Renew	0.719	0.262	0.020	0.307	0.670	0.023	0.551	0.431	0.017
EPP	0.495	0.489	0.017	0.387	0.595	0.018	0.511	0.481	0.008
NI	0.544	0.333	0.123	0.479	0.407	0.113	0.526	0.337	0.137
ECR	0.229	0.684	0.086	0.724	0.180	0.096	0.401	0.493	0.106
ID	0.213	0.671	0.115	0.868	0.057	0.075	0.392	0.448	0.160
Obs.		1,599,240			188,514			361,169	

Note: Share of MEP×vote observations that agree with, disagree with, or abstain relative to the benchmark party's modal vote on each roll call, computed on all domain-relevant roll calls. Benchmark parties: the Greens for climate, ID for immigration, The Left for digitization. Roll calls on which the benchmark's modal choice is abstain or missing are dropped.

Table D.4 – Agreement with the benchmark party's modal vote, main roll calls only

	Climate (vs The Greens)			Immigration (vs ID)			Digitization (vs The Left)		
	Agree	Disagree	Abstain	Agree	Disagree	Abstain	Agree	Disagree	Abstain
The Greens	0.985	0.007	0.009	0.258	0.644	0.099	0.765	0.171	0.064
The Left	0.824	0.070	0.106	0.280	0.568	0.152	0.836	0.056	0.109
S&D	0.921	0.063	0.016	0.301	0.657	0.042	0.694	0.297	0.010
Renew	0.845	0.119	0.036	0.337	0.637	0.025	0.669	0.307	0.024
EPP	0.767	0.176	0.058	0.416	0.557	0.027	0.629	0.354	0.017
NI	0.627	0.238	0.134	0.529	0.377	0.094	0.595	0.312	0.093
ECR	0.358	0.441	0.202	0.695	0.192	0.114	0.466	0.322	0.212
ID	0.305	0.506	0.189	0.858	0.051	0.090	0.431	0.356	0.212
Obs.		105,469			37,590			37,343	

Note: Main-roll-call counterpart of Table D.3.

D.3. Relationship between narratives and legislation

Variance decomposition. We start by comparing how much MEPs within a national party differ in their narratives and votes, given that party discipline may be stronger on votes than on speeches. For each domain we construct two MEP-level shares: the share of the MEP's domain-relevant speeches coded as a tighter narrative, and the share of her domain roll calls cast with the benchmark group's modal choice discussed above.

Table D.5 reports the variance of these MEP-level shares across the full sample of MEPs, the residual variance after demeaning by national party, and the share of total variance absorbed by national-party identity.

Table D.5 – Within-national-party variance: narratives vs votes

	Climate		Immigration		Digitization	
	Narrative (tighter)	Vote (agree w/ Greens)	Narrative (tighter)	Vote (agree w/ ID)	Narrative (tighter)	Vote (agree w/ The Left)
Variance of MEP means	0.1192	0.0738	0.1236	0.0622	0.1151	0.0258
Within-party variance of MEP means	0.0346	0.0013	0.0388	0.0028	0.0679	0.0021
% of variance explained by party	71.0%	98.3%	68.6%	95.5%	41.1%	92.0%
Observations	8,843	1,538,646	2,175	180,605	4,438	344,740
MEPs	687	845	451	840	642	844
National parties	253	284	208	283	244	284

Note: Each cell is an MEP-level share. *Narrative* is the share of the MEP's domain-relevant speeches coded "tighter" on the non-economic dimension; *Vote* is her share of OEIL-matched roll calls in the domain cast with the benchmark group's modal choice (Greens for climate, ID for immigration, The Left for digitization; benchmark-abstain or missing-mode votes dropped). The last row is 1− (within-party variance/ total variance).

National-party identity absorbs 92–98% of the variance in MEP-level vote shares across

the three domains, but only 41–71% of the variance in narrative shares.

Narrative–vote alignment. We now investigate the alignment between narratives and votes. Because roll-call variation is meaningful only between national parties as we have just shown, we conduct the main narrative–vote alignment exercise at the national-party level, collapsing to one observation per national party on both sides. This prevents us from overstating the number of observations that actually generate variation. For completeness, we also show the same analysis at MEP level.

Figure D.2 plots, at the national-party level, each party's share of roll-call votes cast opposite the domain's benchmark group against two narrative shares: the loose-economic narrative (top row) and the loose-non-economic narrative (bottom row), with one column per policy domain. For climate and immigration the slopes are positive and relatively large, while for digitization the relationship is flatter. All six slopes are statistically significant. Figure D.4 shows results when partialing out country fixed effects, and Figure D.3 at the MEP level, both of which have qualitatively similar results.

MEP-level narrative–vote delta scatter. Figure D.5 reports the MEP-level counterpart to Figure 6 of the main text. The construction mirrors the national-party-level design, with the unit of observation replaced by the individual MEP. Both fitted relationships are positive, though flatter than in the national-party-level figure.

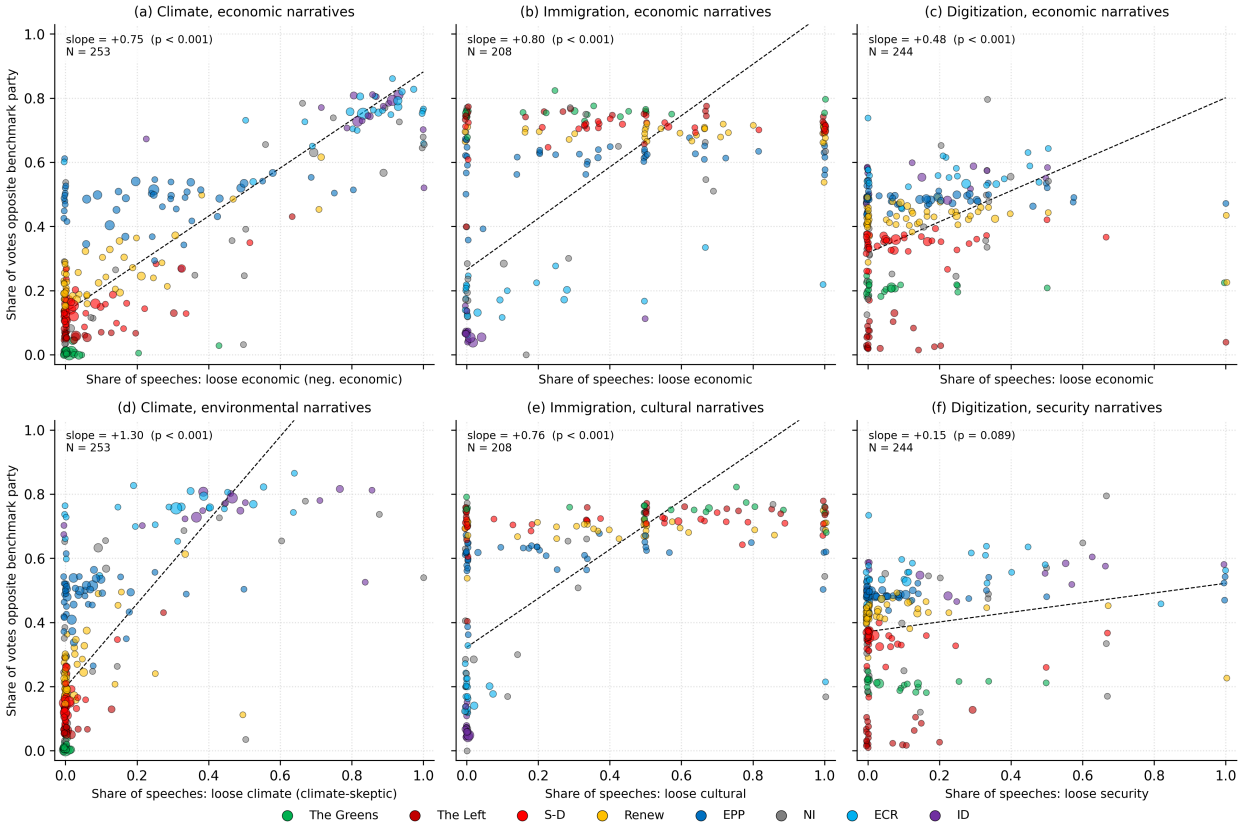

Note: Each point is a national party; point size scales with the square root of number of speeches. The y-axis in every panel is the party's share of votes cast opposite the domain's benchmark group (Greens for climate, ID for immigration, The Left for digitization; benchmark-abstain votes dropped). Columns: climate, immigration, digitization. Top row: x-axis is the party's share of domain-relevant speeches with a loose or no-effect economic narrative; bottom row: x-axis is the share with a loose or no-effect non-economic narrative (climate change for climate policy, culture for immigration, and security for digitization). Red lines are speech-weighted regression lines, and standard errors used to compute p-values are clustered by country.

Figure D.2 – National-party-level narrative–vote alignment across three policy domains

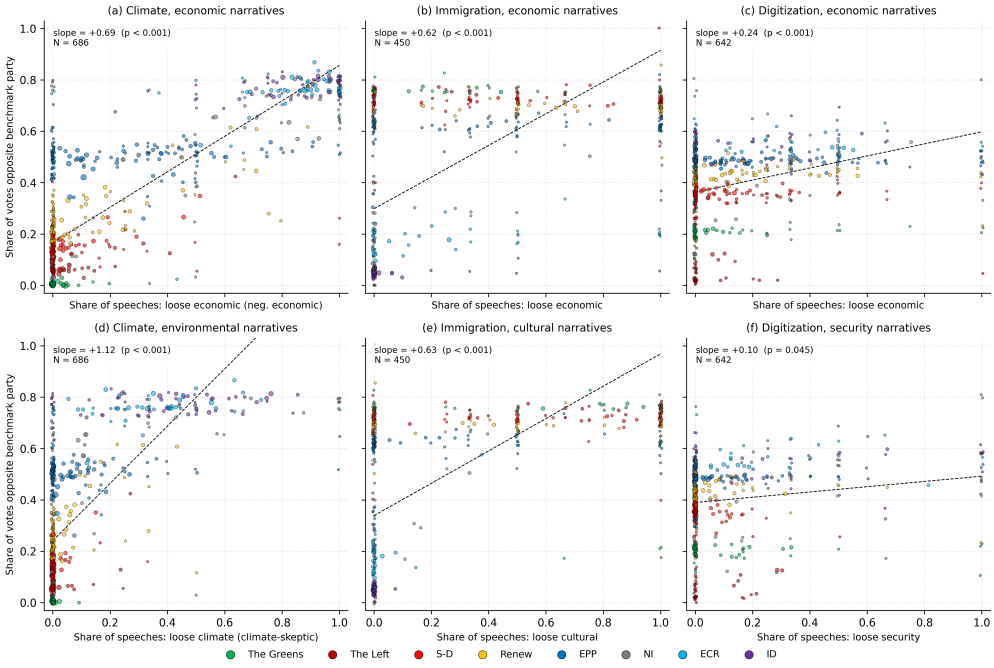

Note: MEP-level counterpart of Figure D.2, with standard errors clustered at country times party group.

Figure D.3 – MEP-level narrative–vote alignment across three policy domains

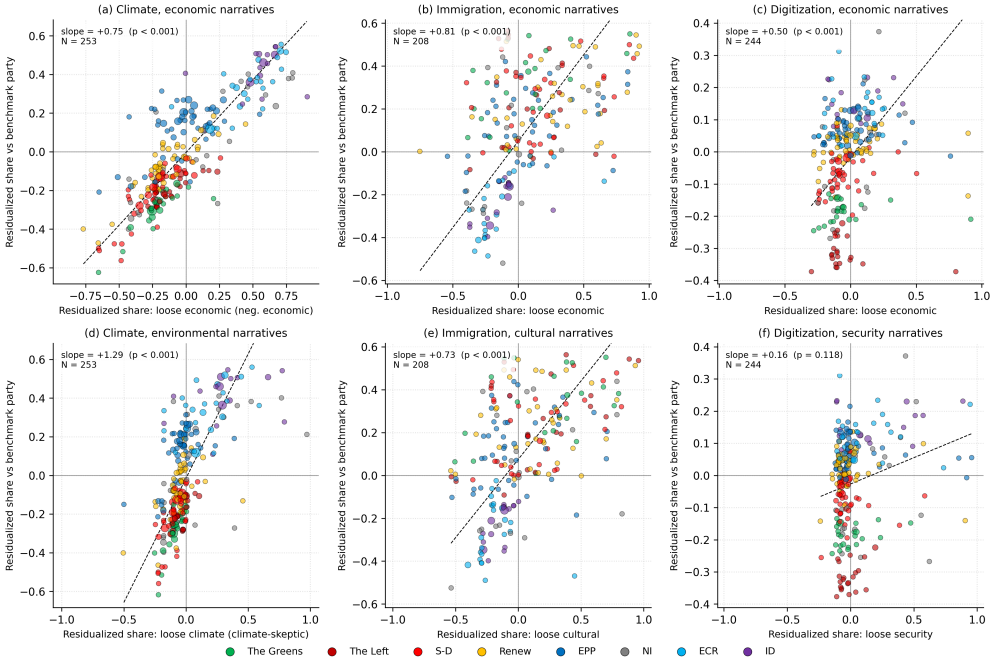

Note: Counterpart of Figure D.2 with country fixed-effects.

Figure D.4 – National-party-level narrative–vote alignment, country fixed effects

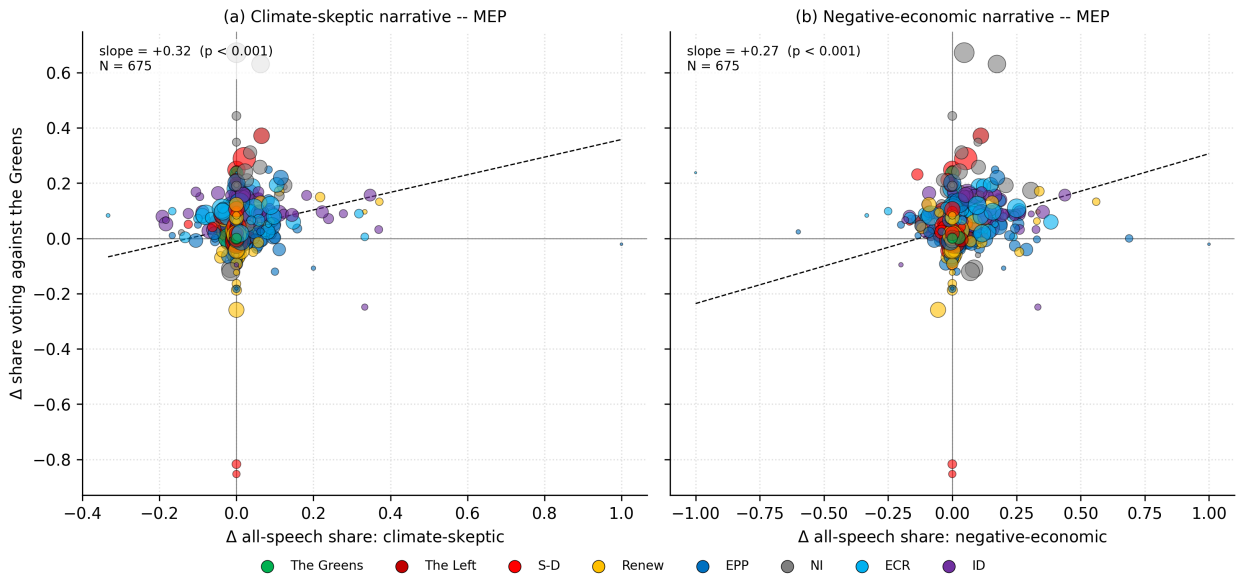

Note: MEP-level counterpart of Figure 6, with standard errors clustered at country times party group.

Figure D.5 – Shifts in narratives versus legislative voting, MEP level

E. Exact classification prompts

We classify each speech six times, once for each ordered policy–narrative pairing: climate policy on climate change, climate policy on the economy, immigration policy on culture, immigration policy on the economy, digitization policy on security, and digitization policy on the economy. All six requests use the same LLM (`gpt-5-mini-2025-08-07`), the same response format (`json_object`), and the same system message:

> You are a political economist specializing in classification of narratives made in political speeches. You will analyze a speech made in the European Parliament in one of the EU's official languages. You will output ONLY valid JSON in English, with no extra text or commentary.

At the end of the prompt, we also briefly specify the exact JSON-output format.

The user message is then constructed from one of the six templates reproduced below, followed by the literal speech text.

E.1. Climate policy and climate change

> Classify the speech below according to its narrative about how climate policy affects climate change, and provide a 1-2 sentence justification for your classification.
>
> **Classification categories:**
>
> - **"tighter"**: This speech implies that tighter climate policy would mitigate climate change.
> - **"no effect"**: This speech implies that climate policy has no effect on climate change.
> - **"looser"**: This speech implies that looser climate policy would mitigate climate change.
> - **"unrelated"**: This speech is not related to the effect of climate policy on climate change.
>
> **Classification rules:**
>
> - Speeches whose primary topic is not climate policy should be classified as "unrelated", even if they contain one or two sentences about climate policy.
> - Speeches that doubt climate change is an issue, or that use sarcasm and irony to dismiss climate action, should be classified as "no effect".
> - Exclusively focus on the link between climate policy and climate change, not on the link between climate policy and the economy or on the link between climate policy and other outcomes. If the speech only discusses the effect of climate

policy on the economy or other outcomes, but not on climate change, it should be classified as "unrelated".

- Capture the speech's overall desired direction for climate policy based on its effect on climate change, not its stance on any single measure. For example, a speech can criticise a particular policy instrument while overall advocating for tighter climate action, in which case it should be classified as "tighter". Or, a speech can endorse a specific green initiative while overall arguing that climate policy should be less ambitious, in which case it should be classified as "looser".

E.2. Climate policy and the economy

Classify the speech below according to its narrative about how climate policy affects the economy, and provide a 1-2 sentence justification for your classification.

Classification categories:

- **"tighter"**: This speech implies that tighter climate policy would have a positive effect on the economy.
- **"no effect"**: This speech implies that climate policy has no effect on the economy.
- **"looser"**: This speech implies that looser climate policy would have a positive effect on the economy.
- **"unrelated"**: This speech is not related to the effect of climate policy on the economy.

Classification rules:

- Speeches whose primary topic is not climate policy should be classified as "unrelated", even if they contain one or two sentences about climate policy.
- Exclusively focus on the link between climate policy and the economy, not on the link between climate policy and climate change or on the link between climate policy and other outcomes. If the speech only discusses the effect of climate policy on climate change or other outcomes, but not on the economy, it should be classified as "unrelated".
- Capture the speech's overall desired direction for climate policy based on its effect on the economy, not its stance on any single measure. For example, a speech can criticise a particular policy instrument while overall arguing that tighter climate policy benefits the economy, in which case it should be classified as "tighter". Or, a speech can endorse a specific green initiative while overall arguing that climate policy is too costly for the economy, in which case it should be classified as "looser".

E.3. Immigration policy and culture

Classify the speech below according to its narrative about how immigration policy affects culture, and provide a 1-2 sentence justification for your classification.

Classification categories:

- **"tighter"**: This speech implies that tighter immigration policy would have a positive effect on culture.
- **"no effect"**: This speech implies that immigration policy has no effect on culture.
- **"looser"**: This speech implies that looser immigration policy would have a positive effect on culture.
- **"unrelated"**: This speech is not related to the effect of immigration policy on culture.

Classification rules:

- Speeches whose primary topic is not immigration policy should be classified as "unrelated", even if they contain one or two sentences about immigration policy.
- Exclusively focus on the link between immigration policy and culture, not on the link between immigration policy and the economy or on the link between immigration policy and other outcomes. If the speech only discusses the effect of immigration policy on the economy or other outcomes, but not on culture, it should be classified as "unrelated".
- Capture the speech's overall desired direction for immigration policy based on its effect on culture, not its stance on any single measure. For example, a speech can criticise a particular tightening measure while overall arguing that tighter immigration policy would benefit culture, in which case it should be classified as "tighter". Or, a speech can endorse a specific tightening measure while overall arguing that looser immigration policy would benefit culture, in which case it should be classified as "looser".

E.4. Immigration policy and the economy

Classify the speech below according to its narrative about how immigration policy affects the economy, and provide a 1-2 sentence justification for your classification.

Classification categories:

- **"tighter"**: This speech implies that tighter immigration policy would have a positive effect on the economy.
- **"no effect"**: This speech implies that immigration policy has no effect on the economy.

- **"looser"**: This speech implies that looser immigration policy would have a positive effect on the economy.
- **"unrelated"**: This speech is not related to the effect of immigration policy on the economy.

Classification rules:

- Speeches whose primary topic is not immigration policy should be classified as "unrelated", even if they contain one or two sentences about immigration policy.
- Exclusively focus on the link between immigration policy and the economy, not on the link between immigration policy and culture or on the link between immigration policy and other outcomes. If the speech only discusses the effect of immigration policy on culture or other outcomes, but not on the economy, it should be classified as "unrelated".
- Capture the speech's overall desired direction for immigration policy based on its effect on the economy, not its stance on any single measure. For example, a speech can criticise a particular tightening measure while overall arguing that tighter immigration policy benefits the economy, in which case it should be classified as "tighter". Or, a speech can endorse a specific tightening measure while overall arguing that loosening immigration policy would benefit the economy, in which case it should be classified as "looser".

E.5. Digitization policy and security

Classify the speech below according to its narrative about how digitization policy affects security, and provide a 1-2 sentence justification for your classification.

Classification categories:

- **"tighter"**: This speech implies that tighter digitization policy would have a positive effect on security.
- **"no effect"**: This speech implies that digitization policy has no effect on security.
- **"looser"**: This speech implies that looser digitization policy would have a positive effect on security.
- **"unrelated"**: This speech is not related to the effect of digitization policy on security.

Classification rules:

- Speeches whose primary topic is not digitization policy should be classified as "unrelated", even if they contain one or two sentences about digitization policy.
- Exclusively focus on the link between digitization policy and security, not on the link between digitization policy and the economy or on the link between

digitization policy and other outcomes. If the speech only discusses the effect of digitization policy on the economy or other outcomes, but not on security, it should be classified as "unrelated".

- Capture the speech's overall desired direction for digitization policy based on its effect on security, not its stance on any single measure. For example, a speech can criticise a particular tightening measure while overall arguing that tighter digitization policy would benefit security, in which case it should be classified as "tighter". Or, a speech can endorse a specific tightening measure while overall arguing that looser digitization policy would benefit security, in which case it should be classified as "looser".

E.6. Digitization policy and the economy

Classify the speech below according to its narrative about how digitization policy affects the economy, and provide a 1-2 sentence justification for your classification.

Classification categories:

- **"tighter"**: This speech implies that tighter digitization policy would have a positive effect on the economy.
- **"no effect"**: This speech implies that digitization policy has no effect on the economy.
- **"looser"**: This speech implies that looser digitization policy would have a positive effect on the economy.
- **"unrelated"**: This speech is not related to the effect of digitization policy on the economy.

Classification rules:

- Speeches whose primary topic is not digitization policy should be classified as "unrelated", even if they contain one or two sentences about digitization policy.
- Exclusively focus on the link between digitization policy and the economy, not on the link between digitization policy and security or on the link between digitization policy and other outcomes. If the speech only discusses the effect of digitization policy on security or other outcomes, but not on the economy, it should be classified as "unrelated".
- Capture the speech's overall desired direction for digitization policy based on its effect on the economy, not its stance on any single measure. For example, a speech can criticise a particular tightening measure while overall arguing that tighter digitization policy benefits the economy, in which case it should be classified as "tighter". Or, a speech can endorse a specific tightening measure while overall

arguing that loosening digitization policy would benefit the economy, in which case it should be classified as "looser".

F. Cross-validation and robustness of measurement

For each robustness exercise we draw a fixed random sample of 1,000 speeches (seed fixed for replication) and re-classify the six policy–narrative pairings. Let the *baseline* refer to the GPT-5-mini classification underlying all results in the main text. Agreement rates are reported both unconditionally and conditional on both the baseline and the alternative classifying the speech as relevant (i.e., voicing a narrative in at least one dimension of the policy domain). We also report Cohen's κ, which adjusts raw agreement for the high share of speeches (78–96% across domains) that are "unrelated" in both classifications.

F.1. Comparison to other models

Table F.1 pools the 4×4 confusion matrices across all six policy–narrative pairings. Panel (a) compares the baseline to the full GPT-5 model. The diagonal accounts for 97.2% of cells and Cohen's κ equals 0.78, indicating substantial chance-corrected agreement. The largest off-diagonal mass — 131 cells out of 6,000 — consists of speeches the baseline labels as relevant but GPT-5 labels as unrelated; the reverse direction (baseline-unrelated but GPT-5-relevant) accounts for only 18 cells. GPT-5 is therefore more conservative than the baseline on topic relevance, but there is no evidence of differential calibration across narratives. Conditional on both classifiers calling the speech relevant, their narrative labels agree on 88–96% of cells across the six pairings, and sign reversals ("tighter" versus "looser") occur in only 16 of 6,000 cells (0.27%).

Panel (b) compares the baseline to GPT-5.4-mini, OpenAI's April 2026 update of the mini family. The diagonal accounts for 95.4% of cells and Cohen's κ equals 0.72, again substantial chance-corrected agreement. Unlike the full GPT-5 comparison, the off-diagonal mass is asymmetric in the opposite direction: 195 cells are flagged as relevant by GPT-5.4-mini but as unrelated by the baseline, against only 48 in the reverse direction. GPT-5.4-mini is therefore slightly more likely to classify a speech as on topic, but once relevance is agreed, narrative labels line up: conditional on both classifiers calling the speech relevant, agreement ranges from 83% to 100% across the six pairings, with sign reversals in 22 of 6,000 cells (0.37%).

F.2. Robustness to other prompts

Table F.2 repeats the exercise with the baseline model (GPT-5-mini) but under three alternative prompts applied to the same 1,000-speech sample.

Table F.1 – Robustness to alternative models: confusion matrices against the GPT-5-mini baseline

(a) Baseline (GPT-5-mini) versus full GPT-5

Robustness classifier	**tighter**	**no effect**	**looser**	**unrelated**	**Total**
tighter	237	3	13	99	352
no effect	0	5	1	3	9
looser	3	2	69	29	103
unrelated	10	3	5	5,518	5,536
Total	**250**	**13**	**88**	**5,649**	**6,000**

(b) Baseline (GPT-5-mini) versus GPT-5.4-mini

Robustness classifier	**tighter**	**no effect**	**looser**	**unrelated**	**Total**
tighter	301	3	10	38	352
no effect	1	2	6	0	9
looser	12	1	80	10	103
unrelated	127	3	65	5,341	5,536
Total	**441**	**9**	**161**	**5,389**	**6,000**

Note: Each panel reports the pooled 4×4 confusion matrix across the six policy–narrative pairings on the 1,000-speech random sample. Rows are the baseline classification, columns are the alternative model.

Panel (a) varies the structure of the prompt. We replace the one-step classification with a two-step prompt that first asks the classifier to decide whether the speech is relevant to the policy–narrative pairing, and only conditional on relevance to assign one of three narrative stances. Overall agreement with the baseline is 97.3% and Cohen's κ is 0.82 on the pooled 6,000 cells; narrative agreement conditional on both classifications calling the speech relevant ranges from 92% to 100% across the six pairings, with sign reversals in 16 of 6,000 cells (0.27%).

Panel (b) varies the vocabulary of the narrative categories. We replace "tighter" and "looser" with "more restrictive" and "less restrictive", holding the rules, the examples, and the output format otherwise fixed. Overall agreement is 97.4% and Cohen's κ is 0.82, essentially identical to Panel (a). Narrative agreement conditional on both classifications calling the speech relevant ranges from 92% to 100%, with 33 sign reversals (0.55%). The classification is therefore not anchored on the particular English vocabulary used for the directional labels.

Panel (c) varies the content of the prompt. We strip the explicit *classification rules* from

the user prompt, retaining only the header, the four-category definition, and the output format. Overall agreement with the baseline is 95.5% and Cohen's κ is 0.75. The disagreement is overwhelmingly one-sided: 239 baseline-unrelated cells are re-classified as relevant, against only 16 in the reverse direction. Conditional on both classifications calling the speech relevant, narrative agreement is 95.7–97.6% across the six pairings, with 13 sign reversals (0.22%). The rules are therefore doing exclusion work — filtering peripheral mentions of a policy that are not the speech's primary topic — but they do not strongly influence the narrative labels. Within the relevant set, the four-category structure is doing the classification on its own.

Table F.2 – Robustness to alternative prompts: confusion matrices against the GPT-5-mini baseline

(a) Baseline versus two-step (relevance then narrative) prompt

Robustness classifier	**tighter**	**no effect**	**looser**	**unrelated**	**Total**
tighter	313	0	5	34	352
no effect	0	8	1	0	9
looser	11	1	84	7	103
unrelated	67	2	35	5,432	5,536
Total	**391**	**11**	**125**	**5,473**	**6,000**

(b) Baseline versus "more/less restrictive" vocabulary prompt

Robustness classifier	**tighter**	**no effect**	**looser**	**unrelated**	**Total**
tighter	273	3	27	49	352
no effect	0	5	2	2	9
looser	6	0	87	10	103
unrelated	28	1	28	5,479	5,536
Total	**307**	**9**	**144**	**5,540**	**6,000**

(c) Baseline versus rules-stripped prompt

Robustness classifier	**tighter**	**no effect**	**looser**	**unrelated**	**Total**
tighter	341	0	1	10	352
no effect	0	8	1	0	9
looser	12	1	84	6	103
unrelated	179	4	56	5,297	5,536
Total	**532**	**13**	**142**	**5,313**	**6,000**

Note: Each panel reports the pooled 4×4 confusion matrix across the six policy–narrative pairings on the 1,000-speech random sample. Rows are the baseline classification, columns are based on alternative prompt with the same model.