Political Cycles over Worldviews: Increments and Backlashes in Complexity

Philippe Jehiel[*] Paul-Henri Moisson[**]

January 9, 2026

Abstract

We develop a dynamic model of political competition in which worldviews and corresponding policies are proposed by politicians, voters differ in their ability to understand worldviews, and vote for the proposed worldview that they can understand and that explains best the observed data. While multiple ergodic distributions can arise, in all of them the complexity of the winning worldview follows a deterministic cycle, composed of steady increments towards higher complexity, interrupted by a backlash towards the simplest worldview which consistently picks the short-termist policy. We consider the implications of our model in terms of voting turnout and welfare.

[*]Paris School of Economics and University College London, jehiel@enpc.fr

[**]Paris School of Economics and Ecole des Ponts ParisTech, paul-henri.moisson@psemail.eu

1 Introduction

Political cycles in which incumbent politicians are repeatedly defeated by some challenger from the opposition can have different motives. As noted long ago by Condorcet (1785), cycles may arise for well specified heterogeneities in voters' preferences over potential candidates (assumed to be more than two). Other motives for political cycles include the possibility that the relative preference for the incumbent politician fades over time, due to either learning with a different model than the one used by the incumbent, or to memory imperfections, making the alternative policies proposed by the challenger look comparatively better over time (see Levy, Razin, and Young (2022) and Levy and Razin (2025) for recent formalizations of the former and the latter, respectively).

We propose a different rationale. We adopt an epistemic perspective: In our model, all voters share the same objective of choosing the best policies. However, while they know the immediate cost attached to the various policies a in the various states ω, they do not a priori know how big the prospective benefit of these are. They must rely on the worldviews disclosed by politicians, which express statistical links between state-policy pairs (ω, a) and the distribution of outcomes (or benefit) y (0-1 in our model). We formalize worldviews as simplified explanations of how policies affect outcomes in the various states, which take the form of partitions of state-policy pairs (ω, a) into "analogy classes", with for each class α, an associated outcome distribution $\beta(\alpha)$. The worldview is interpreted as a theory postulating that for all state-policy pair (ω, a) in the category α, the benefit y is distributed according to $\beta(\alpha)$.[1]

In line with the epistemic perspective, we assume that voters assess the proposed worldviews according to how good they are at explaining the observed data, which take the form of strings of triplets (ω, a, y) describing the past realized benefit/outcome y after policy a was chosen in state ω. While we assume that all past data are available (no memory imperfections), we crucially assume that voters differ in their ability to assess worldviews, which we parameterize by how many categories they can encompass. This can be viewed as a complexity bound we impose on voters, and we assume that voters are heterogeneous in this dimension.[2] More precisely, a voter with sophistication k

[1] Our formulation of worldviews follows that of analogy partitions (as in Jehiel (2005) or Mailath and Samuelson (2020)), but in some special cases of the worldview they can be viewed as expressing causal relations as in the work of Spiegler (2016) (see Jehiel (2022) for a discussion of the link between these various approaches).

[2] Complexity limitations have been documented in a number of experimental works, see in particular Oprea (2020, 2024).

can check the likelihood/plausibility of the observed realized data for a given worldview when this worldview consists of k categories at most. In a given period, voters choose among the incumbent and a challenger politician and vote for the one who proposes the most plausible worldview among those they understand. When no worldview can be assessed (because they are both too complex), the voter abstains. This in turn determines the current-period political competitive framework where we assume that the incumbent cannot change his worldview compared to the one he last proposed and we let the entrant freely choose his worldview so as to maximize his vote share. This results in a political outcome for the current term with the winner implementing the policy he has proposed after the realized state is observed. The corresponding data is added to the database after the benefit outcome for that pair of state-action is observed. We consider the ergodic distributions over worldviews and the corresponding policies that arise from such dynamics.

We observe that there always exists at least one ergodic distribution. In all ergodic distributions (there may be several), the incumbent is always defeated by the entrant and the complexity of the winning worldview follows a deterministic cycle, composed of steady increments towards higher complexity, interrupted by a backlash towards the simplest worldview. Participation in the election decreases as more complex worldviews are proposed. When the simplest worldview wins, the policy is always the short-termist policy in which the cost-minimizing policy is chosen in all states irrespective of the outcome distributions. As more complex worldviews win, other policies can be implemented depending on how the outcome distribution varies with the policy and the state. When the outcome distribution varies more with the state than the policy (which may represent the fate of small countries or global issues), we observe that the short-termist policy prevails for a potentially long phase of the political cycle. This may not be so when the outcome distribution varies more with the policy than the state (which may apply more naturally to large countries or domestic issues). But, when the space of policies becomes dense (with policies varying in the continuum in the limit), the policies adopted in the political cycle converges to the short-termist policy at all phases of the political cycle. We also observe that the efficiency of the policy may not vary monotonically with the complexity of the winning worldview and that a shift of distribution of voters toward more sophisticated ones need not improve the overall efficiency of the political cycle.

The intuition for our results is as follows. Given the cumulated data, the worldview with k categories that best explains the data is the one that minimizes the Kullback-Leibler divergence between the worldview with k categories and the observed data (this is similar to the entropy version of the global clustering considered in Jehiel and Weber (2024) and it can be viewed as being similar in spirit to Schwartzstein and Sunderam (2021) even if the idea of restricting attention to theories with k categories is not present there). Now, suppose that at some stage of the cycle, the incumbent relies on a worldview with complexity k. The entrant faces the following trade-off. By offering the (best) worldview with higher complexity say $k' > k$, he can explain the data better, but such a worldview will only be understood by those voters with sophistication at least k'. It follows that the best the entrant can do with such a strategy is to pick the (best) worldview with complexity $k+1$ as this would allow to explain the data better than the incumbent and get the maximum share of voters to understand it. This explains the incremental increase of complexity of the winning worldview in our cycle. Clearly, starting from the simplest worldview, this strategy of the entrant is optimal and allows him to win with a large margin. When the complexity of the incumbent's worldview becomes sufficiently large, the entrant can consider another strategy. Instead of trying to explain the data better, he can seek to get the support of the voters who do not understand the worldview of the incumbent (because it is too complex). The best such alternative strategy is to propose the simplest one-category worldview as this one is understood by everyone. This explains the backlash toward the simplest (best) worldview in the political cycle.

Regarding the policy implication, when the simplest worldview wins, it means that the distribution of outcomes is not distinguished according to the chosen policy, which immediately translates into the implementation of the short-termist policy. When outcomes are more driven by the state than the policy, the optimal clustering for a number of categories that is no more than the number of states will pool together (ω, a) when a varies, and this will also lead to the choice of the short-termist policy. When this is not the case, other policies can arise sooner in the cycle, and we note that they can sometimes be less efficient than the short-termist policy.[3] When there are many nearby policies, we also note that policies close to the short-termist policy arise because in each category,

[3] To see this most simply, suppose that policies are binary $a = 0, 1$ in all states ω with always $a = 0$ yielding the same (smallest) probability of success (say 0). Then when the winning worldview uses two states, all $a = 1$ from all the states will be assigned to one category, which will lead to a policy that chooses $a = 1$ in those states ω which have a lower cost $c(\omega, a = 1)$. When efficiency would require adopting $a = 1$ in those states where the cost is higher, the resulting policy may be less efficient than the short-termist policy.

only the cost-minimizing policy can be chosen, and with nearby policies this induces an unravelling force toward the short-termist policy.

While our model is stylized, we believe it captures several features observed in actual elections. In a number of countries, phases of rising technocratization and decreasing participation have been interrupted by populist backlashes, associated with simple worldviews and surges in political participation, which is consistent with our cyclical dynamics. For instance, a phase of initial (albeit mild) action against climate change was interrupted by the US withdrawals (2017, 2025) from the 2016 Paris Climate Change Agreement, motivated by the simplest possible worldview ("climate change is a hoax").[4] Such simple worldviews have led to short-termist policies, taking the form of, e.g., inaction (if not reversal of prior action) against climate change (IPCC (2022)), or sharp increases in tariffs or in rent controls (ignoring their longer-term, general-equilibrium consequences). It also captures another feature that we believe is plausible. When politicians use too technocratic a language (which we identify in our setting with politicians using more complex worldviews), many voters disregard those (because they are too hard to assess), and when the whole political offer relies on those, it leads to lower involvement (participation) of the voters. To some extent, the growing distrust toward the EU institutions, and the decreasing support for the EU mainstream parties can be related to this prediction.[5]

Of course, the exact deterministic cycles that arise in our setting may not fit reality perfectly, and we note how adding frictions in the form of noise in the observed dataset and/or noise in the extrinsic preference (valence) for politicians could lead to more stochasticity in the political cycles. We also extend our analysis in several directions. We show that the core dynamics persist when the state follows Markovian transitions – the distribution of today's state depends on yesterday's policy choice and state. The cycles continue, but timing becomes more intricate: A complex worldview may induce a policy that improves the state, but may be followed notwithstanding by a populist backlash. This generates patterns where populist victories coincide with improved conditions due to policies implemented under the prior technocratic regime, suggesting that populist backlashes need not follow economic downturns.

[4] Similarly, the 2016 Brexit referendum marked the end of the UK's integration within the EU, and its voting turnout was the highest ever for a UK-wide referendum, and the highest for any national vote since 1992.

[5] Notwithstanding, we show that a higher worldview complexity does not always imply better policies: more technocratic platforms can lead to worse policy-making. For instance, the multiplication of complex trade agreements may well lead to welfare losses for some social groups or countries (see, e.g., Rodrik (2011)).

We also examine a variant where "intellectuals" compete to supply worldviews while "politicians" separately compete on policy platforms, taking worldviews as given. Then, while complexity dynamics remain cyclical, the winning policy at a given stage may be suboptimal according to the prevailing worldview. Moreover, policy backlashes may precede worldview backlashes. Indeed, when less sophisticated voters become pivotal, winning politicians may cater to their simpler worldviews even as the intellectual discourse remains complex.

1.1 Related literature

Our work builds most directly on recent papers studying how political actors compete by offering simplified models of the world. Izzo, Martin and Callander (2023) develop a framework where political parties with different policy preferences compete by developing "ideologies", namely, linear-regression models that voters use to interpret data. Voters adopt the ideology that best explains their observations and then vote for the party whose policy maximizes utility according to that ideology. Our model differs in five key respects. First, we consider a dynamic framework, instead of a static one. Second, we introduce heterogeneous sophistication among voters, which is absent in their framework but central to our dynamics. Third, we introduce an incumbent-challenger asymmetry through the assumption that incumbents cannot change their worldviews (or equivalently, are strongly penalized for doing so), while entrants are free to choose the optimal worldview given the available data and the incumbent's worldview. Fourth, we take an epistemic approach: in our model, voters (and politicians) have identical preferences over outcomes, and disagree only on the policies' outcome distributions – hence, electoral competition is about worldviews rather than policy preferences. Lastly, we generalize their modeling approach by using analogy-based partitions (Jehiel 2005, 2022) rather than linear regressions, allowing for richer patterns of correlation beyond linear causal relationships.

Closely related, Montiel Olea and Prat (2025) extend the Izzo et al. framework by generalizing the class of statistical models voters can consider, and assume that to evaluate the politicians' promised policies, voters form their own worldview using *both* politicians' worldviews. Consequently, politicians shape their worldviews so as to influence the voters' perception not only of their own promises, but also of their rival's promises. Montiel Olea and Prat (2025) thus provide an elegant decomposition of worldview choice into three components: fit (how well the worldview explains the available data), simplicity (a penalty for possibly overfitting the data), and fear (suggesting that the rival worldview leads to bad outcomes). By contrast, in our (dynamic) framework, voters adopt *the* worldview that best explains the available data, but differ in their sophistication: hence, the entrant politician shapes its worldview and notably its complexity so as to capture the largest share of voters given the latter's sophistication constraints and the incumbent's worldview.

Our paper relates to other strands of literature including that on persuasion through models as initiated by Schwartzstein and Sunderam (2021) and that on competing narratives as initiated by Eliaz and Spiegler (2020). While the former shares with our approach the idea that models are selected on the basis of their plausibility given the observed data, we differ in that we impose constraints on the complexity bound used by voters/agents to assess worldviews and that we focus on the effect of competition and dynamics in contrast to that literature.[6] Regarding Eliaz and Spiegler (2020), they view narratives as describing causal links between variables, which can be related to our modeling of worldviews (except that we adopt the framework of Jehiel (2005) rather than Spiegler (2016)), but a key distinction is that narratives in Eliaz and Spiegler (2020) are selected based on the payoff they promise rather than on their plausibility. As for the models on persuasion through models, our formulation in terms of constraints on the sophistication of voters has no counterpart in Eliaz and Spiegler (2020) and subsequent papers that have built on it (see in particular, Eliaz, Galperti and Spiegler (2024) for a political economy application and Eliaz and Spiegler (2025) for an application to media).

2 Model

Time is discrete and indexed by $t \in (-\infty, +\infty)$. There is a continuum of voters of mass 1. In each period, an election takes place between the incumbent politician, elected in the previous period, and a new entrant. After the election has taken place, the state of the world ω is drawn from a set $\Omega \equiv \{\omega_1, ..., \omega_{N_\Omega}\}$, with $N_\Omega < \infty$, with i.i.d. draws across periods according to a distribution G with full support on Ω.

After observing the state of the world, the elected politician chooses a policy $a \in A \equiv \{a_1, ..., a_{N_A}\}$, with $2 \leq N_A < \infty$. To make things interesting, we assume that $N \equiv N_\Omega N_A \geq 3$. A given state of the world $\omega \in \Omega$ and a given policy $a \in A$ result in a

[6] This also applies to the more recent developments in Aina (2024) and Schwartzstein and Sunderam (2025).

random outcome $\tilde{y} \in \{0,1\}$ (failure or success).[7] For any $(\omega, a) \in \Omega \times A$, we denote by $y(\omega, a)$ the probability that $\tilde{y} = 1$ given (ω, a), measured according to the true outcome distribution.

The distribution G of states of nature is known by the agents.[8] By contrast, the outcome distributions, $y(\cdot,\cdot)$, are unknown to the agents (voters and politicians alike). In each period, the state of the world, the chosen policy and the realized outcome are publicly observed. Voters and politicians have perfect recall.

Elections and worldviews. Each period begins with an election between two politicians, the incumbent (elected in the previous period) and a (new) challenger, who enters the game in that period.[9] Politicians compete by proposing *worldviews.* A worldview $v \equiv (p, (\beta(\alpha))_{\alpha \in p})$ is a partition p of the set $\Omega \times A$ in a finite number of *analogy classes*, together with, for each analogy class $\alpha \in p$, a distribution of outcomes, captured by the conditional probability of success $\beta(\alpha) \in [0,1]$. A worldview thus clusters together couples of states and policies, (ω, a), attributing to all couples in the same class, the same distribution of outcomes. Hence, a worldview can be thought of as a simplified model for outcome distributions.

We say that a worldview v is *k-complex* if its partition of $\Omega \times A$ has k analogy classes. A more complex worldview thus offers a finer model of outcome distributions.

The incumbent cannot change its worldview from one period to the next, and thus re-proposes the worldview it proposed in the previous period. We consider several microfoundations for this assumption, which we discuss in Section 8.[10] By contrast, the entrant is free to choose any worldview it likes.

Voters. Voters have identical preferences over policy outcomes (described below), but

[7]Our analysis extends straightforwardly to more differentiated outcomes, $\tilde{y} \in Y$, with $|Y| \geq 2$.

[8]When G is the uniform distribution over Ω, an alternative interpretation of our model is that the states ω_i represent different (independent) policy issues or domains, and in each period, the elected politician chooses a policy $a \in A$ for each of the domains.

[9]The entrant becomes the next-period incumbent if she wins the election, and leaves the game if she loses the election.

[10]Our two main microfoundations are: (i) voters turn away from politicians who change their stance, consistently with a large empirical evidence (see, e.g., the literature review in Kartik and McAfee, 2007), and thus, in the non-limit case with horizontal shocks, changing its stance would annihilate the incumbent's chances of being reelected; (ii) there is a pool of politicians, each firmly believing in a specific worldview, and a primary within the opposition party selects the politician believing in the most effective worldview to win the election against the incumbent. This second interpretation explains not only why the incumbent does not change their worldview from one period to another, but also why politicians "commit" to select a policy according to their worldview (see below).

differ in their *sophistication* with respect to the assessment of worldviews – or put differently, the time and energy they can devote to understanding politicians' claims, or their (dis)taste for complexity. A voter's sophistication is measured by the maximum number of analogy classes that a worldview can contain for the voter still to understand (or tolerate) it. For any integer $k \geq 1$, we let $\mu_k \in [0,1]$ be the mass of voters with sophistication k, i.e., who understand any worldview with at most k analogy classes.

In each election, voters take into consideration only the worldviews they understand: a voter with sophistication k takes into consideration all worldviews (if any) with at most k analogy classes.

Voters are myopic. In each election, a voter votes for the worldview that best explains the available data (breaking ties randomly) *among the worldviews that the voter considers*, if any. Formally, within the worldviews they consider, a voter votes for a worldview that minimizes the Kullback-Leibler divergence (or relative entropy) between the distribution of observations $(\omega, a, \tilde{y})$ in the available data, d, and the distribution induced by the worldview.[11]

If both the incumbent's and the entrant's worldviews exceed a voter's sophistication, the voter does not participate in the election.

The politician who receives the most votes wins the election.

Politicians' electoral objectives. Politicians are myopic, and aim at maximizing their probability of being elected in the current election. We restrict attention to Markov strategies for the entrants, i.e., that depend only on the current data $d \in \mathcal{D}$ and on the incumbent's worldview $v^I \in \mathcal{P}$. The entrant's electoral strategy is thus a mapping ν from the product set of all possible data and possible worldviews, $\mathcal{D} \times \mathcal{P}$, to the set of distributions over worldviews, $\Delta\mathcal{P}$, such that for current data d and incumbent's worldview v^I, the entrant chooses worldview v with probability $\nu_{d,v^I}(v) \in [0,1]$.

While there may exist several worldviews allowing the entrant to defeat the incumbent with probability 1, we assume that the entrant has lexicographic preferences on firstly, maximizing its vote share,[12] and secondly, explaining the data. Hence, the entrant

[11] The Kullback-Leibler divergence can be interpreted as the average (according to the frequencies observed in the data) of the log of the likelihood ratios between the data and the worldviews' predictions. Analogous results would obtain if instead of the Kullback-Leibler divergence, agents used standard distances over $\Omega \times A$, such as, e.g., the square of the Euclidean distance.

[12] We could equivalently replace the vote share (ratio of own votes over total votes) with the vote ratio (ratio of own votes over rival's votes) as the entrant's vote ratio is a strictly increasing function of its vote share.

chooses a worldview that maximizes its vote share (i.e., the ratio of its own votes over total votes), and among worldviews that maximize its vote share, it chooses a worldview that minimizes the Kullback-Leibler divergence with the data. We microfound the first dimension of the entrant's objective by assuming that a voter's choice is firstly determined by how convincing the politicians worldviews seem to the voter (consistently with our epistemic approach), yet that the voter's understanding of the worldviews is noisy – e.g., due to small "mistakes" from voters in computing the Kullback-Leibler divergences, or from politicians' in describing/explaining their worldviews.[13] We microfound the second dimension of the entrant's preferences by introducing exogenous, idiosyncratic sympathy/hostility shocks (orthogonal to worldviews), realized after politicians have chosen their worldviews but before the election takes place, and that induces voters to either actually vote for the politician offering their preferred worldview, or abstain. We then take the limit as both noises/shocks vanish.

Policy-making. In each period, after the election has taken place, the current-period state of the world $\omega \in \Omega$ is realized, and the elected politician then chooses a policy $a \in A$. Implementing a policy a in state ω entails a direct and immediate cost $c(\omega, a) \in \mathbb{R}_+$. While the outcome distributions are unknown to the agents, the cost function $c : \Omega \times A \to \mathbb{R}_+$ is known by the politicians. Intuitively, the cost function c captures only the direct and immediate costs of a policy – e.g., its direct budgetary cost, excluding any indirect and/or long-term effects. By contrast, the policy outcome $\tilde{y}$ captures the latter, which agents can anticipate only via a model of the world, i.e., a worldview.

Voters share identical preferences over policies and outcomes: Assuming that the payoff difference between success and failure is 1, a voter's current-period payoff, given state of the world ω, policy a and expected success probability $y(\omega, a)$, is equal to

$$y(\omega, a) - c(\omega, a).$$

Once in office, the elected politician's policy-making strategy is a mapping from the set of couples of worldviews and states, $\mathcal{P} \times \Omega$, to the set of distributions over actions ΔA, associating to a worldview $v \in \mathcal{P}$ and a state $\omega \in \Omega$ a distribution $\sigma(\omega|v) \in \Delta A$.[14] We

[13] See Section 8.1 for details.

[14] An equivalent assumption would be that the entrant politician offers during the election both its worldview strategy $\nu \in \Delta\mathcal{P}$ and the corresponding set of policy strategies, $\sigma(\cdot|v)$ for all v in the support of ν.

assume that the elected politician chooses the current-period policy to maximize voters' (current-period) welfare according to the worldview it chose in the election.[15] Hence, only the distributions $\sigma(\omega|v)$ with v in the support of the politician's worldview strategy $\nu \in \Delta\mathcal{P}$ are outcome-relevant.

To ensure the dataset covers all cases in the ergodic distribution,[16] we add a "tremble": With probability $1 - \varepsilon \in (0, 1)$, where $\varepsilon > 0$ is small, the elected politician chooses a policy $a^*(\omega)$ with strictly positive probability only if

$$a^*(\omega) \ \in \ \arg \max_a \ \beta(\alpha(\omega, a)) - c(\omega, a)$$

where $\alpha(\omega, a)$ is the analogy class to which (ω, a) belongs according to the politician's worldview v. By contrast, with the complementary probability, $\varepsilon \in (0, 1)$, the elected politician chooses an action $a \in A$ randomly, with uniform distribution over A.

Lastly, the policy outcome $\tilde{y}$ is realized and publicly observed, which, together with the realized state ω and the chosen policy a, creates a new observation, $(\omega, a, \tilde{y})$, that is added to the data to be explained in the next period.

Equilibrium concept. Our assumption on policy-making determines the set of equilibrium policy strategies for the elected politician after any worldview and state of nature. We thus look for a Markov Perfect equilibrium of the electoral game in which in each period, given data d and the incumbent's worldview v^I, the current-period entrant chooses a Markov strategy $\nu_{d,v^I} \in \Delta\mathcal{P}$ to maximize in a lexicographic order, its vote share in the current election, and the fit of its worldview with the data (minimizing its Kullback-Leibler divergence with the data).

Remark: Noises and trembles. In the elaborate version of our model (see Section 8.1), we assume noise in the voters' understanding of a worldview and in their idiosyncratic preference for one politician or the other, and a tremble for the politician in office when choosing a policy. We then consider the limit in which all noise and trembles disappear.

[15] The most obvious microfoundations for this assumption are that politicians do believe in the worldviews they promote, and/or that any deviation away from the optimal policy according to the worldview (any "inconsistency") will be penalized by voters in the next election.

[16] While we assume a tremble mainly for expositional clarity, it could be realistically microfounded by supposing that in each period the elected politician's objective is subject to (rare) random shocks – e.g., due to personal interests or tastes affecting a politician's willingness to implement a certain policy, a, or to short-term variations in the immediate policy costs, $c(\cdot)$.

These (small) perturbations play distinct roles in our analysis: the noise on the voters' understanding of worldviews and in their idiosyncratic preferences yields for any incumbent's worldview, the selection of a unique best-response for the challenger politician, while the tremble for the politician in office yields that, in equilibrium, all frequencies of observations $(\omega, a, \tilde{y})$ coincide with the probabilities according to the true data generating process. Hence, the noise on the voters' side removes an equilibrium multiplicity stemming from the challenger politician being indifferent over several best-responses, while the tremble on the policymaking removes an equilibrium multiplicity stemming from unobserved frequencies.

3 Elections, worldviews and cycles

3.1 Comparing worldviews

We begin by fixing exogenously the set of available observations, d, to be interpreted by politicians and voters in a given period. We will endogenize this set in Section 4 (it will then consist of all data accumulated over previous electoral terms). As the set of available information, d, is infinitely large (since time runs from $-\infty$ and agents have perfect recall), the frequencies of observations remain constant over any finite number of periods.

For any $(\omega, a, \tilde{y}) \in \Omega \times A \times \{0, 1\}$, we denote by $f(\omega, a, \tilde{y})$ the frequency of $(\omega, a, \tilde{y})$ in d. For simplicity, we assume that $f(\omega, a, 0) + f(\omega, a, 1) > 0$ for any $(\omega, a, \tilde{y}) \in \Omega \times A \times \{0, 1\}$,[17] but we allow $f(\omega, a, \tilde{y})$ to be arbitrarily close to 0. For any set α of couples (ω, a), we denote by $\hat{m}(\alpha)$ the *empirical mass* of the analogy class α in the observed data:

$$\hat{m}(\alpha) \equiv \sum_{(\omega,a)\in\alpha} [f(\omega, a, 0) + f(\omega, a, 1)].$$

When $\alpha = \{(\omega, a)\}$, the empirical mass $\hat{m}(\omega, a)$ represents the frequency of state ω realizing and policy a being chosen in state ω. For any $(\omega, a) \in \Omega \times A$, let

$$\hat{y}(\omega, a) \equiv \frac{f(\omega, a, 1)}{f(\omega, a, 0) + f(\omega, a, 1)}$$

[17]When we endogenize the set of available information as the set of past observations, the "tremble" will ensure that this assumption holds.

denote the *empirical probability of success* of policy a in state ω, and for any analogy class α, let $\hat{b}(\alpha)$ denote the empirical probability of success conditional on $(\omega, a) \in \alpha$, i.e., the mean frequency of successes in α:[18]

$$\hat{b}(\alpha) \equiv \frac{\sum_{(\omega,a)\in\alpha} f(\omega, a, 1)}{\sum_{(\omega,a)\in\alpha} [f(\omega, a, 0) + f(\omega, a, 1)]}.$$

The Kullback-Leibler divergence between the distribution in the observed data and the distribution induced by a worldview $v \equiv (p, (\beta(\alpha))_{\alpha\in p})$ writes as

$$\sum_{(\omega,a)} f(\omega, a, 0) \ln \left(\frac{f(\omega, a, 0)}{[f(\omega, a, 0) + f(\omega, a, 1)](1 - \beta(\alpha(\omega, a)))} \right) + \sum_{(\omega,a)} f(\omega, a, 1) \ln \left(\frac{f(\omega, a, 1)}{[f(\omega, a, 0) + f(\omega, a, 1)]\beta(\alpha(\omega, a))} \right).$$

As mentioned earlier, the Kullback-Leibler divergence corresponds to the average (according to the frequencies observed in the data) of the log of the likelihood ratios between the data and the worldviews' predictions.

Consequently, minimizing the Kullback-Leibler divergence between the two distributions is equivalent to maximizing the following objective:

$$\sum_{\alpha\in p} \hat{m}(\alpha) \Big((1 - \hat{b}(\alpha)) \ln \big(1 - \beta(\alpha)\big) + \hat{b}(\alpha) \ln \big(\beta(\alpha)\big) \Big).$$

We refer to the k-complex worldviews that minimize the KL divergence over all k-complex worldviews as *k-optimal worldviews.* Our first result provides a characterization.

Lemma 1 (**k-optimal worldviews**). *Any k-optimal worldview $v \equiv (p, (\beta(\alpha))_{\alpha\in p})$ with $k \geq 1$ is such that the conditional probabilities of success $(\beta(\alpha))_{\alpha\in p}$ are equal to the empirical probabilities of success:*

$$\beta(\alpha) = \hat{b}(\alpha). \tag{1}$$

[18] In the case of a singleton $\alpha = \{(\omega, a)\}$, $\hat{b}(\alpha) = \hat{y}(\omega, a)$.

and the partition p maximizes over all partitions with k classes the following objective:

$$\sum_{\alpha \in p} \hat{m}(\alpha)\Big((1-\hat{b}(\alpha))\ln\left(1-\hat{b}(\alpha)\right)+\hat{b}(\alpha)\ln\left(\hat{b}(\alpha)\right)\Big), \tag{2}$$

Intuitively, any partition that minimizes the Kullback-Leibler divergence tends to pool in the same analogy class couples (ω, a) that exhibit similar probabilities of success in the data, $\hat{y}(\omega, a)$. In fact, in any optimal partition, all classes are intervals (see, e.g., Banerjee et al, 2005).[19] We refer to Appendix I for additional properties of optimal partitions.

3.2 The dynamics of worldviews and complexity

The main intuition driving complexity dynamics is very simple. Firstly, for any data and any incumbent's worldview, any worldview that is not k-optimal for some k is strictly dominated for the entrant. Therefore, on path, the entrant always faces a k-optimal strategy from the incumbent, for some k.[20] Secondly, as time runs from minus infinity and agents have perfect recall, in any period t, agents have access to an infinitely large set of observations. Consequently, the frequencies of observations do not change from one period to another, and thus, for all $k \geq 1$, the set of k-optimal worldviews remains the same from one period to the other. Therefore, when facing an incumbent with a k-optimal worldview, the entrant cannot outperform the incumbent's worldview with an equally complex worldview.

So, suppose the entrant faces a k-complex incumbent's worldview, with $1 < k < N$. Intuitively, for the entrant, the "most effective" k'-complex worldviews with $k' \geq k+1$ are the $(k+1)$-optimal worldviews. Indeed, such worldviews convince all voters with sophistication at least $k+1$ by outperforming the incumbent's k-complex worldview, and leaves only voters with sophistication exactly k to the incumbent, for whom the entrant's worldview is too complex.[21] Intuitively again, when the distribution of voters'

[19]Given three couples (ω_1, a_1), (ω_2, a_2), (ω_3, a_3) with $y(\omega_1, a_1) < y(\omega_2, a_2) < y(\omega_3, a_3)$ (and $\hat{m}(\omega_i, a_i) > 0$ for $i = 1, 2, 3$), if (ω_1, a_1) and (ω_3, a_3) belong to the same class in an optimal partition, then (ω_2, a_2) belongs to that same class.

[20]Off-path, the entrant may face an incumbent's k-complex worldview that is not k-optimal. The entrant's strict best-response is then to choose a k-optimal worldview as any such worldview yields a vote share equal to 1.

[21]Formally, for all $k' > k+1$,

$$\frac{\mu_{k'} + \mu_{k'+1} + \dots}{\mu_k + \mu_{k+1} \dots} < \frac{\mu_{k+1} + \mu_{k+2} + \dots}{\mu_k + \mu_{k+1} + \dots}.$$

sophistication is not too "irregular", choosing a $(k+1)$-optimal worldview allows the challenger to win the election with probability 1 if there are sufficiently many voters with sophistication at least $k+1$, i.e., if the complexity k of the incumbent's worldview k is not too high.

By contrast, the "most effective" k'-complex worldview with $k' \leq k-1$ is *the* 1-optimal worldview.[22] Indeed, while such a worldview is less convincing than the incumbent's for any voter able to understand both, the entrant's simple worldview is understandable by *all voters*. Hence, it enlists voters with sophistication $1 \leq k' \leq k-1$, for whom the incumbent's worldview is too complex, and leaves voters with sophistication $k' \geq k$ to the incumbent. Choosing the 1-optimal worldview thus allows the challenger to win the election with probability 1 if there are sufficiently many voters with sophistication strictly below k, i.e., if the complexity k of the incumbent's worldview k is high enough.

Therefore, the entrant eventually offers either a slightly more complex, or a drastically simpler worldview than the incumbent's. Its choice depends on the incumbent's worldview complexity, and on the distribution of voters.Specifically, for a k-complex incumbent worldview, the entrant chooses a (distribution over) $(k+1)$-complex worldview(s) if

$$\frac{\mu_{k+1}+\mu_{k+2}+\ldots}{\mu_k+\mu_{k+1}+\mu_{k+2}+\ldots} > \frac{\mu_1+\mu_2+\ldots+\mu_{k-1}}{\mu_1+\mu_2+\ldots+\mu_N},$$

and a 1-complex worldview if the opposite inequality holds.

We assume without loss of generality that $k \leq N_\Omega N_A \equiv N$ (e.g., by regrouping all voters with $k \geq N$ in the category $k = N$), and for simplicity, that $\mu_k > 0$ for all $1 \leq k \leq N$. Let $k^* \geq 2$ be the lowest integer k such that

$$\frac{\mu_{k+1}+\mu_{k+2}+\ldots}{\mu_k+\mu_{k+1}+\mu_{k+2}+\ldots} \leq \mu_1+\mu_2+\ldots+\mu_{k-1},$$

and let us assume for simplicity and genericity that the inequality holds strictly at k^*. By definition of k^*, the mass of voters with sophistication higher than k^* $(\mu_{k^*+1}+\mu_{k^*+2}+\ldots)$ is strictly lower than the mass of voters with sophistication lower than k^* $(\mu_1+\ldots+\mu_{k^*-1})$,

[22]Formally, for all $1 < k' \leq k-1$,

$$\frac{\mu_{k'}+\mu_{k'+1}+\ldots+\mu_{k-1}}{\mu_{k'}+\mu_{k'+1}\ldots} < \frac{\mu_1+\mu_2+\ldots+\mu_{k-1}}{\mu_1+\ldots} = \mu_1+\mu_2+\ldots+\mu_{k-1}.$$

And for any data set d, there exists a unique 1-optimal worldview.

and possibly much lower. As an illustration, consider the case of a uniform distribution of sophistication: $\mu_k = 1/N$ for all k. Then,

$$k^* = \lfloor N - \sqrt{N} + 1 \rfloor.$$

Hence, $k^* > N/2$ for all N, and as N goes to ∞, k^* becomes equivalent to N.[23]

We make two more assumptions before stating our next result. Firstly, we assume that for any $k \leq k^* - 1$,

$$\frac{\mu_{k+1} + \mu_{k+2} + ...}{\mu_k + \mu_{k+1} + \mu_{k+2} + ...} > \frac{1}{2}, \qquad \text{and} \qquad \mu_1 + \mu_2 + ... + \mu_{k^*-1} > \frac{1}{2},$$

and that either $k^* < N$, or $\mu_N < \mu_1 + ... + \mu_{N-1}$. Intuitively, this assumption holds whenever the distribution of μ_k is sufficiently "smooth", which is a realistic benchmark.[24]

Secondly, for simplicity, we assume that there is sufficient variation in the true data generating process that a k-complex worldview with $k < k^*$ cannot explain perfectly the available data set. Specifically, we assume that in the case of an exogenous data set (as in this Section), resp. in the case of an endogenous data set (as will be the case in Section 4), that the success frequencies $\hat{y}(\omega, a)$, resp. $y(\omega, a)$, take at least k^* different values as (ω, a) spans $\Omega \times A$.

Lemma 2 (**Complexity dynamics: Gradual shifts and backlashes**)**.** *When facing an incumbent with a k-optimal worldview where $k \leq k^* - 1$, the entrant offers a $(k+1)$-complex worldview. When facing an incumbent with a k^*-optimal worldview, the entrant offers a worldview with a single analogy class (1-complex). For any incumbent's k-complex worldview, with $1 \leq k \leq k^*$, the entrant wins the election with probability 1.*

Our next result collects the implications of Lemmas 1 & 2.

Proposition 1 (**Exogenous information: Complexity and participation dynamics**)**.** *In any equilibrium, complexity and participation dynamics are deterministic and cyclical:*

[23] As another illustration, if the distribution of sophistication has a unique maximum, $\mu_{\hat{k}} > \max_{k \neq \hat{k}} \mu_k$, and a mass on the right of its maximum at least as high as the one on the left ($\Sigma_{k \geq \hat{k}+1} \mu_k \geq \Sigma_{k \leq \hat{k}-1} \mu_k$), then $k^* > \hat{k}$. In this case as in the case of the uniform distribution, $\mu_1 + ... + \mu_{k^*-1} > 1/2$, and thus, a k^*-complex worldview is taken into consideration by strictly less than half of the voters.

[24] For instance, it is unlikely that the distribution of voters would have a very large (resp. very low) mass on some k, but much lower (resp. much larger) ones on $k-1$ and $k+1$.

(i) A cycle lasts for k^* periods.

(ii) Gradual increases in complexity and backlashes towards maximum simplicity: The complexity of the winning worldview increases gradually until it reaches k^*, at which point it falls down to the minimum level (1), and a new cycle begins.

(iii) Gradual decline in participation and surges: As the complexity of the winning worldview increases, participation decreases. The backlash towards simplicity induces a surge in participation, which reaches its maximum level, and a new cycle begins.

Strikingly, complexity and participation dynamics are deterministic. They exhibit cycles of gradually rising complexity and abstention, interrupted by a backlash towards maximum simplicity and maximum participation, which restarts the cycle.[25] Put differently, a cycle can be interpreted as a gradual shift towards technocracy, interrupted by a populist backlash, followed again by a progressive return towards technocracy.

To illustrate Proposition 1, let us return to the uniform distribution benchmark. Then, $k^* > N/2$ for all N, i.e., at peak complexity (k^*), more than half of the voters do not participate in the election.[26] Moreover, as N goes to infinity, k^*/N converges to 1, i.e., the mass of voters participating in the election at the cycle's peak (k^*-complex winning worldview) converges to zero (the share of abstention converges to one). Furthermore, as N goes to infinity, the length of the cycle (k^* periods) goes to infinity at speed N.

Remark: Non-deterministic cycles. The entrant's lexicographic preferences on firstly, maximizing its vote share, and secondly, explaining the data derive from (exogenous, idiosyncratic) noise in the voters' understanding of worldviews, and horizontal shocks on their sympathy/hostility towards politicians. We then take the limit as noises/shocks vanish. Nonetheless, in the non-limit case in which such noises/shocks are non-zero, non-deterministic cycles would obtain: While the entrant's best responses would be unchanged, the incumbent would win the election with a strictly positive probability thanks to the noises/shocks. Consequently, starting from $k < k^*$, resp. from k^*, the complexity of the winning worldview would either remain in k or move to $k+1$, resp. to 1.

[25]The insight is even more general than Proposition 1 suggests. Suppose for instance that, due to an unforeseen shock, at time t, the incumbent has a k'-complex worldview with k'^*. Then, the complexity of the winning worldview first rises up to some $k'' > k'$, before eventually falling back to 1, and from then on, the "standard" cycle resumes, from 1 to k^*.

[26]The same holds if the distribution of sophistication has a unique maximum, $\mu_{\hat{k}} > \max_{k \neq \hat{k}} \mu_k$, and a mass on the right of its maximum at least as high as the one on the left ($\Sigma_{k \geq \hat{k}+1} \mu_k \geq \Sigma_{k \leq \hat{k}-1} \mu_k$).

4 Endogenous data and ergodicity

Let us now "close" our model, endogenizing the data set available to voters as the result of past policy choices. By Proposition 1, in any equilibrium, complexity dynamics are deterministic and cyclical, and do not depend on the dataset available to voters in a given period. A cycle lasts for k^* periods. Consequently, we restrict attention to cyclical profiles of worldviews and policy strategies, with cycle length k^*: the entrant's electoral strategy thus depends only on the cycle's stage.

We refer to the k-th stage of a cycle as the stage in which the (entrant's) winning worldview is k-complex. At each stage of the cycle, the entrant first chooses a worldview, and then conditional on winning the election and after observing the state of the world, chooses a policy. Let $\Delta\mathcal{P}^k$ denote the set of distributions over the set $\mathcal{P}^k$ of worldviews with partitions of $\Omega \times A$ with (exactly) k elements. We define a worldview profile, $\boldsymbol{\nu} \equiv (\nu^k)_{1\le k\le k^*} \in \Delta\mathcal{P}^1 \times ... \times \Delta\mathcal{P}^{k^*}$ as a vector of worldview strategies, i.e. such that for each $1 \le k \le k^*$ and any k-complex worldview v, the stage-k entrant chooses worldview v with probability $\nu^k(v)$.

Similarly, we define a policy profile, $\boldsymbol{\sigma} \equiv (\sigma^k)_{1\le k\le k^*} \in (\Delta A)^{N_\Omega|\mathcal{P}^1|} \times ... \times (\Delta A)^{N_\Omega|\mathcal{P}^{k^*}|}$, as a vector of policy strategies such that for each $1 \le k \le k^*$, for any $\omega \in \Omega$ and $v \in \mathcal{P}^k$, $\sigma^k(\omega|v) \in \Delta A$ denotes the policy strategy at the k-th stage of the cycle as a function of the current state of nature, $w \in \Omega$, and the chosen worldview, $v \in \mathcal{P}^k$.

Definition 1 (**Optimal worldview-policy profile**). *Fix the available data set d. A worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ is optimal given d if for any $1 \le k \le k^*$, for any worldview $v \in \mathcal{P}^k$, $\nu^k(v) > 0$ only if v is an optimal worldview at stage k of the cycle given data d, and for any $\omega \in \Omega$, $a \in A$ and $v \in \mathcal{P}^k$ such that $\nu^k(v) > 0$, $\sigma^k(\omega|v)(a) > \varepsilon/N_A$ only if policy a is optimal in state ω according to worldview v.*

For any strictly positive tremble $\varepsilon > 0$, a worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ induces a unique ergodic distribution of frequencies of observations $(\omega, a, \tilde{y})$ for all $(\omega, a, \tilde{y}) \in \Omega \times A \times \{0, 1\}$.[27]

Definition 2 (**Consistency of the data**). *A (infinite) data set d is consistent with a worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ if the frequencies of observations $(\omega, a, \tilde{y})$ in d are equal to the frequencies in the ergodic distribution induced by $(\boldsymbol{\nu}, \boldsymbol{\sigma})$.*

[27]Formally, for any $(\omega, a, \tilde{y}) \in \Omega \times A \times \{0, 1\}$, the frequency $f_{(\boldsymbol{\nu},\boldsymbol{\sigma})}(\omega, a, \tilde{y})$ in the ergodic distribution

Definition 3 (**Ergodic worldview-policy profile**)**.** *A worldview-policy profile* $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ *is ergodic if it is optimal given its consistent data set.*

Proposition 2 (**Ergodic worldview-policy profile**)**.** *There exists an ergodic worldview-policy profile.*

In other terms, Proposition 2 yields the existence of an *equilibrium* characterized by a worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ such that, when politicians and voters behaved according to $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ in all previous periods/cycles, a politician's best-reply in the current period/cycle is $(\boldsymbol{\nu}, \boldsymbol{\sigma})$. We note that our definition allows for mixing both about worldviews and policies, as this is required to guarantee existence. Some illustrations of this will be appear below.

5 Implications

Let us study a few simple examples to illustrate the implications of our framework. To alleviate the notation when describing ergodic worldview-policy profiles, we henceforth omit the full description of policy strategies and restrict attention to the policy strategies $\sigma^k(\omega|v) \in \Delta A$ such that v lies in the support of ν^k. (As before, we take the limit of the ergodic profiles as the policy-making tremble vanishes: $\varepsilon \to 0$.)

5.1 Cost-minimization (short-termism): State-driven vs policy-driven outcomes

For any state ω and class α, we refer to the *cost-minimizing policy* in state ω and class α as the policy a such that $(\omega, a) \in \alpha$ and $a \in \arg \min_{a'|(\omega,a')\in\alpha} c(\omega, a')$. We refer to the (global) cost-minimizing policy in state ω as the policy $a \in A$ such that $a \in \arg \min_{a'\in A} c(\omega, a')$.

---

induced by $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ is given by

$$f_{(\boldsymbol{\nu},\boldsymbol{\sigma})}(\omega, a, 1) = \frac{1}{k^*} \sum_{k=1}^{k^*} \sum_{v\in P^k} \nu^k(v) G(\omega) \sigma^k(\omega|v)(a) y(\omega, a),$$

$$f_{(\boldsymbol{\nu},\boldsymbol{\sigma})}(\omega, a, 0) = \frac{1}{k^*} \sum_{k=1}^{k^*} \sum_{v\in P^k} \nu^k(v) G(\omega) \sigma^k(\omega|v)(a) [1 - y(\omega, a)].$$

The tremble $\varepsilon > 0$ ensures that $\sigma^k(\omega|v)(a) > 0$ for all k, ω, v, a, and thus ensures the uniqueness of the ergodic distribution.

Because any 1-complex worldview attributes the same distribution of outcomes to all couples (ω, a), any 1-complex worldview induces the ruling politician to choose the global cost-minimizing policy in each state.

Any worldview not complex enough to distinguish all couples (ω, a) with different outcome distributions in the data (empirical probabilities of success $\hat{y}(\omega, a)$) must pool together several couples (ω, a) with different empirical probabilities of success. As noted above, any partition that minimizes the Kullback-Leibler divergence tends to pool in the same analogy class couples (ω, a) that exhibit similar probabilities of success in the data. When most of the variation in the observed data stems from the state of the world ω, the "most convincing" simple partitions are those that tend to be based on the state of the world ω, i.e., such that two couples (ω, a), (ω', a') belong to the same class if $\omega = \omega'$.

By a similar argument, when most of the variation in the data stems from the policy a, the most convincing partitions are those that tend to be based on the policy a, i.e., such that two couples (ω, a), (ω', a') belong to the same class if $a = a'$. In other words, partitions based on the state of the world attribute the same conditional probability of success to all possible actions in a given state, whereas partitions based on the policy attribute the same conditional probability of success to that policy in all states. Partitions based on the state of the world are thus likely to emerge for, e.g., a small country when the policy issue is related to global warming (on which the small country's individual policies have little impact), whereas partitions based on the action are likely to emerge for, e.g., a large country facing a domestic issue.

Formally, we say that outcomes are *state-driven* if for any $k \leq N_\Omega$, any k-optimal worldview given the true frequencies ($\hat{y}(\omega, a) = y(\omega, a)$) is such that for any $\omega \in \Omega$, any two couples $(\omega, a), (\omega, a')$ with $a, a' \in A$, belong to the same analogy class. Similarly, we say that outcomes are *policy-driven* if for any $k \leq N_A$, any k-optimal worldview given the true frequencies ($\hat{y}(\omega, a) = y(\omega, a)$) is such that for any $a \in A$, any two couples $(\omega, a), (\omega', a)$ with $\omega, \omega' \in \Omega$ belong to the same analogy class. We refer correspondingly to *state-driven*, resp. *policy-driven* worldviews as worldviews that partition couples (ω, a) based on the state of the world ω, resp. based on the policy a.

Consider for instance an environment in which outcome distributions are lexicographically ordered according to the state: for any $\omega, \omega' \in \Omega$,

$$\max_a y(\omega, a) < \min_a y(\omega', a) \qquad \text{or} \qquad \min_a y(\omega, a) > \max_a y(\omega', a).$$

In such an environment, for any $k \leq N_\Omega$, any k-optimal worldview is state-driven.[28]

The next result follows straightforwardly from the definition of state-driven worldviews.

Proposition 3 (State-driven worldviews and cost minimization). *State-driven worldviews lead to cost minimization.*

Indeed, as a state-driven worldview clusters in the same class all policies for a given state, then in that state, the politician in office attributes the same probability of success to each policy, and thus if it chooses a policy in that class, it chooses a cost-minimizing one.[29] Quite intuitively, a state-driven worldview thus leads to short-termism – e.g., inaction in the face of climate change.

By contrast, policy-driven worldviews can lead to excessive costs, due to a form of "overconfidence" in the effectiveness of a given policy.[30] We explore this possibility in more details in the next Sections.

Remark: Fatalistic vs empowering worldviews. Eliaz and Spiegler (2020, 2025) and Eliaz, Galperti and Spiegler (2024) refer to "fatalistic" narratives as narratives in which policies have no causal impact on outcomes, and "empowering" narratives as narratives in which policies have a (strong) causal impact on outcomes. Hence, our state-driven worldviews may be considered as *fatalistic*, while our policy-driven worldviews may be considered as *empowering*. However, as we noted, policy-driven worldviews in our environment need not lead to excessive action.[31]

[28]Similarly, if outcome distributions are lexicographically ordered according to the policy: for any $a, a' \in A$,

$$\max_\omega y(\omega, a) < \min_\omega y(\omega, a') \qquad \text{or} \qquad \min_\omega y(\omega, a) > \max_\omega y(\omega, a'),$$

then for any $k \leq N_A$, any k-optimal worldview is policy-driven.

[29]Formally, a politician choosing an action according to a fatalistic worldview (with $\alpha(\omega, a) = \alpha(\omega, a'^f(\omega)$ for all a, a') solves

$$\max_a \beta(\alpha(\omega, a)) - c(\omega, a) \ = \ \beta(\alpha^f(\omega)) - \min_a c(\omega, a).$$

[30]Formally, a politician choosing an action according to an empowering worldview (with $\alpha(\omega, a) = \alpha(\omega'^e(a)$ for all ω, ω') solves

$$\max_a \beta(\alpha(\omega, a)) - c(\omega, a) \ = \ \max_a \beta(\alpha^e(a)) - c(\omega, a),$$

and may thus choose a policy a with either excessively high or excessively low cost with respect to the efficient policy.

[31]Besides, in our environment, whether the prevailing worldviews are fatalistic or empowering is determined by the available data and the worldview's credibility in explaining the data.

5.2 Cost minimization (short termism): Intermediate policy options

In many environments, policies can be altered by marginal changes – typically, by marginally varying their budget. Hence, the set A of policy options from which politicians choose, and that worldviews must model, may in practice be quite "dense", i.e., with numerous intermediate policies.

We show that more numerous intermediate policy options can foster short-termism (cost minimization). Indeed, more numerous intermediate (close) policy options have a direct effect leading to cost-minimization: when policy options have similar probabilities of success y, the politicians' worldviews pool them in same analogy class, inducing a "locally fatalistic worldview", and hence, within that analogy class, the ruling politician will only consider and possibly choose the cost-minimizing policy. Yet in addition to this direct effect, there is a feedback effect created by the endogeneity of data: if lower-cost options entail lower probabilities of success, then the conditional probability of success for a class, which is determined by the one of the cost-minimizing policy in the class, is equal to the lowest empirical probability of success in that class. Hence, to reduce the Kullback-Leibler divergence of the worldview, policies with even lower costs and probabilities of success may be optimally clustered in that class. Then again, only the cost-minimizing, and thus success-minimizing policy will be considered and possibly chosen, leading to unravelling.

As an illustration, consider the following environment.

Intermediate-policies environment. There is a single state $\Omega = \{\omega\}$. The policy set A is given by $A = \{a_0, a_1, a_2, ..., a_N\} \subset \mathbb{R}_+$ such that for any $i \geq 0$,

(i) $0 \leq \underline{c} = c(\omega, a_0) < c(\omega, a_i) < c(\omega, a_{i+1}) < c(\omega, a_N) = \overline{c}$ and $0 \leq \underline{y} = y(\omega, a_0) < y(\omega, a_i) < y(\omega, a_{i+1}) < y(\omega, a_N) = \overline{y} < 1$,[32]

(ii) $y(\omega, a_{i+1}) - y(\omega, a_i) = \delta$, with δ small, and $c(\omega, a_i) = h(y(\omega, a_i))$ where the function h is continuous and increasing.[33]

(iii) $0 = y(\omega, a_0) - c(\omega, a_0) < y(\omega, a_i) - c(\omega, a_i) < y(\omega, a_{i+1}) - c(\omega, a_{i+1})$ for $i \geq 1$.

[32] We fix the bounds $\underline{c}, \overline{c}$ and $\underline{y}, \overline{y}$ as we focus on the role of *intermediate* policy options. An interpretation of the assumption that $\overline{y} < 1$ is that, while it is always easy to cut short-term costs at the expense of (longer-term) success, it may not be feasible to achieve success for sure, regardless of the short-term costs incurred.

[33] As intuitive, the cost of achieving a given probability of success (according to the true data generating process) does not depend on δ, i.e., does not depend on the set of available policies.

To build intuition, let us consider the case $k^* = 2$ and *pure* ergodic worldview-policy profiles. Then, in any such profile, the highest efficiency achieved at any stage of the cycle goes to zero as δ goes to zero. Indeed, the 1-complex optimal worldview always selects policy a_0. Consider a 2-complex optimal worldview, with partition $\{a_0, ..., a_{i^*-1}\}, \{a_{i^*}, ..., a_N\}$ (any 2-optimal worldview must be of this form). Then, in class $\{a_{i^*}, ..., a_N\}$, the ruling politician only considers a_{i^*}, and in class $\{a_0, ..., a_{i^*-1}\}$, a_0.[34] The conditional probabilities of successes are thus

$$\beta(\{a_0, ..., a_{i^*-1}\}) = y(\omega, a_0), \quad \text{and} \quad \beta(\{a_{i^*}, ..., a_N\}) = y(\omega, a_{i^*})$$

Suppose by contradiction that $y(\omega, a_{i^*}) - c(\omega, a_{i^*})$ does not go to zero as δ goes to zero. Then, $y(\omega, a_{i^*})$ is bounded away from $y(\omega, a_0)$ as δ goes to zero (and remains bounded above by $\overline{y} < 1$). However, by construction, $y(\omega, a_{i^*}) - y(\omega, a_{i^*-1})$ converges to zero, and so, for δ sufficiently small, clustering a_{i^*-1} with $\{a_{i^*}, ..., a_N\}$ yields a strictly lower KL divergence (see Appendix I). As a consequence, $y(\omega, a_{i^*}) - c(\omega, a_{i^*})$ goes to zero, and thus the highest efficiency at any stage of the cycle goes to zero too as δ goes to zero.

We extend this intuition to any $k^* \geq 2$ and to *mixed* ergodic worldview-policy profiles to derive our next result.

Proposition 4 (**Policy options and cost minimization**). *More numerous intermediate policy options foster cost minimization and short-termism: In the intermediate-policies environment, as δ goes to zero, in any ergodic worldview-policy profile, the efficiency at any stage of the cycle goes to zero.*

Cost minimization, efficiency and extreme options. Cost minimization need not imply efficiency – as the above insight shows. In particular, when cost minimization is detrimental to efficiency – e.g., inaction in the face of climate change –, the above example shows that restricting the set of policies available can improve efficiency. This implication echoes Dewatripont and Tirole (1999), and Szalay (2005), although from a different logic: in their frameworks, restricting the set of policies available (or directions to be investigated) induces a higher effort to collect information about such policies' consequences, and the more informed, albeit constrained policy choice that ensues can be more efficient than the less informed, albeit unconstrained one. In our environment, efficiency

[34]Because the ergodic worldview-policy profile is pure, the policies played along the cycle are a_0 at stage 1, and a_i^* at stage 2.

is determined by the difference between a policy's probability of success (y) and cost (c). While the cost is always observed, probabilities of success are subject to worldviews, and the more numerous the intermediate policy options available to politicians, the more the prevailing worldviews are "locally cost-minimizing", which we show has both a direct effect and a feedback effect via the endogeneity of the data, leading to "quais-global cost-minimization".

Consequently, restricting the set of policies available can induce more "extreme" policies, and thus more "extreme" observation, leading to worldviews that put more emphasis on the policies' probability of success rather than on their short-term costs, leading possibly to a more efficient choice.

5.3 Worldview complexity and efficiency: A non-monotone relation

Let us define *efficiency* as maximizing the voters' expected payoff, $\mathbb{E}_\omega[y(\omega, a) - c(\omega, a)]$. What is the relation between worldview complexity and efficiency?

A sufficiently complex worldview can distinguish any two couples (ω, a), (ω', a') that yield different probabilities of success $y(\omega, a) \neq y(\omega', a')$. This guarantees that in any state of nature ω, a politician choosing a policy according to such a worldview solves

$$\max_a \beta(\alpha(\omega, a)) - c(\omega, a) \;=\; \max_a \hat{y}(\omega, a) - c(\omega, a),$$

and thus, when empirical frequencies coincide with the true probabilities ($\hat{y}(\omega, a) = y(\omega, a)$, the politician selects the welfare-maximizing policy. By contrast, as already noted, the cycle's simplest worldview always leads to the the global cost-minimizing policy, which can obviously be welfare-inferior in many cases.

Besides the welfare comparison between extreme worldviews, we now observe that the relation between complexity and efficiency can be non-monotone along a cycle.

Proposition 5 (**Complexity and efficiency: Intermediate worldviews**). *More complex worldviews can lead to less efficient policy choices. In particular, for $k^* \geq 3$, efficiency can first decrease then increase, or conversely, first increase then decrease during the phase of rising complexity. Lastly, the simplicity backlash can improve efficiency: efficiency can be higher at stage 1 than at stage k^*.*

To establish the first part of Proposition 5, consider the following environment, which

Stage	Worldview	Policy
$k=3$	$p_3 = \{(\omega_G,0)\},\{(\omega_B,0)\},\{\{(\omega_G,a),(\omega_B,a)\}\} \equiv \{\alpha_0,\alpha_1,\alpha_2\}$ $\beta_3(\alpha_0) = y(\omega_G,0)$, $\beta_3(\alpha_1) = y(\omega_B,0)$ and $\beta_3(\alpha_2) = 1$	0 in state ω_G a in state ω_B
$k=2$	$p_2 = \{\{(\omega_G,0),(\omega_B,0)\},\{(\omega_G,a),(\omega_B,a)\}\} \equiv \{\alpha_3,\alpha_4\}$ $\beta_2(\alpha_3) = \frac{y(\omega_G,0)+y(\omega_B,0)}{2}$ and $\beta_2(\alpha_4) = 1$	a in state ω_G, 0 in state ω_B
$k=1$	$p_1 = \{\{(\omega_G,0),(\omega_G,a),(\omega_B,0),(\omega_B,a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5) = \frac{y(\omega_G,0)+y(\omega_B,0)+1}{3}$	0 in state ω_G, 0 in state ω_B

Table 1: Ergodic worldview-policy profile.

we will use repeatedly.

Example I (lead example): Let $\Omega = \{\omega_G, \omega_B\}$ with ω_G and ω_B equally likely, and $A = \{0, a\}$, so that $N = 4$. For instance, ω_G, ω_B may denote respectively a "good" and a "bad" state – e.g., a more or less strong/widespread variant of a given virus –, while $0, a$ may denote respectively inaction/status quo and action – e.g., not vaccinating anyone, vs vaccinating the whole population at risk.

Suppose moreover that

(i) $y(\omega_B,0) < y(\omega_G,0) < y(\omega_G,a) = y(\omega_B,a) = 1$,

(ii) $c(\omega,0) = 0$ for all $\omega \in \Omega$, while $c(\omega_G,a) < c(\omega_B,a)$,

(iii) $y(\omega_G,a) - c(\omega_G,a) < y(\omega_G,0)$ and $y(\omega_B,0) < y(\omega_B,a) - c(\omega_B,a)$.

(iv) $y(\omega_B,a) - c(\omega_B,a) < \frac{y(\omega_G,0)+y(\omega_B,0)}{2} < y(\omega_G,a) - c(\omega_G,a)$.

Lastly, suppose that $k^* = 3$.

Table 1 describes a (pure) ergodic worldview-profile in this environment. Let us flesh out its composition. Since $y(\omega,a) = 1$ for all $\omega \in \Omega$, three classes are enough for a partition to perfectly explain the data. Simpler worldviews offer coarser explanations, still bundling (ω_G,a) and (ω_B,a) together. The conditional probabilities of success are then computed from (1) and the frequencies in the induced ergodic data set. Consider for instance $\beta_2(\alpha_3)$:

- observations with $(\omega_G,0)$ arise 2/6 of the time over the whole cycle (when $k = 1,3$ and $\omega = \omega_G$), and so do observations with $(\omega_B,0)$ (when $k = 1,2$ and $\omega = \omega_B$);
- hence, the frequencies of $(\omega_B,0,1)$ and $(\omega_G,0,1)$ in the ergodic data set are equal to $\frac{y(\omega_B,0)}{3}$ and $\frac{y(\omega_G,0)}{3}$, with the total frequency $\sum_{\tilde{y}\in\{0,1\}}[f(\omega_B,0,\tilde{y}) + f(\omega_G,0,\tilde{y})]$ given by $\frac{2}{3}$,

Stage	Worldview	Policy
$k = 3$	$p_3 = \{(\omega_G, 0)\}, \{(\omega_B, 0)\}, \{\{(\omega_G, a), (\omega_B, a)\}\} \equiv \{\alpha_0, \alpha_1, \alpha_2\}$ $\beta_3(\alpha_0) = y(\omega_G, 0)$, $\beta_3(\alpha_1) = y(\omega_B, 0)$ and $\beta_3(\alpha_2) = 1$	0 in state ω_G a in state ω_B
$k = 2$	$p_2 = \{\{(\omega_G, 0), (\omega_B, 0)\}, \{(\omega_G, a), (\omega_B, a)\}\} \equiv \{\alpha_3, \alpha_4\}$ $\beta_2(\alpha_3) = \frac{3y(\omega_G,0)+2y(\omega_B,0)}{5}$ and $\beta_2(\alpha_4) = 1$	0 in state ω_G, 0 in state ω_B
$k = 1$	$p_1 = \{\{(\omega_G, 0), (\omega_G, a), (\omega_B, 0), (\omega_B, a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5) = \frac{3y(\omega_G,0)+2y(\omega_B,0)+1}{6}$	0 in state ω_G, 0 in state ω_B

Table 2: Ergodic worldview-policy profile (blue: differences with respect to the ergodic profile in Table 1).

- lastly, $\beta_2(\alpha_3)$ follows from (1): $\beta_2(\alpha_3) = \frac{y(\omega_G,0)+y(\omega_B,0)}{2}$.

It can then be checked that this worldview-policy profile is optimal, i.e., that at each stage of the cycle and in each state of the world, the policy prescribed by the profile is optimal given the stage-winning worldview.

As announced, efficiency does not increase with complexity along the cycle:

- for $k = 3$, the efficient policy is always selected,
- for $k = 2$, the efficient policy is never selected (the elected politician selects the "wrong" policy in each state),
- for $k = 1$, the efficient policy is selected in state ω_G, but not in state ω_B (hence, half the time).

In fact, efficiency is minimized for $k = 2$.[35]

Let us describe the logic. At stage $k = 2$, the worldview-profile features an "empowering" worldview. It pools $(\omega_B, 0)$ and $(\omega_G, 0)$ in the same class, and so the class conditional probability of success given the status-quo/inaction policy lies between $y(\omega_B, 0)$ and $y(\omega_G, 0)$. The frequencies in the data, induced by the ergodic policy choices, determine the weights on $y(\omega_B, 0)$ and $y(\omega_G, 0)$. Here, the resulting conditional probability of success in the class is (a) sufficiently high to make status quo/inaction the optimal policy (according to the worldview) in state ω_B given the low cost of this policy, but (b) too low to make action (a) optimal in state ω_G given its high cost.

This ergodic worldview-policy profile is the unique pure ergodic worldview-policy profile if $y(\omega_G, a)-c(\omega_G, a) > \frac{3y(\omega_G,0)+2y(\omega_B,0)}{5}$. However, if $y(\omega_G, a)-c(\omega_G, a) \leq \frac{3y(\omega_G,0)+2y(\omega_B,0)}{5}$,

[35]Consistently with our previous observation, efficiency is maximized for $k = 3$, which allows to separate all payoff-relevant states.

Stage	Worldview	Policy
$k=3$	$p_3 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0)\},\{(\omega_1,a),(\omega_2,a)\},\{(\omega_3,a)\}\}$ $\beta_3(\alpha_0)=0$, $\beta_3(\alpha_1)=\frac{2y(\omega_1,a)+y(\omega_2,a)}{3}$ and $\beta_3(\alpha_2)=y(\omega_3,a)$	a in states ω_1,ω_3 0 in state ω_2
$k=2$	$p_2 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0)\},\{(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_2(\alpha_3)=0$ and $\beta_2(\alpha_4)=\frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{4}$	a in states ω_1,ω_2, 0 in state ω_3
$k=1$	$p_1 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0),(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_1(\alpha_5)=\frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{9}$	0 in state $\omega_1,\omega_2,\omega_3$,

Table 3: Ergodic worldview-policy profile when $c(\omega_1,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)}{3} < c(\omega_2,a)$ and $c(\omega_2,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{4} < c(\omega_3,a)$.

there exists a second ergodic worldview-policy profile, described in Table 2. This second worlview-policy profile features the same partitions at all stages as the profile described before. In particular, at stage $k=2$, it pools $(\omega_B,0)$ and $(\omega_G,0)$ in the same class. Yet, here, along the full cycle, policy 0 is chosen more often in state ω_G than in state ω_B. Consequently, the conditional probability of success of the class composed of $(\omega_B,0)$ and $(\omega_G,0)$ is now closer to $y(\omega_G,0)$, sufficiently so to be higher than $y(\omega_G,a)-c(\omega_G,a)$. This induces the elected politician to choose 0 in both states ω_B and ω_G. As a result, in this worldview-policy profile, the policy choice in the first two stages ($k=1,2$) is inefficient with an equal probability (1/2).

To establish the second part of Observation 5 – i.e., that efficiency can first increase then decrease during the phase of rising complexity –, consider the following environment:
Example II: $\Omega = \{\omega_1,\omega_2,\omega_3\}$ with ω_1, ω_2, ω_3 equally likely, and $A=\{0,a\}$. Suppose moreover that

(i) $y(\omega,0)=0$ for all $\omega\in\Omega$, while $0<y(\omega_1,a)<y(\omega_2,a)<y(\omega_3,a)$, and $y(\omega_3,a)-y(\omega_2,a)\gg y(\omega_2,a)-y(\omega_1,a)$, $y(\omega_3,a)-y(\omega_2,a)\gg y(\omega_1,a)$,

(ii) $c(\omega,0)=0$ for all $\omega\in\Omega$, while $0<c(\omega_1,a)<c(\omega_2,a)<c(\omega_3,a)$,

(iii) $y(\omega,a)-c(\omega,a)>0$ for all ω, with $y(\omega_1,a)-c(\omega_1,a)$ and $y(\omega_3,a)-c(\omega_3,a)$ close to 0.

(iv) $c(\omega_1,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)}{3} < c(\omega_2,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{4} < c(\omega_3,a)$.

Lastly, suppose that $k^*=3$.

Table 3 describes an ergodic worldview-policy profile, in which efficiency at stage

$k = k^* = 3$ is worse than at stage $k = 2$.[36] In this example, a more complex worldview allows to separate/isolate an extreme state (ω_3) with high probability of success, which makes the worldview more conservative/more pessimistic regarding the intermediary states. However, welfare gains are low in the extreme state (as costs are then high), while they can be much higher in intermediary states (due to very low costs).

As an illustration, consider climate change, with the recurring states $\omega_1, \omega_2, \omega_3$ correspond to technological breakthroughs: the extreme state ω_3 corresponds to a new, highly effective technology (e.g., to produce energy, or store GHG, etc.) being available, but which involves very large costs to be operated (policy a), while intermediate states (ω_1, ω_2) correspond to more standard, less effective technologies being available, but which can be operated at much lower costs. A more complex worldview then leads to an excessive focus on the possibility of the new, highly effective technology being available, neglecting the efficiency gains that can be achieved when only standard technologies are available.

Lastly, the last part of Observation 5 follows from tweaking Example II: consider the same environment except for (iii), assuming instead that:

(iii') $y(\omega_i, a) - c(\omega_i, a) > 0$ for all $i = 2, 3$, but $y(\omega_1, a) - c(\omega_1, a) < 0$, and $y(\omega_3, a) - c(\omega_3, a)$ close to 0.

It can then be checked that the ergodic worldview-policy profile described in Table 3 still exists, and the expected utility at stage k^* is strictly negative, while it is zero at stage 1.

The role of (un)sophisticated voters. As noted above, sufficiently complex worldviews lead to efficient policy choices. Hence, sufficiently sophisticated voters can lead to efficient policy choices in a given stage if they are pivotal at that stage. Then, in the ergodic state, the efficient policy choices participate in shaping the data set of observations that is available to voters at any stage of the cycle. How does the influence of sophisticated voters (or conversely, the influence of unsophisticated voters) propagate along the cycle?

As for the relation between complexity and efficiency (Observation 5), we show that adding more sophisticated voters (increasing k^*) or "educating" unsophisticated vot-

[36] It is the unique pure ergodic worldview-policy profile if $c(\omega_2, a) > \frac{y(\omega_1,a)+y(\omega_2,a)}{2}$.

ers/reducing their complexity aversion (reducing μ_1 down to zero) may increase or decrease ergodic efficiency depending on the environment. We refer to Appendix E for details and complements.

Observation 1 (The role of sophisticated voters). *A higher k^* can lead to either a higher or a lower ergodic efficiency.*

5.4 Causes of equilibrium multiplicity

There are two potential sources of multiplicity of ergodic worldview-policy profiles:

(i) Multiplicity with same partitions at all stages, but different conditional probabilities of success. Then, the different conditional probabilities of success determine different policy choices, thus different frequencies in the data, themselves consistent with the different conditional probabilities of success. The environment corresponding to the ergodic profiles described in Tables 1 and 2 illutrates this multiplicity.[37]

(ii) Multiplicity with different partitions. This type of multiplicity can arise in particular in environments in which outcomes are state-driven. We illustrate this observation in Appendix H. Intuitively, when $k^* = N_\Omega + 1$, the k^*-optimal worldviews separate some actions for a given state, while all "previous" k-optimal worldviews with $k \leq k^* - 1$ pool all policies together for a given state. Hence, k-optimal worldviews with $k \leq k^* - 1$ lead only to cost-minimizing policies (for certain states), while by contrast, a k^*-optimal worldview may lead to a non-cost-minimizing policy (for a given state). Since such a policy is then the unique non-cost-minimizing policy ever played along the cycle, different (pure) ergodic worldview-policy profiles can coexist, each with a different partition at stage k^* and leading to different non-cost-minimizing policies at that stage.

By contrast, in state-driven environments, when $k^* \leq N_\omega$, all ergodic worldview-policy profiles share the same partitions at all stages of the cycle, and any policy played on cycle with a strictly positive probability is a cost-minimizing policy.

[37] This point is subtler than it may first appear. Indeed, the tremble ε ensures that in an ergodic profile, all individual empirical probabilities of success are correct, i.e., that $\hat{y}(\omega, a) = y(\omega, a)$ for all (ω, a), and so the multiplicity does not arise from arbitrary beliefs "off-path". Instead, the multiplicity is sustained as the class probabilities of success, $\hat{b}(\alpha)$, depend also on the empirical masses, $f(\omega, a, 0) + f(\omega, a, 1)$, of the couples (ω, a) (how often state ω realizes and policy a is chosen in state ω). The multiplicity arises precisely through these empirical masses.

6 Markovian state transitions

While we assume in our baseline model that states of nature are i.i.d. across periods, our analysis can be extended to Markovian state transitions, in which the distribution of states of nature in a given period depends on the previous-period state and policy. This path-dependency of states of nature induces an additional endogeneity in the set of observations available to the agents.[38]

To remove some multiplicity, and in the same spirit as with strategies, we assume that there is a "tremble" in the state transitions, so that for any (ω, a), the probability of reaching $\omega' \in \Omega$ from (ω, a) is strictly positive. We then take the limit case in which both trembles (on state transitions and policies) go to zero. The same arguments as in the proofs of Propositions 1 and 2 apply, yielding the following result.

Proposition 6 (**Markovian state transitions**). *There exists an ergodic worldview-policy profile. Moreover, in any ergodic profile, complexity and participation dynamics are deterministic and exhibit the same cycles as in our baseline model.*

For the sake of concreteness, we provide an illustration, extending our running example by assuming that choosing a certain (known) policy leads to the good state in the next period with probability (almost) 1, while choosing the other ("wrong") policy leads to the bad state. Specifically, consider our running example: $\Omega = \{\omega_G, \omega_B\}$ with ω_G and ω_B equally likely, $A = \{0, a\}$, $k^* = 3$ and assumptions (i)-(iv). Suppose in addition that the current state depends on the previous-period policy as follows (taking the limit in which the tremble on state transitions vanishes):

$$\omega_t = \begin{cases} \omega_G & \text{if } a_{t-1} = a, \\ \omega_B & \text{if } a_{t-1} = 0. \end{cases}$$

Then, there exists a unique pure ergodic worldview-policy profile, described in Table 4.

In this ergodic profile, the good state ω_G is reached only at stage $k = 1$, thanks to the stage-k^* policy: here, the backlash against complexity coincides with an improvement in the state, but due to the previous-period policy choice (inspired by a complex worldview).[39] This timing of events is consistent with evidence that populist backlashes are

[38] Voters and politicians are (still) myopic.

[39] While the elected politician at stage $k = 1$ finds itself in the good state, it chooses the "wrong" policy (as it is the cost-minimizing), and so does the elected politician at the next stage ($k = 2$), inducing the bad state both at stages 2 and 3 of the cycle.

Stage	Worldview	Policy
$k = 3$	$p_3 = \{(\omega_G, 0)\}, \{(\omega_B, 0)\}, \{\{(\omega_G, a), (\omega_B, a)\}\} \equiv \{\alpha_0, \alpha_1, \alpha_2\}$ $\beta_3(\alpha_0) = y(\omega_G, 0)$, $\beta_3(\alpha_1) = y(\omega_B, 0)$ and $\beta_3(\alpha_2) = 1$	0 in state ω_G **$\boldsymbol{a}$ in state $\boldsymbol{\omega}_B$**
$k = 2$	$p_2 = \{\{(\omega_G, 0), (\omega_B, 0)\}, \{(\omega_G, a), (\omega_B, a)\}\} \equiv \{\alpha_3, \alpha_4\}$ $\beta_2(\alpha_3) = \frac{y(\omega_G,0)+y(\omega_B,0)}{2}$ and $\beta_2(\alpha_4) = 1$	a in state ω_G **0 in state $\boldsymbol{\omega}_B$**
$k = 1$	$p_1 = \{\{(\omega_G, 0), (\omega_G, a), (\omega_B, 0), (\omega_B, a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5) = \frac{y(\omega_G,0)+y(\omega_B,0)+1}{3}$	**0 in state $\boldsymbol{\omega}_G$** 0 in state ω_B

Table 4: Ergodic worldview-policy profile (**bold**: realized states and policies).

not necessarily caused/preceded by economic downturns.

7 Intellectuals and politicians

In our baseline model, complexity and participation are tied as voters do not vote when they understand neither the incumbent's nor the entrant's worldviews. We relax this assumption in this Section, assuming instead that when they understand neither the incumbent's nor the entrant's worldviews, voters recall the last worldview they understood, and that voters vote for the politician who promises the highest expected utility, measuring the politician's promises according to the most convincing worldview they understand or recall.[40]

Moreover, while we assume in our baseline model that politicians propose both worldviews and policies, we assume in this Section that worldviews and policies are proposed by different actors: (i) worldviews are first proposed by "intellectuals"; then (ii) policies are proposed by politicians, who take the worldviews as given and commit to a policy function (that is, a policy contingent on the state of nature).

More formally,

- Voters adopt the worldview they find most convincing: namely, and as before, the worldview they understand that best explains the observed data. Voters then vote

[40]Hence, while less sophisticated voters who do not understand current worldviews resort to the last worldview they understood, more sophisticated voters who understand at leat one such worldview opt for a current worldview they understand, regardless of any past worldview they may remember. A microfoundation relies on imperfect memory, assuming that period after period, voters' recall of a given worldview falters, which reduces its plausability, so much so that even a much simpler, but recent worldview performs better in their eyes than a more complex, but older worldview. Alternatively, voters may have a *recency* bias that induces them to underestimate worldviews that are not offered by the current politicians, or voters may face a cost for remembering a worldview (relative to adopting one proposed by a current politician) and they are only willing to incur this cost when no understandable worldview is proposed by the current politicians.

for the politician who offers the highest expected utility as evaluated through the lens of the worldview they have adopted.

- *Worldview competition:* In each period, the intellectual who won the previous-period worldview competition faces a new challenger; the incumbent intellectual cannot change their worldview from the previous period, while the challenger is free to choose any worldview they want. Intellectuals aim firstly at maximizing their audience share (i.e., ratio of mass of voters adopting their worldview over the total mass of voters adopting either their own or the incumbent intellectual's worldview), and secondly, among worldviews that maximize their audience share, at minimizing the Kullback-Leibler divergence between the data and the distribution induced by their worldview.[41]

- *Electoral competition:* After the worldview competition has ended, the incumbent politician (elected in the previous-period) faces a new challenger in the electoral competition. The incumbent cannot change its policy profile from the previous period, while the challenger is free to commit to any policy profile. Politicians (still) aim at maximizing their vote share.

Our first general result is that on path, worldviews still follow the complexity dynamics described in Proposition 1. Indeed, the same arguments as in our baseline model apply, replacing politicians with intellectuals.

Proposition 7 (**Intellectuals and politicians: Worldview dynamics**). *In equilibrium, worldviews' complexity dynamics are deterministic and cyclical:*

(i) A cycle lasts for k^* periods.

(ii) Gradual increases in complexity and backlashes towards maximum simplicity: The complexity of the winning worldview increases gradually until it reaches k^*, at which point it falls down to the minimum level (1), and a new cycle begins.

Our second general result is the existence of ergodic worldview-policy profiles: the deterministic dynamics of worldviews' complexity allow us to use the same arguments as in the proof of Proposition 2.

[41] We rely on the same microfoundations as for politicians in the baseline model (see Section 8.1): voters first select a worldview based on how convincing it seems, but make small mistakes in their assessment (or the intellectuals' explanations are noisy); then, voters actually praise the intellectual whose worldview they selected subject to an exogenous, idiosyncratic sympathy/hostility shock.

Proposition 8 (**Intellectuals and politicians: Ergodic worldview-policy profile**). *There exists an ergodic worldview-policy profile.*

In contrast to our baseline environment, when intellectuals propose worldviews and politicians policies, the winning policy at a given stage can differ from the optimal policy according to the prevailing worldview at that stage. As an illustration, consider our running example, leading to the ergodic worldview-policy profiles in Tables 1-2 when politicians propose both worldviews and policies, and voters do not vote when they understand neither the incumbent's nor the entrant's worldview.

The worldview-policy profiles described in Table 1-2 are no longer ergodic profiles. Indeed, in both profiles, at stage 3, voters who understand only 1-complex worldviews prefer the policy strategy of the incumbent against the challenger's strategy, and thus vote for the incumbent.

By contrast, there exists a pure ergodic worldview-policy profile in which at all stages, the entrant politician offers a policy strategy choosing 0 in both states, winning the election with probability 1/2. Even more strikingly, if $\frac{2y(\omega_G,0)+3y(\omega_B,0)}{5} > y(\omega_B, a) - c(\omega_B, a)$, there exists a pure ergodic worldview-policy profile, described in Table 5, in which the stage-1 and stage-2 entrant politicians choose the policy strategies in Table 1, while the stage-3 entrant politician chooses the same policy strategy as the stage-1 entrant politician. The stage-3 entrant politician thus wins the stage-3 election by catering to the preferences of voters who understand only 1-complex worldviews. Hence, at stage k^*, the "stage policy" (winning policy at a given stage) is not optimal according to the "stage worldview" (winning worldview at a given stage) – in fact, it is suboptimal according to both worldviews competing at that stage (the incumbent intellectual's and the entrant intellectual's). Put differently, at stage k^*, the stage policy features a "simplicity backlash", while the stage worldview is the most complex of the cycle (and gradually more complex than the previous-period stage worldview). In other words, *the simplicity backlash in policies precedes the simplicity backlash in worldviews.*

As a last observation on this ergodic profile, note that at stage $k = 1$, the stage policy coincides with the optimal policy according to the stage worldview: this coincidence is in fact a general result. Indeed, at stage $k = 1$, all voters understand the entrant's 1-optimal worldview, and a majority of voters selects this worldview.

Proposition 9 (**Intellectuals and politicians: Optimal worldviews and (sub)optimal policies**). *While at stage* 1 *of the worldviews' complexity cycle, the win-*

Stage	Worldview	Policy
$k=3$	$p_3 = \{(\omega_G,0)\},\{(\omega_B,0)\},\{\{(\omega_G,a),(\omega_B,a)\}\} \equiv \{\alpha_0,\alpha_1,\alpha_2\}$ $\beta_3(\alpha_0) = y(\omega_G,0)$, $\beta_3(\alpha_1) = y(\omega_B,0)$ and $\beta_3(\alpha_2) = 1$	0 in state ω_G 0 in state ω_B
$k=2$	$p_2 = \{\{(\omega_G,0),(\omega_B,0)\},\{(\omega_G,a),(\omega_B,a)\}\} \equiv \{\alpha_3,\alpha_4\}$ $\beta_2(\alpha_3) = \frac{2y(\omega_G,0)+3y(\omega_B,0)}{5}$ and $\beta_2(\alpha_4) = 1$	a in state ω_G, 0 in state ω_B
$k=1$	$p_1 = \{\{(\omega_G,0),(\omega_G,a),(\omega_B,0),(\omega_B,a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5) = \frac{2y(\omega_G,0)+3y(\omega_B,0)+1}{6}$	0 in state ω_G, 0 in state ω_B

Table 5: Ergodic worldview-policy profile.

ning policy is always the cost-minimizing policy, which is optimal according to the winning simplest worldview, this coincidence may fail at other stages. Then, the winning policy may not be optimal according to the winning worldview. In particular, the simplicity backlash in policies may precede the simplicity backlash in worldviews.

8 Discussion

8.1 The electoral game: Microfoundations

Single entrant. We assume that in each period, there is a single entrant. A microfoundation is that there are large entry costs in the electoral competition (e.g., to run a political campaign), so that the opposition focuses on a single platform/candidate.[42]

For instance, suppose that prior to the main election, a primary – or a form of deliberation/collective agreement – takes place within the opposition to determine its candidate and worldview for the main election. Candidates in the primary compete by offering worldviews, which they will have to keep in the main election if they win the primary.[43] Primary voters understand the (current-period) election game and want the opposition party to win the main election: hence, they vote strategically for the worldview that maximizes the party's vote share in that election. It is then optimal for all candidates in the primary to offer the worldview that maximizes the probability of defeating the incumbent, and thus in equilibrium, the primary winner is chosen based on horizontal

[42]Another microfoundation would be via a "citizen-candidate" model (Besley and Coate, 1997)), in which entry into the electoral competition would take the form of a preemption race among potential challengers (any challenger who has entered is free to adjust his worldview prior to the election after observing who else has entered), and so, with entry costs sufficiently large, a single entrant enters on path.

[43]This consistency condition can be motivated by the same argument as the consistency condition for the incumbent politician (see below).

traits (or "valence").[44]

The entrant's electoral objective. We assume that the entrant has lexicographic preferences on firstly, maximizing its vote share, and secondly, explaining the data. Hence, the entrant chooses a worldview that maximizes its vote share (i.e., the ratio of its own votes over total votes), and among worldviews that maximize its vote share, it chooses a worldview that minimizes the Kullback-Leibler divergence with the distribution observed in the data.

We microfound these preferences as follows. A voter's choice is (a) firstly determined by the politicians' worldviews, yet the voter's understanding of the worldviews is noisy and voters ignore worldviews they find not convincing enough, and (b) secondly, that whether a voter actually votes for the politician offering the worldview the voter finds most convincing (and convincing enough) depends on an idiosyncratic, exogenous "valence" term capturing sympathy (or aversion) to a particular politician – e.g., due to some "horizontal" traits of the candidates (identity, tastes, etc.).

Formally, in each period, the timing in the full-fledged electoral game is as follows:

(i) The entrant politician publicly announces a worldview,

(ii) Each voter then determines the worldview they find most convincing among those they understand (if any), i.e., the worldview that best explains the available data subject to the voter's sophistication constraint. Yet, the transmission and/or the evaluation of worldviews is noisy: letting $\kappa(v, d)$ denote the (correct) Kullback-Leibler divergence of worldview v with respect to data d, the voters' estimation of this divergence (for all voters who understand the worldview) is equal to $\kappa(v, d) + \epsilon_I$ if v is the incumbent's worldview, resp. $\kappa(v, d) + \epsilon_E$ if it is the entrant's, where ϵ_I, ϵ_E are independent, normally distributed, with mean 0.

Moreover, voters choose to ignore worldviews that they believe achieve an exceedingly high Kullback-Leibler divergence: there exists $\overline{k} > 0$ such that a voter who understands a worldview v decides to ignore it (either voting for the other worldview if understandable, or not voting at all) if $\kappa(v, d) + \epsilon > \overline{k}$, where $\overline{k}$ is sufficiently large that for any data d and the corresponding 1-optimal worldview v_1, $\kappa(v_1, d) < \overline{k}$

[44] An alternative microfoundation is thus that the primary is not about specific worldviews (the latter is chosen only later on in the main election), and participants to the primary simply choose their favorite candidate based on horizontal traits.

(and thus the same inequality holds for any k-optimal worldview for all $k \geq 1$).[45]

(iii) For any voter who has a most convincing and convincing enough worldview v, a politician-specific "valence" shock realizes, and influences the voter's participation: voters favoring the incumbent's worldview, resp. the entrant's worldview, actually vote for the incumbent, resp. the entrant, with probability $1 - \varepsilon_I \in [0, 1]$, resp. $1 - \varepsilon_E \in [0, 1]$, where $\varepsilon_I, \varepsilon_E$ are independent, randomly distributed with full support on $[0, 1]$, and otherwise abstain.[46]

The entrant politician's objective is to maximize its probability of being elected in the current period.

Our basic model thus corresponds to the limit of this general environment as both noises/shocks vanish.

The "consistency" constraints. We make two *consistency* assumptions on the (elected) politicians' worldview and policy choices: we assume that (a) the incumbent cannot change its worldview from one period to another, and that (b) when choosing policies, elected politicians stick to the worldviews they promoted during the election.

These two assumptions are in line with the large empirical evidence that voters tend to turn away from politicians who change their stance (see, e.g., Kartik and McAfee (2007) and references therein). Then, any change of worldview from one period to another, or any deviation away from the optimal policies according to the worldview promoted by the politician will be penalized by voters in the next election.[47]

These two assumptions are also in line with the interpretation of the model in which politicians believe faithfully in the worldview they promote (i.e., believe that the worldview they promote corresponds to the true data generating process). Then, the entrant is selected from a pool of politicians, each believing faithfully in a specific worldview, and

[45] This assumption realistically implies that optimal worldviews can still convince voters given the cap $\overline{k}$. It also implies that the entrant (still) wants to select an optimal worldview when she selects a simpler worldview than the incumbent's – otherwise, voters who understand only the entrant's worldview would not react to a worldview being non-optimal, and the entrant would maximize her vote share by picking *any* 1-complex worldview.

[46] The "valence" shock (e.g., due to idiosyncratic sympathy/aversion) affects only voters who have a favorite worldview running in the election. A rationale, consistent with our general stance, is that voters who understand/take into consideration neither the incumbent's nor the entrant's worldview as both are too complex, lose interest in the current-period electoral competition and thus in the politicians' characters.

[47] Note that, except in the limit as all noises vanish, the incumbent has a strictly positive probability of being reelected.

the selection mechanism (e.g., a primary) picks the politician believing in the worldview that is most likely to win the election.[48]

Voters' sophistication and participation. While we refer to voters' heterogeneity in terms of sophistication, we aim at capturing more generally any constraint on a voter's ability to understand politicians' claims – e.g., not only education, but also time and energy currently available, quality of the media they have access to, etc.

We assume that a voter does not participate in the election when both the incumbent's and the entrant's worldview exceed the voter's sophistication. This assumption can be microfounded in several ways. For instance, suppose that a voter who understands neither the incumbent's nor the entrant's worldviews is indifferent between the two, and so, if there is a (possibly very small) positive cost of voting, strictly prefers abstaining.[49]

By contrast, when voters understand at least one of the competing worldviews, they vote for the most convincing worldview among those they understand, ruling out the worldview they do not understand (if any). A possible microfoundation may be that voters are maximally suspicious of any worldview they do not understand.[50]

8.2 Deterministic asymmetric cycles: A comparison to models in industrial organizations

While our cycles arise from worldview competition with heterogeneous voter sophistication, the asymmetric pattern of gradual change punctuated by sudden reversals appears in the IO literature through entirely different mechanisms. Notably, Maskin and Tirole (1988) characterize Edgeworth cycles in duopoly pricing where firms alternate setting prices in staggered periods. Prices decrease gradually until reaching the static Nash level, at which point one firm discontinuously raises its price well above this level, trig-

[48]Besides, in our environment, a forward-looking politician would have no incentive to bias its current-period policy choice to influence the next-period election. Indeed, since at any time $t \in (-\infty, +\infty)$, voters have access to an infinitely large data set and thus empirical frequencies do not change from one period to the other, a forward-looking politician could not change the empirical frequencies in the next period (and thus its probability of reelection) via its choice of policy.

[49]If there is no cost of voting, the voter may vote at random for one or the other candidate, or following a small idiosyncratic "sympathy shock" for one or the other candidate. With our general microfoundations, the entrant politician aims at maximizing its probability of being elected in the current period, and thus such random votes would lead to the same electoral dynamics as our main assumption, except for the constant participation/turnout.

[50]Such maximal suspicion may arise, e.g. from motivated beliefs and in particular, "self-esteem maintenance" (choosing to believe that any worldview they do not understand must be flawed/bad to preserve their self-esteem).

gering the other firm to follow and restart the cycle. Though driven by different forces – price competition rather than worldview complexity –, their dynamics share the same structural feature of slow drifts followed by large jumps as our cycles.[51]

Indeed, in these and following works as in ours, such asymmetric cycles arise due to two key properties of the environment: (i) a single player moves in each period (in our model, this stems from the structure of the political contest: a single entrant/challenger faces the incumbent who is constrained to stick to the worldview he promoted when he got elected); (ii) the "action space" is discrete (which in our model stems from the structure of worldviews and their finite number of analogy classes).

9 Conclusion

This paper develops a theory of political cycles driven by the interplay between worldview complexity and voter sophistication. We show that when voters assess competing worldviews based on their empirical plausibility, but differ in their ability to understand complex explanations, electoral competition generates deterministic cycles with a distinctive asymmetric structure: worldview complexity rises gradually, then collapses suddenly to the simplest possible worldview (a *simplicity backlash*). The cycles are accompanied by declining participation as complexity increases, followed by surges in participation when simplified worldviews bring all voters back into the political process.

We derive several key insights about the relationship between worldview complexity and policy outcomes. First, we identify conditions under which political cycles systematically favor short-termism: when the simplest worldview wins, cost-minimizing policies are implemented regardless of long-term benefits; when outcomes vary more with states than policies, short-termist policies can persist even at higher complexity levels; and when policy spaces are dense with many intermediate options, the worldviews' local clustering of "close" policies, and the policy-makers' (local) cost-minimization among "close" policies,

[51]Pesendorfer (1995) analyzes fashion cycles driven by signaling in matching markets, where a monopolist designer periodically creates new designs that consumers use as signals in a "dating game". Initially, only high-type consumers purchase the new design at premium prices to signal their type. Over time, prices fall and the design spreads across the population as more consumers adopt it. Once sufficiently many consumers own the design, its signaling value is destroyed, making it profitable for the monopolist to introduce a new design that renders the old one obsolete, restarting the cycle. While the mechanism differs entirely from ours, the dynamics share an same asymmetric pattern, albeit reversed: gradual diffusion as designs spread and prices fall (vs gradual narrowing of participation as complexity increases in our framework), followed by sudden obsolescence and narrower adoption when new designs are introduced (vs surge in participation when complexity drops).

create an unraveling leading to *global* cost-minimization. Second, we demonstrate that the relationship between worldview complexity and policy efficiency is non-monotonic: Intermediate complexity levels can generate worse outcomes than either simple or highly complex worldviews, as moderately complex worldviews may cluster state-policy pairs in ways that lead to perverse policy choices. Third, we show that increasing voter sophistication has ambiguous welfare effects: While more sophisticated voters can break persistent short-termism by supporting complex worldviews, they also lengthen cycles and can shift the distribution of policies in efficiency-reducing directions.

Our framework extends naturally to environments with path-dependent states and to settings where intellectual and political competition are separated. When states evolve according to past policies, the core cycle dynamics persist but new patterns can arise: We provide an illustration in which populist (simplistic) backlashes coincide with improved conditions due to prior technocratic (complex) policies. When intellectuals supply worldviews while politicians compete on policy platforms, policy backlashes can precede worldview backlashes, with politicians catering to less sophisticated voters even as the dominant intellectual discourse remains complex.

Taken together, our insights suggest that the cyclical dynamics between (technocratic) complexity and (populist) simplicity represent an inherent feature of democratic competition when voters face heterogeneous sophistication constraints.

References

Aina, C. (2024). "Tailored Stories". Working paper.

Banerjee, A., Merugu, S., Dhillon, I. S., and Ghosh, J. (2005). "Clustering with Bregman Divergences". *Journal of Machine Learning Research*, 6:1705–1749.

Besley, T. and Coate, S. (1997). "An Economic Model of Representative Democracy". *Quarterly Journal of Economics*, 112(1):85–114.

Condorcet, N. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie royale, Paris.

Dewatripont, M. and Tirole, J. (1999). "Advocates". *Journal of Political Economy*, 107(1):1–39.

Eliaz, K., Galperti, S., and Spiegler, R. (2024). "False Narratives and Political Mobilization". *Journal of the European Economic Association*, page jvae047.

Eliaz, K. and Spiegler, R. (2020). "A Model of Competing Narratives". *American Economic Review*, 110(12).

Eliaz, K. and Spiegler, R. (2025). "News Media as Suppliers of Narratives (and Information)". Working paper.

IPCC (2022). *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. [P.R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, J. Malley, (eds.)] Cambridge University Press, Cambridge, UK and New York, NY, USA.

Izzo, F., Martin, G. J., and Callander, S. (2023). "Ideological Competition". *American Journal of Political Science*, 67(3):687–700.

Jehiel, P. (2005). "Analogy-Based Expectation Equilibrium". *Journal of Economic Theory*, 123(2):81–104.

Jehiel, P. (2022). "Analogy-Based Expectation Equilibrium and Related Concepts: Theory, Applications, and Beyond". Working paper.

Jehiel, P. and Weber, G. (2024). "Endogenous Clustering and Analogy-Based Expectation Equilibrium". Working paper.

Kartik, N. and McAfee, R. P. (2007). "Signaling Character in Electoral Competition". *American Economic Review*, 97(3).

Levy, G. and Razin, R. (2025). "Political Social Learning: Short-Term Memory and Cycles of Polarization". *American Economic Review*, 115(2).

Levy, G., Razin, R., and Young, A. (2022). "Misspecified Politics and the Recurrence of Populism". *American Economic Review*, 112(3).

Mailath, G. J. and Samuelson, L. (2020). "Learning under Diverse World Views: Model-Based Inference". *American Economic Review*, 110(5):1464–1501.

Maskin, E. and Tirole, J. (1988). "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles". *Econometrica*, 56(3):571–599.

Montiel Olea, J. L., Ortoleva, P., Pai, M. M., and Prat, A. (2022). "Competing Models". *Quarterly Journal of Economics*, 137(4):2419–2457.

Montiel Olea, J. L. and Prat, A. (2025). "Competing Ideologies: Fit, Simplicity and Fear". Working paper.

Oprea, R. (2020). "What Makes a Rule Complex?". *American Economic Review*, 110(12).

Oprea, R. (2024). "Decisions under Risk Are Decisions under Complexity". *American Economic Review*, 114(12).

Pesendorfer, W. (1995). "Design Innovation and Fashion Cycles". *American Economic Review*, 85(4):771–792.

Rodrik, D. (2011). *The Globalization Paradox: Democracy and the Future of the World Economy.* W.W. Norton, New York and London.

Schwartzstein, J. and Sunderam, A. (2021). "Using Models to Persuade". *American Economic Review*, 111(1):276–323.

Schwartzstein, J. and Sunderam, A. (2025). "Sharing Models to Interpret Data". Working paper.

Spiegler, R. (2016). "Bayesian Networks and Boundedly Rational Expectations". *Quarterly Journal of Economics*, 131(3):1243–1290.

Szalay, D. (2005). "The Economics of Clear Advice and Extreme Options". *Review of Economic Studies*, 72(4):1173–1198.

Appendix

A Proof of Lemma 1

The first part of the result obtains by making explicit the expression of the Kullback-Leibler divergence between the distribution observed in the data – where the probability of an outcome $(\omega, a, \tilde{y})$ is its frequency $f(\omega, a, \tilde{y})$ –, and the distribution induced by the worldview given the data – where the probability of an outcome $(\omega, a, \tilde{y})$ is the product of the frequency of (ω, a) in the data, times the conditional probability of observing $\tilde{y}$ given (ω, a) according to the worldview ($\beta(\alpha(\omega, a))$ if $\tilde{y} = 1$, resp. $1 - \beta(\alpha(\omega, a))$ if $\tilde{y} = 0$). The second part of the lemma, i.e., the expression for the conditional probabilities of success ($\beta(\alpha)$), follows from the first-order condition – which is both necessary and sufficient.[52]

B Proof of Lemma 2

For any $k \leq N - 1$, any $(k+1)$-optimal worldview achieves a strictly lower Kullback-Leibler divergence than any k-optimal worldview. Consider an optimal k-complex worldview, with partition p, and consider the $(k+1)$-complex worldview obtained by (i) splitting in two an analogy class $\alpha \in p$ with at least two elements (ω, a), (ω', a') with $\hat{y}(\omega, a) \neq \hat{y}(\omega, a)$,[53] isolating the couple (ω, a) in a new class (thus with a single element), and (ii) computing the conditional probabilities of sucess as the mean frequencies of successes given the new partition via (1). As $f(\omega, a, 0) + f(\omega, a, 1) > 0$ for any $(\omega, a, \tilde{y})$, Lemma 1 implies that the new worldview achieves a strictly lower Kullback-Leibler divergence.

Voters vote for the most complex worldview they can understand, if any, and does not vote otherwise (if both the incumbent's and the entrant's worldview are too complex). Hence, as noted in the text, given two competing worldviews respectively k-optimal and k'-optimal with $k < k'$, the k-optimal worldview receives $\mu_k + ... + \mu_{k'-1}$ votes, while the

[52]For any α, the politician chooses $\beta(\alpha)$ that minimizes

$$-\left(\sum_{(\omega,a)\in\alpha} f(\omega, a, 0)\right) \ln(1 - \beta(\alpha)) - \left(\sum_{(\omega,a)\in\alpha} f(\omega, a, 1)\right) \ln \beta(\alpha),$$

which is a strictly convex function of $\beta(\alpha) \in (0, 1)$, with a unique interior minimum.

[53]There exists at least one such class as we assume that there is sufficient variation in the data that $\hat{y}(\omega, a)$ takes at least N different values as (ω, a) spans $\Omega \times A$.

k'-optimal worldview receives $\mu_{k'} + \mu_{k'+1} + ... + \mu_N$ votes (and a mass $\mu_1 + ... + \mu_{k-1}$ of voters does not vote). Similarly, when the entrant faces a k-complex worldview that is not k-optimal, then she maximizes her vote share by choosing a k-optimal worldview (achieving a vote share of 1).

The result follows as the entrant chooses a worldview that maximizes her vote share, and among such worldviews, a worldview that minimizes the Kullback-Leibler divergence with the data.

C Proof of Proposition 2

Consider a sequence of worldview-policy profiles, $(\boldsymbol{\nu}_n, \boldsymbol{\sigma}_n)_n$, converging to a worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$. The tremble implies that for any index n, and any stage k, state of the world ω and worldview v, $\sigma_n^k(\omega|v)$ and $\sigma^k(\omega|v)$ have full support over A.

For any n, we denote by $d_n \equiv (f_n(\omega, a, y))_{(\omega,a,y)}$ the ergodic data set induced by the worldview-policy profile $(\boldsymbol{\nu}_n, \boldsymbol{\sigma}_n)$. From the previous step, each dataset d_n has strictly positive frequencies for all (ω, a), uniformly bounded away from zero by the tremble: $f_n(\omega, a, 0) + f_n(\omega, a, 1) \geq \varepsilon > 0$. The sequence of datasets $(d_n)_n$ thus converges to some dataset $d = (f(\omega, a, y))_{(\omega,a,y)}$, with strictly positive frequencies for all (ω, a).

Let $\mathcal{P}_n^k$, resp. $\mathcal{P}^k$ denote the set of k-optimal worldviews given data d_n, resp. given data d. Consider a sequence of worldviews $(v_n)_n = (p_n, \beta_n)_n$ such that $v_n \in \mathcal{P}_n^k$. By Lemma 1, $\beta_n(\alpha)$ is given by (1) for any $\alpha \in p_n$, as $p_n \in \mathcal{P}_n^k$. To ensure that all β_n are defined on the same set $2^{\Omega \times A}$, we let for any index n and class $\alpha \subseteq \Omega \times A$, $\beta_n(\alpha) = 0$ whenever $\alpha \notin p_n$.

Suppose that the sequence $(v_n)_n$ converges to $v = (p, \beta)$. Hence, there exists m such that for any $n \geq m$, $p_n = p$.[54] Moreover, by Lemma 1, for any worldview $v_n \in \mathcal{P}_n^k$, the conditional probability of success $\beta_n(\alpha_n)$ for any class $\alpha_n \in p_n$ is given by (1), which is a continuous function of the frequencies $f_n(\omega, a, y)$ in data d_n. Hence, for any $\alpha \in p$, $\beta(\alpha)$ satisfies (1) with the frequencies $f(\omega, a, y)$ in d. By Lemma 1 again and equation (2), the KL divergence of any worldview with partition p and conditional probability of success given by (1) is a continuous function of the frequencies $f_n(\omega, a, y)$. By construction, for all n, the partition p_n achieves the maximum in (2) given the frequencies $f_n(\omega, a, y)$ in d_n. Therefore, since for all $n \geq m$, $p_n = p$, the partition p achieves the maximum in (2)

[54]Note that unless there exists m such that for any $n \geq m$, $d_n = d_m$, the sequence $(\beta_n)_n$ remains non-constant after any n.

given the frequencies $f(\omega, a, y)$. Hence, $v \in \mathcal{P}^k$.

As for all $n \geq m$, $p_n = p$ and $\beta_n(\alpha)$ converges to $\beta(\alpha)$ for all $\alpha \in p$, there exists $m' \geq m$ such that for any $n \geq m'$ and any $\omega \in \Omega$, for any sequence of policies $(a_n)_{n \geq m'}$ such that for all $n \geq m'$, $a_n \in \arg\max_a \beta_n(\alpha_n(\omega, a)) - c(\omega, a)$ (i.e., policy a_n is optimal in state ω according to worldview v_n), and $\lim_n a_n = a^*$, policy a^* is optimal in state ω according to worldview v:

$$a \in \arg\max_a \beta(\alpha(\omega, a)) - c(\omega, a).$$

For any worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$, inducing an ergodic data set $d_{(\boldsymbol{\nu}, \boldsymbol{\sigma})}$, we denote by $\Gamma(\boldsymbol{\nu}, \boldsymbol{\sigma})$ the set of optimal worldview-policy profiles given data $d_{(\boldsymbol{\nu}, \boldsymbol{\sigma})}$. Consider the correspondence $\tilde{\Gamma}$ that maps any worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$ to the convex hull of $\Gamma(\boldsymbol{\nu}, \boldsymbol{\sigma})$. By the previous steps, the correspondence $\tilde{\Gamma}$ has a closed graph.

Moreover, for any worldview-policy profile $(\boldsymbol{\nu}, \boldsymbol{\sigma})$, $\tilde{\Gamma}(\boldsymbol{\nu}, \boldsymbol{\sigma})$ is non-empty and convex. Kakutani's fixed-point theorem thus implies that $\tilde{\Gamma}$ admits a fixed point, which concludes the proof.

D Proof of Proposition 4

For a given (pure or mixed) ergodic worldview-policy profile, we denote the set of policies played with strictly positive probability (on top of the tremble) by $\{a_{\phi(1)}, ..., a_{\phi(m)}\}$ for some $m \geq 1$, where $\phi(i) < \phi(i+1)$ for all $1 \leq i \leq m-1$. To alleviate the notation, we drop the dependence on the (single) state ω, letting $y(a) \equiv y(\omega, a)$ and $c(a) \equiv c(\omega, a)$.

Let us suppose by contradiction that there exists $\rho > 0$ and a sequence $(\delta_n)_{n \geq 0}$, with $\delta_n \in (0, 1)$ for all $n \geq 0$ and $\lim \delta_n = 0$, such that there exists a sequence of corresponding ergodic worldview-policy profiles (one for each δ_n), where for each $n \geq 0$, the set of policies played with strictly positive probability, denoted by $\{a_{\phi_n(1)}, ..., a_{\phi_n(m_n)}\}$, satisfies

$$y(a_{\phi_n(m_n)}) - c(a_{\phi_n(m_n)}) \geq \rho.$$

Consider a given point in the sequence. As $a_{\phi_n(m_n)}$ is played with strictly positive probability, there exists a stage k_n of the complexity cycle at which there exists a k_n-optimal worldview, v_n, chosen with strictly positive probability by the entrant, such that its partition includes the class $\{a_{\phi_n(m_n)}, a_{\phi_n(m_n)+1}, ..., a_N\}$,[55] and such that choosing $a_{\phi_n(m_n)}$ maximizes efficiency according to worldview v_n. Since $a_{\phi_n(m_n)}$ is the highest action played with a strictly positive probability along the cycle,

$$\begin{cases} \beta\Big(\{a_{\phi_n(m_n)}, a_{\phi_n(m_n)+1}, ..., a_N\}\Big) = y(a_{\phi_n(m_n)}), \\ \hat{m}\Big(\{a_{\phi_n(m_n)}, ..., a_N\}\Big) = \hat{m}\Big(a_{\phi_n(m_n)}\Big), \end{cases}$$

and so $y(a_{\phi_n(m_n)}) - c(a_{\phi_n(m_n)})$ is the maximum efficiency achieved by any policy a according to worldview v_n.

Let $\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}$ denote the second-highest class in that worldview's partition. A necessary condition for worldview v_n to be k_n-optimal is that fixing all classes but the highest two, the partition with $\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}$, $\{a_{\phi_n(m_n)}, a_{\phi_n(m_n)+1}, ..., a_N\}$ achieves a lower KL-divergence than the partition $\{a_{l_n}, ..., a_{l'-1}\}$, $\{a_{l'}, ..., a_N\}$ for any l' with $l_n+1 \leq l' \leq \phi_n(m_n) - 1$ (updating accordingly the conditional probabilities of success β with (1)).[56]

Suppose by contradiction that the second-highest class does not "shrink" as n goes to $+\infty$ (and thus as δ_n goes to zero), i.e., suppose by contradiction that there exists a subsequence (which we re-index to alleviate the notation) such that

$$\exists \rho' > 0 \mid \forall n \geq 0, \ y(a_{\phi_n(m_n)}) - y(a_{l_n}) \geq \rho'.$$

Hence in particular, $\phi_n(m_n) - l_n$ goes to $+\infty$ as n goes to $+\infty$. Moreover,

$$\exists \rho'' > 0 \mid \forall n \geq 0, \ c(a_{\phi_n(m_n)}) - c(a_{l_n}) \geq \rho''.$$

Let us show that the difference between the probability of success of the second-highest class, $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\})$, where $\beta(\cdot)$ is computed via v_n according to (1), and $y(a_{\phi_n(m_n)})$ goes to zero as n goes to $+\infty$. Consider policy $a_{\phi_n(m_n)-1}$. Since the worldview v_n is k_n-optimal, clustering $a_{\phi_n(m_n)-1}$ with the policies $\{a_{l_n}, ..., a_{\phi_n(m_n)-2}\}$ yields a (weakly) lower KL divergence than clustering $a_{\phi_n(m_n)-1}$ with the policies $\{a_{\phi_n(m_n)}, ..., a_N\}$. As

[55] This step follows from the same logic as before: for $a_{\phi_n(m_n)}$ to be played with strictly positive probability according to a given worldview, it must be the cost-minimizing policy within a class.

[56] If $l_n = \phi_n(m_n) - 1$, then we consider the third- and second-highest class (instead of the second- and first-highest), and the necessary condition becomes that the partition with $\{a_{l_n}, ..., a_{\phi_n(m_n)-2}\}$, $\{a_{\phi_n(m_n)-1}\}$ achieves a lower KL-divergence than the partition $\{a_{l_n}, ..., a_{l'-1}\}$, $\{a_{l'}, ..., a_{\phi_n(m_n)-1}\}$ for any l' with $l_n + 1 \leq l' \leq \phi_n(m_n) - 2$. The argument goes through (iterating if the third-highest class is also a singleton, etc.)

$y(a_{\phi_n(m_n)}) - y(a_{\phi_n(m_n)-1}) = \delta_n$ goes to zero as n goes to $+\infty$, our characterization of optimal partitions (see Appendix I) implies that either:

(i) $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)})$ converges to zero as n goes to $+\infty$ (by observation (i), "clustering couples with similar empirical probabilities of success"), or

(ii) the ratio $\hat{m}(\{a_{l_n}, ..., a_{\phi_n(m_n)-2}\}) \big/ \hat{m}(a_{\phi_n(m_n)})$ goes to zero sufficiently fast as n goes to $+\infty$, while the ratio $\hat{m}(a_{\phi_n(m_n)-1}) \big/ \hat{m}(a_{\phi_n(m_n)})$ remains bounded away from zero (by observation (iii), "no class with zero mass/balancing mass across classes").

However, if the ratio $\hat{m}(\{a_{l_n}, ..., a_{\phi_n(m_n)-2}\}) \big/ \hat{m}(a_{\phi_n(m_n)})$ goes to zero while the ratio $\hat{m}(a_{\phi_n(m_n)-1}) \big/ \hat{m}(a_{\phi_n(m_n)})$ remains bounded away from zero, then, as the class probabilities of success in an optimal partition are computed according to (1), this implies that $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)-1})$ converges to zero as n goes to $+\infty$. Therefore, in both cases, $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)})$ converges to zero as n goes to $+\infty$.[57]

As a consequence, the worldview v_n attributes to policy a_{l_n} an efficiency equal to

$$\begin{aligned}&\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - c(a_{l_n})\\ &\quad= y(a_{\phi_n(m_n)}) - c(a_{\phi_n(m_n)}) + \Big(c(a_{\phi_n(m_n)}) - c(a_{l_n})\Big) + \Big(\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)})\Big),\end{aligned}$$

where $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)})$ converges to zero, while $c(a_{\phi_n(m_n)}) - c(a_{l_n})$ remains bounded away from zero as n goes to $+\infty$. Hence, for n sufficiently large, according to the worldview v_n, the policy a_{l_n} achieves a strictly higher efficiency than the policy $a_{\phi_n(m_n)}$, and therefore the latter cannot be played with strictly positive probability, a contradiction.

Therefore, the second-highest class "shrinks" as n goes to $+\infty$ (and thus as δ_n goes to zero): $y(a_{\phi_n(m_n)}) - y(a_{l_n})$ goes to zero as n goes to $+\infty$. We then iterate the argument, considering the third-highest class and showing with the same logic that it shrinks as n goes to $+\infty$. As a consequence, since any partition has at most k^* classes (and k^* does not depend on δ_n), $y(a_{\phi_n(m_n)})$ goes to zero as n goes to $+\infty$.

[57] Note that this step suffices to reach a contradiction when there exists a subsequence with $m_n = 2$ (e.g., with $k_n^* = 2$ and pure ergodic profiles), as then $\beta(\{a_{l_n}, ..., a_{\phi_n(m_n)-1}\}) = y(a_0) = \underline{y}$ for all $n \geq 0$. This step also suffices to reach a contradiction when there exists a subsequence with $k_n = 2$ (with pure or mixed ergodic profiles), as policy a_0 is always played at stage 1 of the cycle and thus $\hat{m}(a_0) \geq 1/k_n^*$, which prevents $\beta(\{a_0, ..., a_{\phi_n(m_n)-1}\}) - y(a_{\phi_n(m_n)})$ from converging to zero.

Stage	Worldview	Policy
$k = 3$	$p_3 = \{\{a_1\}, \{a_2\}, \{a_3, a_4\}\}$ $\beta_3(\{a_1\}) = y(a_1)$, $\beta_3(\{a_2\}) = y(a_2)$ and $\beta_3(\{a_3, a_4\}) = y(a_3)$	a_3
$k = 2$	$p_2 = \{\{a_1, a_2\}, \{a_3, a_4\}\}$ $\beta_2(\{a_1, a_2\}) = y(a_1)$ and $\beta_2(\{a_3, a_4\}) = y(a_3)$	a_3
$k = 1$	$p_1 = \{\{a_1, a_2, a_3, a_4\}\}$ $\beta_1(\{a_1, a_2, a_3, a_4\}) = \frac{y(a_1)+2y(a_3)}{3}$	a_1

Table 6: Ergodic worldview-profile when $k^* = 3$.

Stage	Worldview	Policy
$k = 3$	$p_3 = \{\{a_1\}, \{a_2\}, \{a_3, a_4\}\}$ $\beta_3(\{a_1\}) = y(a_1)$, $\beta_3(\{a_2\}) = y(a_2)$ and $\beta_3(\{a_3, a_4\}) = \frac{y(a_3)+y(a_4)}{2}$	a_1
$k = 2$	$p_2 = \{\{a_1, a_2\}, \{a_3, a_4\}\}$ $\beta_2(\{a_1, a_2\}) = y(a_1)$ and $\beta_2(\{a_3, a_4\}) = \frac{y(a_3)+y(a_4)}{2}$	a_1
$k = 1$	$p_1 = \{\{a_1, a_2, a_3, a_4\}\}$ $\beta_1(\{a_1, a_2, a_3, a_4\}) = y(a_1)$	a_1

Table 7: Ergodic worldview-profile when $k^* = 3$.

E Complements on efficiency and the role of (un)sophisticated voters

Let us begin with an example in which raising k^* leads to a higher ergodic efficiency. Consider the following example: $\Omega = \{\omega\}$ (we will omit the dependence on ω in $y(\cdot)$ and $c(\cdot)$ to alleviate the notation), $A = \{a_1, a_2, a_3, a_4\}$ such that

(i) $c(a_1) < c(a_2) < c(a_3) < c(a_4)$,

(ii) $y(a_1) < y(a_2) < y(a_4) < y(a_3)$ with $\frac{y(a_3)-y(a_4)}{y(a_4)-y(a_2)} \ll \frac{y(a_2)-y(a_1)}{y(a_4)-y(a_2)} \ll 1$.

(iii) $y(a_3) - c(a_3) > y(a_1) - c(a_1) > \max_{i \in \{2,4\}} y(a_i) - c(a_i)$, so that a_3 is the efficient policy.

(iv) $y(a_1) - c(a_1) > \frac{y(a_3)+y(a_4)}{2} - c(a_3)$.

Suppose first that $k^* = 3$. Tables 6 & 7 then describe the two pure ergodic worldview-policy profiles. While the profile in Table 6 selects the efficient policy (a_3) at stages 2 and 3, the profile in Table 7 nevers selects it, yielding instead the cost-minimizing, yet inefficient policy (a_1) at stages 2 and 3.

Suppose now by contrast that $k^* = 4$. Table 8 then describes *the unique* pure ergodic worldview-policy profile. Strikingly, this ergodic worldview-policy profile with $k^* = 4$ has the same worldview partitions for $k = 1, 2, 3$ than the profiles in Tables 6 & 7. However,

Stage	Worldview	Policy
$k=4$	$p_4 = \{\{a_1\},\{a_2\},\{a_3\},\{a_4\}\}$ $\beta_3(\{a_1\}) = y(a_1)$, $\beta_3(\{a_2\}) = y(a_2)$, $\beta_3(\{a_3\}) = y(a_3)$ and $\beta_3(\{a_4\}) = y(a_4)$	a_3
$k=3$	$p_3 = \{\{a_1\},\{a_2\},\{a_3,a_4\}\}$ $\beta_3(\{a_1\}) = y(a_1)$, $\beta_3(\{a_2\}) = y(a_2)$ and $\beta_3(\{a_3,a_4\}) = y(a_3)$	a_3
$k=2$	$p_2 = \{\{a_1,a_2\},\{a_3,a_4\}\}$ $\beta_2(\{a_1,a_2\}) = y(a_1)$ and $\beta_2(\{a_3,a_4\}) = y(a_3)$	a_3
$k=1$	$p_1 = \{\{a_1,a_2,a_3,a_4\}\}$ $\beta_1(\{a_1,a_2,a_3,a_4\}) = \frac{y(a_1)+3y(a_3)}{4}$	a_1

Table 8: Ergodic worldview-profile when $k^* = 4$.

the conditional probabilities of success differ from the "inefficient" profile of Table 6. Intuitively, sufficiently sophisticated voters (with $k = 4$) generate observations of the efficient policy (a_3, which they choose when they are pivotal, at stage $k = 4$). This choice in turn influences less complex worldviews (via their conditional probabilities of success), and thus improves the policy choices of simpler worldviews (except for the simplest one, $k = 1$). Indeed, the efficient policy is now chosen at all stages of the cycle, except at stage 1 ($k = 1$). In other words, the arrival of sufficiently sophisticated voters yield the "unique implementation" of the efficient policy at stages 2 and 3, along with adding a new period in the cycle ($k = 4$) at which the efficient policy is also selected. Hence, the efficient policy is selected 3/4 of the time, instead of 2/3 of the time in the profile of Table 6 and never in the one of Table 7.

Observation 2 (**Sophisticated voters, history and efficiency**). *Sufficiently sophisticated voters can lead to higher efficiency at all but the first ($k = 1$) stage of the cycle.*

Let us consider the other end of the distribution of voters: Suppose that investing in public education, or in the quality of the media improves voters' sophistication, so that $\mu_1 = 0$, while k^* remains equal to 3. Then, if $\frac{y(a_1)+y(a_2)}{2} - c(a_1) < y(a_3) - c(a_3)$ (which is allowed by (i)-(iv)), the worldview-policy profile described in Table 9 is an ergodic profile. The logic is the converse of the one of adding sophisticated voters: now removing observations of the inefficient policy (a_1) changes the conditional probabilities of success, favoring the adoption of the efficient policy (a_3). Notwithstanding, there also exists another ergodic worldview-policy profile, in which policy a_1 is chosen at all stages ($k = 2$ and $k = 3$), yielding the same expected efficiency along the cycle as the profile in Table 7, and a strictly lower expected efficiency tha the profile in Table 6 (both with $\mu_1 > 0$).

Stage	Worldview	Policy
$k=3$	$p_3 = \{\{a_1\},\{a_2\},\{a_3,a_4\}\}$ $\beta_3(\{a_1\}) = y(a_1)$, $\beta_3(\{a_2\}) = y(a_2)$ and $\beta_3(\{a_3,a_4\}) = y(a_3)$	a_3
$k=2$	$p_2 = \{\{a_1,a_2\},\{a_3,a_4\}\}$ $\beta_2(\{a_1,a_2\}) = \frac{y(a_1)+y(a_2}{2}$ and $\beta_2(\{a_3,a_4\}) = y(a_3)$	a_3

Table 9: Ergodic worldview-profile when $k^* = 3$, and $\mu_1 = 0$.

Remark. The simplicity backlash goes "down to" $k = 1$ whenever $\mu_1 > 0$, even if μ_1 is arbitrarily small. Hence, from a practical perspective, improving education and/or media quality to avoid the stage $k = 1$ of the cycle may not be feasible. However, raising k^* (from 3 to 4) may be achieved by increasing μ_4 or decreasing μ_1, and can generate efficiency gains.

Nonetheless, a higher k^* can also generate a lower ergodic efficiency. To establish the second part of Observation 1, we reconsider the environment introduced in Section 5.3 : $\Omega = \{\omega_1, \omega_2, \omega_3\}$ with ω_1, ω_2, ω_3 equally likely, and $A = \{0, a\}$. Suppose moreover that

(i) $y(\omega, 0) = 0$ for all $\omega \in \Omega$, while $0 < y(\omega_1, a) < y(\omega_2, a) < y(\omega_3, a)$, and $y(\omega_3, a) - y(\omega_2, a) \gg y(\omega_2, a) - y(\omega_1, a)$, $y(\omega_3, a) - y(\omega_2, a) \gg y(\omega_1, a)$,

(ii) $c(\omega, 0) = 0$ for all $\omega \in \Omega$, while $0 < c(\omega_1, a) < c(\omega_2, a) < c(\omega_3, a)$,

(iii) $y(\omega, a) - c(\omega, a) > 0$ for all ω, with $y(\omega_1, a) - c(\omega_1, a)$ and $y(\omega_3, a) - c(\omega_3, a)$ close to 0.

(iv) $c(\omega_1, a) < \frac{2y(\omega_1,a)+y(\omega_2,a)}{3} < c(\omega_2, a) < \frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{4} < c(\omega_3, a)$.

We show in Section 5.3 that, when $k^* = 3$, there exists an ergodic worldview-policy profile is such that efficiency is higher at stage $k = 2$ than at stage $k^* = 3$ (see Table 3).

Let us now suppose that $k^* = 2$. Table 10 describes the unique (pure) ergodic worldview-policy profile when, in addition to the above assumptions, $c(\omega_2, a) < \frac{y(\omega_1,a)+y(\omega_2,a)}{2}$ and $\frac{y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{3} < c(\omega_3, a)$. The partitions and policies at stage $k = 1$ and $k = 2$ are the same in this profile (with $k^* = 2$) as those in the profile described in Table 3 (with $k^* = 3$).

The ergodic efficiency in this worldview-policy profile with $k^* = 2$ is thus strictly higher than the ergodic efficiency in the worldview-policy profile with $k^* = 3$ (described

Stage	Worldview	Policy
$k=2$	$p_2 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0)\},\{(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_2(\alpha_3)=0$ and $\beta_2(\alpha_4)=\frac{y(\omega_1,a)+y(\omega_2,a)}{2}$	a in states ω_1,ω_2, 0 in state ω_3
$k=1$	$p_1 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0),(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_1(\alpha_5)=\frac{y(\omega_1,a)+y(\omega_2,a)}{6}$	0 in state $\omega_1,\omega_2,\omega_3$,

Table 10: Ergodic worldview-policy profile when $c(\omega_2,a) < \frac{y(\omega_1,a)+y(\omega_2,a)}{2}$ and $\frac{y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{3} < c(\omega_3,a)$.

Stage	Worldview	Policy
$k=3$	$p_3 = \{\{(\omega_1,0),(\omega_2,0),(\omega_2,0)\},\{(\omega_1,a),(\omega_2,a)\},\{(\omega_3,a)\}\}$ $\beta_3(\alpha_0)=0$, $\beta_3(\alpha_1)=\frac{y(\omega_1,a)+y(\omega_2,a)}{2}$ and $\beta_3(\alpha_2)=y(\omega_3,a)$	a in states $\omega_1,\omega_2,\omega_3$
$k=2$	$p_2 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0)\},\{(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_2(\alpha_3)=0$ and $\beta_2(\alpha_4)=\frac{2y(\omega_1,a)+2y(\omega_2,a)+y(\omega_3,a)}{5}$	a in states ω_1,ω_2, 0 in state ω_3
$k=1$	$p_1 = \{\{(\omega_1,0),(\omega_2,0),(\omega_3,0),(\omega_1,a),(\omega_2,a),(\omega_3,a)\}\}$ $\beta_1(\alpha_5)=\frac{2y(\omega_1,a)+2y(\omega_2,a)+y(\omega_3,a)}{9}$	0 in state $\omega_1,\omega_2,\omega_3$,

Table 11: Ergodic worldview-policy profile when $c(\omega_1,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)}{3} < c(\omega_2,a) < \frac{y(\omega_1,a)+y(\omega_2,a)}{2}$, and $c(\omega_2,a) < \frac{2y(\omega_1,a)+y(\omega_2,a)+y(\omega_3,a)}{4} < c(\omega_3,a)$, and $c(\omega_3,a) > \frac{2y(\omega_1,a)+2y(\omega_2,a)+y(\omega_3,a)}{5}$.

in Table 3).[58] Intuitively, increasing k^* from 2 to 3 introduces inefficiencies at the stage $k = k^* = 3$.

In this environment, when $k^* = 3$ and $c(\omega_2,a) < \frac{y(\omega_1,a)+y(\omega_2,a)}{2}$ and $c(\omega_3,a) > \frac{2y(\omega_1,a)+2y(\omega_2,a)+y(\omega_3,a)}{5}$, there exists another (pure) ergodic worldview-policy profile, described in Table 11. In fact, this profile yields a higher ergodic efficiency than the (unique pure) worldview-policy profile for $k^* = 2$ described in Table 10.

F Proof of Proposition 6

The proof follows by replicating the proofs of Proposition 1 (see Section 3.2 and Appendices A-B) and Proposition 2 (see Appendix C). Indeed, the tremble in the state transitions yields that for any (ω,a), the probability of reaching $\omega' \in \Omega$ from (ω,a) is strictly positive. This tremble, together with the tremble on policies, ensure that the data available to the agents at any time t contains observations for all (ω,a), and that the empirical frequencies are equal to the true probabilities: $\hat{y}(\omega,a) = y(\omega,a)$ for all

[58]Indeed,

$$\frac{1}{2}[y(\omega_1,a)-c(\omega_1,a)+y(\omega_2,a)-c(\omega_2,a)] > \frac{1}{3}[2y(\omega_1,a)-c(\omega_1,a)+y(\omega_2,a)-c(\omega_2,a)+y(\omega_3,a)-c(\omega_3,a)]$$

$(\omega, a) \in \Omega \times A$.

The same arguments as in the proof of Proposition 2 then apply, as politicians and voters are myopic and therefore do not internalize the consequences of their choices in the current-period on the state and outcomes in the next period.

G Proofs of Propositions 7 and 8

G.1 Proof of Propositions 7

The proof of Proposition 7 follows from the same arguments as the one of Proposition 1 (see Section 3.2 and Appendices A-B). Indeed, intellectuals choose worldviews with the same objective as politicians do in the baseline model leading to Proposition 1. And as the tremble ensures observations for all (ω, a) and empirical frequencies equal to the true probabilities ($\hat{y}(\omega, a) = y(\omega, a)$ for all $(\omega, a) \in \Omega \times A$), the politicians' choices at the policy-making stage do not affect the complexity of the worldviews chosen by intellectuals, as the latter remains determined on path only by the complexity of the incumbent's worldview and the voters' distribution of sophistication.

G.2 Proof of Proposition 8

The proof follows by replicating the proof of Proposition 2 (see Appendix C). Indeed, Proposition 7 ensures that in any ergodic worldview-policy profile, complexity dynamics are deterministic and as described by Proposition 1. The same arguments as in the proof of Proposition 2 thus yield the existence of an ergodic worldview-policy profile.

H Multiplicity of ergodic worldview-policy profiles in state-driven environments

Consider the following example of partition-based multiplicity of ergodic worldview-policy profiles. Let $\Omega = \{\omega_1, \omega_2\}$, with ω_1 and ω_2 equally likely, and $A = \{0, a\}$. Suppose moreover that

(i) $y(\omega_1, 0) < y(\omega_1, a) < y(\omega_2, 0) < y(\omega_2, a)$, with $\frac{y(\omega_1,a)-y(\omega_1,0)}{y(\omega_2,0)-y(\omega_1,a)} = \frac{y(\omega_2,a)-y(\omega_2,0)}{y(\omega_2,0)-y(\omega_1,a)} \ll 1$,

(ii) $y(\omega_i, a) - c(\omega_i, a) > 0 > y(\omega_i, 0) - c(\omega_i, 0)$ for $i \in \{1, 2\}$.

Cycle stage	Worldview	Policy
$k=3$	$p_3 = \{\{(\omega_1,0)\},\{(\omega_1,a)\},\{(\omega_2,0),(\omega_2,a)\}\} \equiv \{\alpha_0,\alpha_1,\alpha_2\}$ $\beta_3(\alpha_0)=y(\omega_1,0)$, $\beta_3(\alpha_1)=y(\omega_1,a)$ and $\beta_3(\alpha_2)=y(\omega_2,0)$	a in state ω_1 0 in state ω_2
$k=2$	$p_2 = \{\{(\omega_1,0,\omega_1,a)\},\{(\omega_2,0),(\omega_2,a)\}\} \equiv \{\alpha_3,\alpha_4\}$ $\beta_2(\alpha_3)=\frac{2y(\omega_1,0)+y(\omega_1,a)}{3}$ and $\beta_2(\alpha_4)=y(\omega_2,0)$	0 in state ω_1, 0 in state ω_2
$k=1$	$p_1 = \{\{(\omega_1,0),(\omega_1,a),(\omega_2,0),(\omega_2,a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5)=\frac{2y(\omega_1,0)+y(\omega_1,a)+3y(\omega_2,0)}{6}$	0 in state ω_1, 0 in state ω_2

Table 12: Ergodic worldview-policy profile

Cycle stage	Worldview	Policy
$k=3$	$p_3 = \{\{(\omega_1,0),(\omega_1,a)\},\{(\omega_2,0)\},\{(\omega_2,a)\}\} \equiv \{\alpha_0,\alpha_1,\alpha_2\}$ $\beta_3(\alpha_0)=y(\omega_1,0)$, $\beta_3(\alpha_1)=y(\omega_2,0)$ and $\beta_3(\alpha_2)=y(\omega_2,a)$	0 in state ω_1 a in state ω_2
$k=2$	$p_2 = \{\{(\omega_1,0,\omega_1,a)\},\{(\omega_2,0),(\omega_2,a)\}\} \equiv \{\alpha_3,\alpha_4\}$ $\beta_2(\alpha_3)=y(\omega_1,0)$ and $\beta_2(\alpha_4)=\frac{2y(\omega_2,0)+y(\omega_2,a)}{3}$	0 in state ω_1, 0 in state ω_2
$k=1$	$p_1 = \{\{(\omega_1,0),(\omega_1,a),(\omega_2,0),(\omega_2,a)\}\} \equiv \{\alpha_5\}$ $\beta_1(\alpha_5)=\frac{3y(\omega_1,0)+2y(\omega_2,0)+y(\omega_2,a)}{6}$	0 in state ω_1, 0 in state ω_2

Table 13: Ergodic worldview-policy profile

Lastly, suppose that $k^* = 3$.

There exist (exactly) two pure ergodic worldview profiles, described in Tables 12 & 13. Both feature fatalistic worldviews at stage $k = 2$, leading to the cost-minimizing policy (0), which is, in both examples, the inefficient choice in both states.

The multiplicity in Tables 12-13 arises more generally in fatalistic environments. Indeed, suppose that $k^* = |\Omega| + 1$ and that (for any of the data sets that will arise) for all $k \leq k^* - 1$, the k-optimal worldviews are fatalistic worldviews.[59] Then, there can exist multiple ergodic worldview-policy profiles, which differ in their partition and policy choice at stage k^* (and thus on probabilities of successes at all stages): each ergodic worldview-policy profile isolates at stage k^* exactly one couple (ω, a) such that policy a is chosen in state ω at stage k^*, but policy a_0 remains chosen in state ω at stages $k \leq k^* - 1$.

I Complements on optimal clustering with KL-divergence

Consider the following clustering problem: A couple $(\omega, a) \equiv X$ must belong to exactly one of two classes, A and B, with $\hat{b}(A) < \hat{b}(B)$. The conditional probabilities of success over classes are given by (1), which is a necessary condition for optimality (Lemma 1).

[59]At stage $k^* - 1$, the $(k^* - 1)$-optimal worldview thus distinguishes all states.

To which class should X belong to minimize the KL-divergence between the observed frequencies in the data and the distribution induced by the worldview?

Three properties of the optimal partition(s) thus arise:

(i) *Clustering couples with similar empirical probabilities of success:* For fixed strictly positive empirical masses $(\hat{m}(A), \hat{m}(B), \hat{m}(X) > 0)$, X is clustered with A if $\hat{y}(X)$ is sufficiently close to $\hat{b}(A) > 0$, resp. clustered with A, resp. clustered with B if $\hat{y}(X)$ is sufficiently close to $\hat{b}(B) < 1$.

(ii) *"Isolating extremes":* If the empirical mass of X, $\hat{m}(X)$, is close to zero, X is clustered with A if $\hat{b}(A) > 0$ and $\hat{b}(B)$ is close to 1, resp. clustered with B if $\hat{b}(B) < 1$ and $\hat{b}(A)$ is close to 0.[60]

(iii) *No class with zero mass (balancing mass across classes):* Fixing $\hat{m}(X) > 0$ and $\hat{m}(B) > 0$ (resp. $\hat{m}(A) > 0$), if the empirical mass of A (resp. B) is close to zero, then X is clustered with A (resp. B). Intuitively, if A has very little mass, it weights little in the KL-divergence, and it is thus more efficient to distinguish couples with more important masses.

Proof. To alleviate the notation, let for any class α (possibly a singleton), $\hat{m}_\alpha = \hat{m}(\alpha)$ and $\hat{y}_\alpha \equiv \hat{b}(\alpha)$. The difference between the KL-divergence of the worldview that clusters X with B, and the one that clusters X with A (with all conditional probabilities of success equal to the empirical probability of success) is equal to

$$
\begin{aligned}
&[\hat{m}_A(1-\hat{y}_A)+\hat{m}_X(1-\hat{y}_X)]\ln\left(\frac{\hat{m}_A(1-\hat{y}_A)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_A+\hat{m}_X}\right)\\
&\quad+[\hat{m}_A\hat{y}_A+\hat{m}_X\hat{y}_X]\ln\left(\frac{\hat{m}_A\hat{y}_A+\hat{m}_X\hat{y}_X}{\hat{m}_A+\hat{m}_X}\right)+\hat{m}_B(1-\hat{y}_B)\ln(1-\hat{y}_B)+\hat{m}_B\hat{y}_B\ln(\hat{y}_B)\\
&-[\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)]\ln\left(\frac{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_B+\hat{m}_X}\right)\\
&\quad-[\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X]\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X}{\hat{m}_B+\hat{m}_X}\right)-\hat{m}_A(1-\hat{y}_A)\ln(1-\hat{y}_A)-\hat{m}_A\hat{y}_A\ln(\hat{y}_A)\\
=&\hat{m}_A(1-\hat{y}_A)\left[\ln\left(\frac{\hat{m}_A(1-\hat{y}_A)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_A+\hat{m}_X}\right)-\ln(1-\hat{y}_A)\right]+\hat{m}_A\hat{y}_A\left[\ln\left(\frac{\hat{m}_A\hat{y}_A+\hat{m}_X\hat{y}_X}{\hat{m}_A+\hat{m}_X}\right)-\ln(\hat{y}_A)\right]\\
&+\hat{m}_X(1-\hat{y}_X)\left[\ln\left(\frac{\hat{m}_A(1-\hat{y}_A)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_A+\hat{m}_X}\right)-\ln\left(\frac{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_B+\hat{m}_X}\right)\right]\\
&+\hat{m}_X\hat{y}_X\left[\ln\left(\frac{\hat{m}_A\hat{y}_A+\hat{m}_X\hat{y}_X}{\hat{m}_A+\hat{m}_X}\right)-\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X}{\hat{m}_B+\hat{m}_X}\right)\right]\\
&-\hat{m}_B(1-\hat{y}_B)\left[\ln\left(\frac{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_B+\hat{m}_X}\right)-\ln(1-\hat{y}_B)\right]-\hat{m}_B\hat{y}_B\left[\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X}{\hat{m}_B+\hat{m}_X}\right)-\ln(\hat{y}_B)\right]
\end{aligned}
$$

[60]The result holds when the mass of X is small relative to the mass of either A or B, so that its clustering with one or the other does not affect too much the entropy of the resulting class.

(i) Clustering couples with similar empirical probabilities of success. For $\hat{y}_X$ close to $\hat{y}_A$, the above difference is equal to

$$\begin{aligned}
&\hat{m}_X(1-\hat{y}_A)\ln(1-\hat{y}_A) + \hat{m}_B(1-\hat{y}_B)\ln(1-\hat{y}_B) \\
&- [\hat{m}_X(1-\hat{y}_A) + \hat{m}_B(1-\hat{y}_B)]\ln\left(\frac{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_A)}{\hat{m}_B+\hat{m}_X}\right) \\
&+ \hat{m}_X\hat{y}_A\ln(\hat{y}_A) + \hat{m}_B\hat{y}_B\ln(\hat{y}_B) - [\hat{m}_X\hat{y}_A + \hat{m}_B\hat{y}_B]\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_A}{\hat{m}_B+\hat{m}_X}\right) + O(|\hat{y}_X-\hat{y}_A|),
\end{aligned}$$

which is strictly positive for $\hat{y}_X$ sufficiently close to $\hat{y}_A$, since $\hat{y}_B > \hat{y}_A$.[61] Consequently, for $\hat{y}_X$ close to $\hat{y}_A$, X is clustered with A in the optimal worldview. Similar computations yield that for $\hat{y}_X$ close to $\hat{y}_B$, X is clustered with B in the optimal worldview.

(ii) Isolating extremes. For $\hat{m}_X = 0$, the difference between the KL-divergences of the two worldviews is equal to zero. Its partial derivative with respect to $\hat{m}_X$, taken at $\hat{m}_X = 0$, is equal to

$$(1-\hat{y}_X)\ln\left(\frac{1-\hat{y}_A}{1-\hat{y}_B}\right) + \hat{y}_X\ln\left(\frac{\hat{y}_A}{\hat{y}_B}\right), \tag{3}$$

which is strictly positive for $\hat{y}_B$ sufficiently close to 1 (fixing $\hat{y}_X$ and $\hat{y}_A > 0$). Hence, for $\hat{m}_X$ close to zero and $\hat{y}_B$ close to 1, X is clustered with A in the optimal worldview. Conversely, fixing $\hat{y}_X$ and $\hat{y}_B < 1$, for $\hat{m}_X$ close to zero and $\hat{y}_A$ close to 0, (3) is strictly negative, and thus X is clustered with B in the optimal worldview.

(iii) No class with zero mass. For $\hat{m}_A$ close to zero, the difference between the KL-divergences of the two worldviews is equal to

$$\begin{aligned}
&\hat{m}_X(1-\hat{y}_X)\ln(1-\hat{y}_X) + \hat{m}_B(1-\hat{y}_B)\ln(1-\hat{y}_B) \\
&- [\hat{m}_X(1-\hat{y}_X) + \hat{m}_B(1-\hat{y}_B)]\ln\left(\frac{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)}{\hat{m}_B+\hat{m}_X}\right)] \\
&+ \hat{m}_X\hat{y}_X\ln(\hat{y}_X) + \hat{m}_B\hat{y}_B\ln(\hat{y}_B) - [\hat{m}_X\hat{y}_X + \hat{m}_B\hat{y}_B]\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X}{\hat{m}_B+\hat{m}_X}\right) + O(\hat{m}_A).
\end{aligned}$$

which is positive for $\hat{m}_A$ sufficiently close to zero, strictly so whenever $\hat{y}_X \neq \hat{y}_B$.[62] Consequently, for $\hat{m}_A$ close to zero, X is clustered with A in the optimal worldview. □

[61] Writing this difference as $\varphi(\hat{y}_A, \hat{y}_B) + O(|\hat{y}_X - \hat{y}_A|)$, $\varphi(\hat{y}_A, \hat{y}_A) = 0$, and the partial derivative of φ with respect to $\hat{y}_B$ is equal to

$$\hat{m}_B\left[\ln\left(\frac{\hat{y}_B}{1-\hat{y}_B}\right) - \ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_A}{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_A)}\right)\right],$$

which is strictly positive for $\hat{y}_B > \hat{y}_A$. Hence, $\varphi(\hat{y}_A, \hat{y}_B) > 0$.

[62] Writing this difference as $\varphi(\hat{y}_X, \hat{y}_B) + O(\hat{m}_A)$, $\varphi(\hat{y}_B, \hat{y}_B) = 0$, and the partial derivative of φ with

respect to $\hat{y}_X$ is equal to

$$\hat{m}_X\left[\ln\left(\frac{\hat{y}_X}{1-\hat{y}_X}\right)-\ln\left(\frac{\hat{m}_B\hat{y}_B+\hat{m}_X\hat{y}_X}{\hat{m}_B(1-\hat{y}_B)+\hat{m}_X(1-\hat{y}_X)}\right)\right],$$

which is strictly negative for $\hat{y}_X < \hat{y}_B$ and strictly positive for $\hat{y}_X > \hat{y}_B$. Hence, $\varphi(\hat{y}_X, \hat{y}_B) \geq 0$, with strict inequality whenever $\hat{y}_X \neq \hat{y}_B$.