Closed Codes
Three Paths to the Downfall of an Elite*

Steffen Huck
WZB & UCL

24 October 2026

Abstract

An elite controls a screen and may show its own code or a common code. Part of the audience cannot read the elite's own code. A citizen who misses an episode asks her neighbour whether she got it, and a season in which nobody around her got it leaves a grievance with the elite. The elite may get increasingly infatuated with its own code and rely more heavily on it over time. I examine two channels through which those not familiar with the elite's code may acquire it: leakage through acquaintances and public schooling. I characterise the elite's season-by-season choice of what to show on the screen and the three routes that lead to its downfall. Learning is not monotone for the elite. In the limit of slow learning a little is never better than none, and it can be fatal where none is not.

1 Introduction

In a school class there are three groups of kids, and they all like jokes. Everybody in the class understands the teacher jokes. But there are also two smaller groups of kids who have their own private jokes: the football kids and the horse kids.

When a football kid says "nutmeg!" the football kids laugh. They know it means the ball went between somebody's legs, which is the most embarrassing thing that can

*I wrote this paper with the assistance of two AI models, Claude Fable and ChatGPT Astra. I used Astra mainly as an adversarial referee and Claude as a companion throughout—to develop and refine the model and proofs, to draft and to redraft. The opening fable is one example of this joint production; it took many iterations before it converged. The idea is mine, and so is all responsibility for what the paper says.

happen to you on a pitch. Everyone else hears a spice and wonders what cake has to do with anything.

When a horse kid says "she's got a sock!" the horse kids laugh. They know it means a pony with one white leg. Everyone else wonders why anyone is talking about laundry.

When someone says "Mr Burns is on fire today," the whole class laughs. Everybody gets the teacher jokes.

Nobody is cleverer. They just know different things.

The class has a class rep. Right now it's Lena, from the football group. The class rep decides what goes on the big wall at the back of the classroom. That's the one thing being class rep gives you.

Lena can put teacher jokes on the wall, which everybody gets. Or she can put up football jokes. She likes football jokes, and her football friends love seeing them on the wall.

When a kid who isn't a football kid looks at a joke on the wall and doesn't get it, she doesn't know why. Maybe she just missed it, and she'd have got it on another day. Or maybe it was a football joke, and it was never meant for her.

So she asks the kids around her. If they got it, she thinks: I really just missed it. If they didn't get it either, she thinks: that was a football joke. The wall is not for us.

A kid is not put off by one bad day. She is put off by an entire bad school term with too many jokes she didn't get. Then she goes home and says "Lena only cares about the football kids." And when, at the start of the next term, a new class rep has to be chosen, she may vote for Paul, from the horse group.

None of the kids who only get the teacher jokes ever runs against Lena or Paul. They have no jokes of their own, and nobody thinks of them as a group, not even they themselves. So when a kid is fed up with Lena, Paul is who there is.

Lena knows this. So she puts up some football jokes, but not too many. How many depends on how much she loves football, how much she wants to stay class rep, and how many kids aren't football kids.

Two more things. The more football jokes Lena puts up, the fonder she gets of them. They are just too funny, aren't they? And some kids sit with football kids, at mixed tables, and after a while they may get the football jokes too. Then Lena sees more kids laughing and puts up even more football jokes. But at the tables where there are no football kids to learn from they get even angrier about Lena when the wall becomes nearly all football.

Then, when the next class rep vote is up, Paul may win. He puts horse jokes on the wall. Now it's everybody else who doesn't get them, who asks the kids around them, and who goes home saying "Paul only cares about the horse kids."

This is what the paper is about. The rest of it says the same thing with a model, so that one can see what has to be assumed and what follows.

The wall is a screen, and the class rep is whoever is in power and commands it. The idea for this paper goes back to the late 2000s when—thanks to a new satellite dish—I was able to catch up on *Tatort*, Germany's most popular crime show on television. I hadn't followed the show in over a decade and was struck by how much it had changed. In the late 1990s, when I had last watched it, the inspectors went home to living rooms with a *Schrankwand*, the wall cabinet with a hole for the television. The people whose crimes they solved had the same cabinet. All of this had gone some time in the new millennium. The inspectors, I now saw, lived in fancy lofts or shared flats, sparsely furnished but invariably with a couple of Bauhaus objects or an Eames chair, and the criminals had moved either to villas at the top or to squalor at the bottom. Someone had decided that police inspectors would live like this, and had even taken away the possibility that people with a *Schrankwand* could fall victim to a crime or commit one.

Would that not possibly alienate a rather large chunk of viewers? I stored that thought and only revisited it a decade later when the AfD, the German right-wing populist party that once began out of discontent with Europe and the euro, started to rise in the polls. With Maja Adena I embarked on an ambitious empirical project in 2019—in many ways a project about furniture and the furnishing of right-wing attitudes.

We ran a pre-registered survey of about 5,000 German voters in December 2019, exploring familiarity with modernism. Respondents see, for example, a Wagenfeld lamp and are asked which of three objects pairs with it: another Bauhaus piece, an ordinary lamp, or something unrelated. The ordinary lamp is easy. Spotting the Bauhaus connection requires understanding of a code. As it turned out, missing the Bauhaus pairing predicts support for the AfD, and it keeps predicting it after all the usual socioeconomic controls, including educational attainment, income and wealth. The test was one of recognition. It did not ask whether the respondent liked any of the lamps. It simply probed the understanding of a code.

This is the theory companion to that empirical work. It asks how an elite fills a screen with a code that part of its audience cannot read; why it keeps doing so; and what that does to its hold on power. The answer the model gives has nothing to do with modernism in particular. It is about how an elite that commands the screen may use an exclusionary *closed* code and how increasing infatuation with it (remember Lena?) may be one path to its downfall.

If the mechanism were a link between the right and modernism as a code, history would refute it. Italian fascism embraced modernism, deeply and with lasting architectural results. In stark contrast, the Nazis declared modernism degenerate and closed the Bauhaus in Dessau in 1932.

There are echoes of this in contemporary politics. In September 2026 the AfD won the state election in Sachsen-Anhalt, which contains Dessau, with 39 of 83 seats, and an AfD-led government is in prospect. The party has the Bauhaus in its sights again. Its

programme calls the school contested and without national roots, and in 2024 its Landtag group moved for a critical reckoning with it under the title "Irrweg der Moderne", which is Schultze-Naumburg's phrase from *Kunst und Rasse* of 1928.[1] So the link is not between the right and modernism as such. Whatever it is, it has to allow modernism to change sides.

The model has three groups, three codes, one screen, and time. A code is a language in which things are shown. Everybody knows the common code, the teacher jokes from the fable above. Two of the groups, to which I will refer, for obvious reasons, as milieus, have a code of their own as well; the third, which may be the largest, knows the common code only. A citizen either knows a code or does not, and whoever knows it may still miss an episode (a joke) now and then. There is some noise.

Whoever holds the screen is the elite, and the elite is a role: whoever won the last political contest and controls the screen. Candidates come from the two milieus, because a code is what makes a group a group, and the third (those who only get the teacher jokes) has none; the paper says why this is the right assumption and what it costs.

The elite likes its code but has to fear the next vote. Each season the commander of the screen chooses what share of it is in her own milieu's code. Citizens watch, miss or get, ask their neighbours, and remember their bad seasons. Nobody is cultivated by assumption. The only asymmetry in the model is who holds the screen.

I model this as a dynamic problem from Lena's perspective, the current elite controlling the screen. Lena is both the one who picks the jokes and the one who stands for election; in the world those are often two people, the editor and the minister, and Section 6 says what seating them in one chair assumes. There are three stocks that move. Grievances accumulate among the excluded and only fade slowly. The elite's taste for its own code grows with use. And the code leaks: it can be learned, either from acquaintances or through public schooling.

I should flag that throughout the paper I assume that Lena is myopic. She does not fully appreciate how her own tastes will change through the repeated use of football jokes, nor how the stock of grievances builds, nor that the class is learning.

Four things come out. First, any season's problem—the choice between the codes—has a unique solution, and it is interior for a wide range of parameters without any assumption on the curvature of preferences. The convexity that makes it so comes from the seasons. A citizen is not put off by one bad Monday, after a Sunday on which she did not get the *Tatort* episode; it takes a full season of bad Mondays.

[1] Landtag von Sachsen-Anhalt, Drucksache 8/4681, debated in October 2024 and rejected by all other groups (Landes, 2025). The programme passages and the origin of the phrase are reported by Brierley (2026). The Landtag elected on 6 September 2026 convened on 6 October; at the time of writing no government has been formed, and the AfD's leader has announced that he will stand for Minister-President in December (*Stuttgarter Zeitung*, 6 October 2026).

Second, an elite that starts with a majority can lose its grip on power in three ways, and the model says which parameters select which. Slowly, when grievances are not easily forgotten. Faster through too much infatuation with one's own code. And, possibly even faster, when the leakage of codes backfires in a surprising fashion—a counterintuitive scenario to which I will devote considerable space below.

Indeed, this is the third takeaway: the elite is not necessarily safer when its code diffuses more easily, be it through acquaintances or schooling. With more learning, usage of its own preferred code becomes cheaper for the elite, and when Lena does not fully account for how her preferences will become ever more biased, learning through leakage or schooling can be part of her eventual downfall. The result is sharper than that. The elite is safe with no learning, ousted with a little, and safe again with a lot; in the limit of slow leakage, a little learning is never better for the elite than none, and that is a theorem. Half-done cultural education is worse than none.

Fourth, when the screen changes hands the new elite faces the same problem with exchanged roles. That is the model's answer to Italy and to Dessau: the challenger drops whatever was on the screen while she was not holding it, and shows what is in her own code.

I proceed as follows. Section 2 places the model in the literature. Section 3 sets out the model and solves one season. Section 4 offers the three paths to an elite's eventual downfall and examines comparative statics. Section 5 examines what happens after the downfall (when Paul takes over). Section 6 offers discussions and interpretations, Section 7 names extensions, and Section 8 is a coda.

2 Related Literature

Distinction, and why the elite does not move on. Bourdieu (1984) is the canonical ancestor, and the ancestry is direct: *Distinction* describes cultural consumption as decoding from its first pages, where "a beholder who lacks the specific code feels lost in a chaos of sounds and rhythms, colours and lines, without rhyme or reason" (p. 2). His distinction between the pure gaze, which sees form, and the popular aesthetic, which sees function, is the Wagenfeld puzzle in operation. The use of the term *milieu* is a nod to the German milieu research that develops Bourdieu's analysis of class, habitus and lifestyle (Vester et al., 2001). Simmel (1904) supplies the cyclical nature of what I briefly discuss in the extension section: a code may only distinguish as long as it is not imitated by the many, but when that happens the elite will find new codes. Economics has formalised the cycle of imitation and distinction in Karni and Schmeidler (1990) and Pesendorfer (1995), where a design loses its value once it has diffused and the designer innovates, and Corneo and Jeanne (1997) model the status consumption that drives imitation, with snobs and bandwagons. In the cycle models the elite moves on once imitated. The elite in this

paper's main model digs in. It has no other closed code to move to, and it grows fonder of the one it has. Imitation is accounted for by leakage: the illiterate may acquire the code through acquaintances. Smaldino, Flamson and McElreath (2018) model covert signalling, jokes that are understood by the intended audience and obscure to everyone else, as a way of not alienating outsiders; this paper is about what happens when the outsiders notice.

Identity and cultural conflict. Akerlof and Kranton (2000) put identity in the utility function of economic models and Shayo (2009) let voters choose whom to identify with. Bonomi, Gennaioli and Tabellini (2021) and Gennaioli and Tabellini (2025) have culture crowd out economics as the political cleavage, through the salience of stereotypes and through parties that voice them; Levy, Razin and Young (2022) have rule by the sophisticated worsen the simple group's misreading of the world and produce cycles. These are the paper's nearest formal neighbours and they share some of its questions. They do not share its mechanism. In none of them does the elite's own cultural output alienate, season by season, a group that cannot choose whether it understands. Identity enters this paper in one place only, and from the other side: a shared code is what makes a group a group and lets it field a candidate, and the people who share only the common code do not perceive of themselves as a group and field none; in Bourdieu's (1985) terms, they are a classificatory group and not a mobilised one. The empirical side of this intellectual neighbourhood, Gidron and Hall (2017), Norris and Inglehart (2019), Noordzij, de Koster and van der Waal (2021), Gethin, Martínez-Toledano and Piketty (2022), and the survey of Guriev and Papaioannou (2022), examines the correlation between cultural status, perceived distance from elites and populist support. These studies connect subjective status and perceived cultural distance from politicians with anti-establishment attitudes, and document the rising political weight of sociocultural cleavages. They motivate the mechanism studied here; none of them identifies the effect of exclusionary cultural output or the accumulation of grievances, and the model is a proposal for what the generating process might be.

Learning a code. Leakage is the acquisition of a competence through contact, and its closest relative is Lazear (1999), where a minority learns the majority's language and learns it faster where the majority is more prevalent. That comparative static will play out when the model introduces geographical regions. Lazear's learning is a deliberate investment to widen trade. In my model this is, in contrast, captured through mere contagion, either through random contacts or through schooling; the law of motion is the contagion of Bass (1969), which Young (2009) distinguishes from social learning proper. Bisin and Verdier (2000, 2001) examine the transmission of preferences across generations. This is related but what leaks here is not a preference but the ability to understand. Alesina and Fuchs-Schündeln (2007) find East German preferences converging towards Western ones and extrapolate a horizon of one to two generations; Bondar and Fuchs-Schündeln (2023) find the convergence slowing. This paper has a competence rather than a preference

converge slowly, and shows that the slowness can be fatal to the incumbent.

Grievances that accumulate. Passarelli and Tabellini (2017) make political unrest a function of grievances against the current state. In my model there is no explicitly modelled unrest but once grievances reach a threshold, the elite will be ousted. Mueller (1970) found presidential approval falling with time in office and read it, following Downs (1957), as a coalition of minorities alienated one after another. The grievance of this paper is a stock of unforgotten bad seasons, so that the length of memory is going to be a parameter of interest.

Infatuation. The habit stock is Pollak (1970): myopic habit formation, not the rational addiction of Becker and Murphy (1988), whose addict foresees the habit and whose beneficial consumption capital is the case this paper's taste corresponds to, although I will address the role of foresight in my extensions, where it cuts both ways: foresight of the habit raises the share, foresight of the grievances lowers it, and which wins depends on what office is worth. Lévy-Garboua and Montmarquette (1996) have theatre-goers discover a taste they did not know they had by going; the taste here is not discovered but grown, and it grows through the choice of supply. The elite gets ever more infatuated with what it produces for the screen.

Citizens who reason from samples. A citizen who is put off by a season of bad Mondays is a procedural reasoner in the sense of Osborne and Rubinstein (1998, 2003), Spiegler (2006) and Salant and Cherry (2020): she reasons from a small sample of her own experience rather than solving the elite's problem. The counting rule of this paper is one specific member of that family. Remark 1 says what a citizen who did solve it would do to the model, and why the counting citizen appears to be the right assumption. Gerbner and Gross (1976) and Tuchman (1978) supply the idea that absence from the screen is a status fact, which they call symbolic annihilation; Lenski (1954) supplies status inconsistency, the rich citizen who cannot read the code. Durante, Pinotti and Tesei (2019) show that what entertainment television showed shaped who later voted for whom, through a different channel from the one here.

Carmichael (2025) models electoral competition over substantive against symbolic representation. What this paper adds, and what I have not found elsewhere in combination, is an incumbent who controls the cultural output, an audience that cannot all read it, grievances that accumulate from failed decoding, and a code that diffuses; and it is dynamic. It has an elite that chooses and may be willing to bear the costs, even if it means its ultimate downfall.

3 The Model

3.1 One season

Groups and regions. Three groups live in two regions, that I shall call, with Germany in mind, West and East, with population shares μ_W and μ_E. Two of the groups are milieus with a code of their own, A and B; the third, C, has none. In region r a share α_r of the population belongs to A and a share γ_r to C, with the rest belonging to B; the West has more of A, $\alpha_W > \alpha_E$. A holds the screen, and its national mass is $m_A = \sum_r \mu_r \alpha_r$. Everyone who is not of milieu A I call a citizen of $\bar{A}$, the others; while A holds the screen nothing distinguishes a citizen of milieu B from a citizen of group C, and γ_r enters only in Section 5, which examines what happens after a change in power. A citizen's acquaintances are drawn from her own region with probability $1 - \omega$ and from the whole country with probability ω. The main analysis examines $\omega = 0$, sealed regions; the general case is kept for the extensions in Section 7.

Codes: knowing and getting. There are three codes, the common code c and the milieu codes a and b. Knowing a code is a state, and a citizen knows which codes she knows. Everybody knows code c. Citizens of milieu A know code a, because it is theirs, and citizens of milieu B know code b for the same reason; citizens of group C know code c only. A citizen outside milieu A may nonetheless have learned code a. Among the citizens of region r who are not of milieu A, the fraction who know code a is $\ell_r \in [0, 1]$, the region's literacy; I call the rest illiterate. Getting an episode is a draw. An episode in a code one knows is missed with probability $\lambda \in (0, 1)$; an episode in a code one does not know is missed with probability $\varphi \in (\lambda, 1]$. So a citizen of milieu A misses an own-code episode now and then, like anyone.

The screen. The elite chooses the share $\pi \in [0, 1]$ of the season's episodes shown in code a; the rest are in code c. The share is a house style, an instruction to the editors about what fraction of the season's episodes are in code a; which particular episodes are in it is nobody's choice. For analytical convenience, I assume that a season has many episodes, a continuum of them. Each viewer watches $m \geq 2$ of them, one per Sunday, drawn at random; in the calibration $m = 2$. The continuum does two things. It makes the share a continuous choice, where a finite schedule would let the elite pick only whole episodes. And it makes the season's aggregates deterministic: each viewer draws her own episodes, so the law of large numbers runs across viewers. At the same time, not everybody shares everybody else's luck in understanding; citizens have different draws. A reader who prefers to think of a season as two screened episodes, each in code a with probability π, gets the same expected rate of bad seasons and the same expected objective for the elite in any one season. What the continuum removes is the season's common luck. With two shared

broadcasts the paths are random, and an elite that is safe below can lose its majority after a long enough run of own-code seasons, which eventually arrives. The deterministic paths and the notion of safety that follow use the continuum. Each episode is watched by a random sample of the region, and each Sunday the viewer gets her episode or misses it. A season is the unit over which bad Mondays are counted and nothing more.

Mondays. A viewer who missed Sunday's episode asks someone who watched it whether he got it. The conversation is about last night's *Tatort*, and the person one talks to about it is someone who saw it, drawn at random from the episode's audience, which is itself a random sample of the region. He answers yes or no and says nothing else. Each Sunday has its own episode and its own audience, which keeps the Mondays of a season independent. A *bad Monday* for an illiterate citizen of region r outside milieu A is a Monday on which she missed the episode and so did the person she asked. Its probability is

$$g_r(\pi) = \pi\,\varphi\,F_r + (1-\pi)\,\lambda^2, \qquad F_r = \nu_r\varphi + (1-\nu_r)\lambda, \qquad \nu_r = (1-\alpha_r)(1-\ell_r), \quad (1)$$

where ν_r is the chance that the person asked is an illiterate citizen outside milieu A and F_r the chance that he misses an a-episode. With $\omega = 0$ both are regional. g_r is affine in π: one episode is one draw of the code, and conditional on the code nothing anyone sees depends on π. Asking n people instead of one puts the exponent n on F_r and λ in (1); Section 4.5 says what that does.

Bad seasons. A citizen is not put off by one bad Monday. A *bad season* is a season in which every Monday was bad, and its probability is

$$a_r(\pi) = g_r(\pi)^m. \quad (2)$$

This is strictly convex in π for $m \geq 2$, a product of m increasing affine functions. The convexity comes from two things together, and both are modelling choices. The calendar supplies m independent draws of the code, and the rule that every Monday must be bad makes their product the object. A rule under which two bad Mondays out of three suffice gives $3g^2 - 2g^3$, which is convex only below $g = \frac{1}{2}$; Section 7 checks that rule.

Who draws the conclusion. A bad season is possible under a neutral screen: $a_r(0) = \lambda^{2m} > 0$. Citizens of milieu A have bad seasons too, and so do literate citizens outside milieu A, at rates of the same form. The model gives grievances only to the illiterate, and the reason is what a bad season can mean. A citizen who knows every code on the screen knows that her bad season was bad luck; there is nothing else it could have been. A citizen who does not know code a cannot tell a season of bad luck from a season of exclusion, and the rule she follows is to count it as exclusion. The grievance is in the inference, and only those who might be excluded make it. Two things follow. A neutral screen does not buy

the incumbent a clean slate: at $\pi = 0$ every bad season is bad luck, and every one of them is counted.

Votes. Each citizen draws an attachment ξ to the incumbent, uniform on $[-\frac{1}{2}+\beta, \frac{1}{2}+\beta]$ with $\beta \geq 0$, independently of everything else. A citizen of milieu A gains $\rho > 0$ when the episode she watched was in her own code. An illiterate citizen outside milieu A carries a grievance stock S_r, defined in Section 3.2, of which this season's bad season, if it was one, is part; each unit of it costs the incumbent $\varepsilon > 0$ in her vote. Literate citizens outside milieu A carry nothing. A citizen votes for the incumbent if ξ plus these terms is nonnegative, so that the incumbent's approval at the end of the season is

$$V = \tfrac{1}{2} + \beta + m_A \rho\, \pi - \varepsilon \sum_r \mu_r (1-\alpha_r)(1-\ell_r)\, S'_r, \tag{3}$$

where S'_r is the stock after the season. The β is a baseline majority: with a neutral screen and no grievances the incumbent has $\frac{1}{2} + \beta$. Citizens of milieu A are not a bloc. Their attachment is the same draw as everyone's, and the screen moves them only by ρ; so the elite's milieu may be a majority of the country and the elite may still lose its majority, and in the calibration below it is a minority. For (3) to be exact every individual's vote probability must stay inside the unit interval. A citizen of milieu A needs $\rho \leq \frac{1}{2} - \beta$. An illiterate citizen outside milieu A can carry at most $1/\theta$ unforgotten bad seasons, where θ is the forgetting rate of Section 3.2, so she needs $\varepsilon/\theta \leq \frac{1}{2} + \beta$. I maintain both throughout, and the second is the binding one. A long memory must come with a light weight per grievance.

The occupant of the role. The screen is held by a leader of A or of B. Citizens of group C field no candidate: a code is what makes a group a group, and C has none; Section 6 defends the assumption and says what it costs. The elite of season t maximises $\psi\pi + RV$, where $\psi \geq 0$ is the weight she places on her milieu's recognition and $R > 0$ is the value she places on a unit of approval. Approval is what the next contest will be fought on, and the paper studies its path. The contest itself is not modelled, and RV is a preference over approval, not an expected rent from a specified election. She chooses for the season and takes the stocks as given. Whether the same person holds the role season after season, as Lena does, or a succession of editors, makes no difference to the model: either way the occupant of season t sets the share that is best for season t. Since the stock of past grievances enters V additively, it drops out of her choice: an occupant whose audience's patience has been used up chooses exactly as one whose audience is fresh. Three things are ignored in this way. This season's bad seasons stay in the stock, so a unit of own code that costs her G now costs G/θ over the audience's memory, undiscounted. This season's share feeds the habit, which raises what every later share is worth to whoever shows it. And this season's share teaches the audience. Only the first is a cost she would cut if she saw it; the other two are benefits to later seasons, and each of them, seen on its

own, pulls the share up. Section 7 computes what an occupant who saw all three does. Write

$$G = \psi/R + m_A\rho, \qquad c_r = \varepsilon\,\mu_r(1-\alpha_r)(1-\ell_r).$$

G collects what pulls her towards her own code, and c_r is the weight of region r's illiterate citizens in the cost.

The first proposition of this paper can now be stated.

Proposition 1 (The season)**.** *Suppose some region has illiterate citizens outside milieu A, so that $\sum_r c_r > 0$. The season's problem is strictly concave in π and has a unique solution,*

$$\pi^* = \begin{cases} 0 & \text{if } G \le \sum_r c_r\, a_r'(0), \\ 1 & \text{if } G \ge \sum_r c_r\, a_r'(1), \\ \text{the unique root of } G = \sum_r c_r\, a_r'(\pi) & \text{otherwise,} \end{cases}$$

with $a_r'(\pi) = m\, g_r(\pi)^{m-1} g_r'$ and $g_r' = \varphi F_r - \lambda^2 > 0$. It depends on the state only through literacy and taste, and it is weakly increasing in ψ, ρ and each ℓ_r, weakly decreasing in ε and in each cost weight c_r holding G and the g_r fixed, and weakly decreasing in R when $\psi > 0$, strictly in each case at an interior solution. For a single region and $m = 2$,

$$\pi^* = \frac{G/(2c\,g') - \lambda^2}{g'} \quad \text{clipped to } [0,1].$$

Proof. Dropping what does not involve π, the objective is $G\pi - \sum_r c_r a_r(\pi)$. Since g_r is affine increasing and $a_r = g_r^m$, $a_r'' = m(m-1)g_r^{m-2}(g_r')^2 > 0$, and some $c_r > 0$, so the derivative $G - \sum_r c_r a_r'(\pi)$ is strictly decreasing; the corner tests and the root follow. The grievance stock enters the objective additively and drops out. For the comparative statics: ψ, ρ and R enter only through G; ε only through c_r; and a rise in c_r with G and the g_r fixed raises $\sum_r c_r a_r'$ pointwise. (The regional primitives μ_r and α_r enter m_A and ν_r as well as c_r; at fixed μ_r a lower α_r lowers G and raises g_r, so all three channels push the share the same way.) Literacy ℓ_r lowers c_r and lowers ν_r, hence F_r, hence g_r and g_r' at every π, so it lowers $\sum_r c_r a_r'$ pointwise, and the root rises; at a corner the tests move the same way. The explicit form is $G = 2c\,g(\pi)g'$ solved for π. □

Two observations. (i) A little of one's own code is almost free: $a_r'(0) = m\lambda^{2(m-1)} g_r'$, so the region of "none" is of order $\lambda^{2(m-1)}$. (ii) Literacy makes own code cheaper, which is the engine of Section 4.

The citizen who counts her bad Mondays is the procedural reasoner of Section 2: she acts on her own small sample and does not solve the elite's problem. The remark that follows says what a citizen who inferred instead of counted would do to the model, and why I do not use her.

Remark 1 (Why not a Bayesian who knows the share). Suppose instead that the citizen knows π, asks about her own episode, and forms the posterior that it was in code a; and that this posterior enters her vote linearly. Then the posterior, conditional on her own miss, has expectation $\pi\varphi/[\pi\varphi + (1-\pi)\lambda]$ whatever the neighbours say, because the expectation of a posterior is the prior; multiplied by the mass of citizens who miss, the exclusion term in approval is $\pi\varphi$. The elite's cost is affine in π and her choice is a corner. Neighbours move the belief from some citizens to others and do not change its total. A citizen who knows the share has nothing to learn about the wall; what she learns is about her episode. The citizen of this paper does not know the share, does not need to, and is not put off by a belief but by a season. This is the one place where the paper's behavioural assumption does work.

3.2 The stocks

Grievances, the elite's own taste for its code, and literacy are stocks that move from one season, indexed by t, to the next.

Grievances. $S_{r,t}$ is the number of unforgotten bad seasons per illiterate citizen of region r outside milieu A. A bad season adds one, and each season a share $\theta \in (0,1)$ of the stock is forgotten:

$$S_{r,t+1} = (1-\theta)S_{r,t} + a_r(\pi_t), \qquad S_{r,0} = 0. \tag{4}$$

The vote of season t counts this season's bad season, so it uses $S_{r,t+1}$. Two things are assumed. Those who learn the code in a season are drawn at random among the illiterate, so learning leaves the stock per illiterate citizen unchanged. And a citizen who has learned the code stops minding; her grievances lapse with her illiteracy. Section 7 checks what happens when they do not.

Taste. The elite's weight on recognition is $\psi_t = \psi_0 + h_t$, where h is a habit stock, fed by showing the code and fading with disuse:

$$h_{t+1} = (1-\zeta)h_t + \eta\,\pi_t, \qquad h_0 = 0, \tag{5}$$

with $\zeta \in (0,1]$ and $\eta \geq 0$. Under a constant share the stock settles at $\eta\pi/\zeta$. The current elite, the holder of the screen (Lena in the fable) takes ψ_t as given. She does not see that she is feeding it; whether she would restrain herself if she did is a question for Section 7.

Literacy. Leakage passes the code, not merely the understanding of a single episode. What an illiterate citizen learns from an acquaintance who knows the code is not what

last night's joke meant but what a nutmeg is, and once she knows it she gets the next one herself. In parallel, schools teach regardless of how neighbourhoods are composed.

The law of motion of literacy is

$$\ell_{r,t+1} = \ell_{r,t} + (1 - \ell_{r,t})\big(\kappa\, \pi_t\, k_r + s\big), \qquad k_r = \alpha_r + (1 - \alpha_r)\ell_{r,t}, \tag{6}$$

where k_r is the chance that an acquaintance knows code a, $\kappa \geq 0$ is the rate of social leakage and $s \geq 0$ the rate of schooling, with $\kappa + s \leq 1$ so that literacy stays in the unit interval. At the start nobody outside milieu A knows code a, so $\ell_{r,0} = 0$ in both regions. With sealed regions the East leaks at the pace its own citizens of milieu A set: few to learn from at first, then, as local readers accumulate, a contagion that catches up with the West.

Timing and ousting. In season t the occupant observes (ℓ_t, S_t, h_t) and chooses $\pi_t = \pi^*(\ell_t, \psi_t)$ by Proposition 1. Citizens watch, ask, and vote; approval V_t is (3) with $S_{r,t+1}$. The stocks update by (4)–(6). I say that the elite *loses its majority* in the first season with $V_t < \frac{1}{2}$, and for brevity call that the season of ousting, or the downfall. An elite I call *safe* is one that never loses its majority. The paper is a theory of majority approval along a path conditional on continued incumbency. The figures that follow report that path and continue it past the season of ousting as if the elite were still there. This shows how an elite may hold on to power, and spread its code, in a world where downfalls are not deterministic: what could have been, or might yet be, if an elite survives a severe crisis.

Before the three paths, one lemma says how the stocks move along any path. Literacy never falls, and the West is always ahead of the East; the share and the habit never fall either. The lemma is used in every proof that follows. It is what makes the paths monotone where they are, and it is what the certificate of safety behind the figures rests on.

Lemma 1 (Orderings)**.** *(i) $\ell_{r,t}$ is nondecreasing in t for each r, and $\ell_{W,t} \geq \ell_{E,t}$ for all t. (ii) π_t is nondecreasing in t. (iii) h_t is nondecreasing in t.*

Proof. (i) The map $\ell \mapsto \ell + (1-\ell)(\kappa\pi[\alpha + (1-\alpha)\ell] + s)$ is nondecreasing in ℓ when $\kappa\pi + s \leq 1$ (its derivative is at least $1 - \kappa\pi - s$) and increasing in α. Both regions start at zero and face the same π_t, so by induction $\ell_{W,t+1} = f(\ell_{W,t}; \alpha_W) \geq f(\ell_{E,t}; \alpha_W) \geq f(\ell_{E,t}; \alpha_E) = \ell_{E,t+1}$. Monotonicity in t is immediate from (6). (ii) and (iii) together, by induction: if $\ell_t \geq \ell_{t-1}$ componentwise and $h_t \geq h_{t-1}$, then $\pi_t = \pi^*(\ell_t, \psi_0 + h_t) \geq \pi^*(\ell_{t-1}, \psi_0 + h_{t-1}) = \pi_{t-1}$ by Proposition 1, and then $h_{t+1} - h_t = (1 - \zeta)(h_t - h_{t-1}) + \eta(\pi_t - \pi_{t-1}) \geq 0$. The base case is $h_1 = \eta\pi_0 \geq 0 = h_0$ and $\ell_1 \geq \ell_0$. □

The share never falls. Each season the occupant finds a more literate audience and a fonder taste than the season before, and shows at least as much.

4 Three Paths to Downfall

An elite using its own closed codes in what it puts on the screen is susceptible to downfall. Below I chart three paths that gave this paper its subtitle. For now, keep in mind that these paths are derived with the elite being myopic, as discussed above. I will later revisit how these paths may or may not survive with a more forward-looking Lena.

4.1 Slowly, or if lucky, never

Proposition 2 (No leakage, no infatuation)**.** *Let $\kappa = s = \eta = 0$. The share is constant at $\pi^* = \pi^*(\ell_0, \psi_0)$, each $S_{r,t}$ rises monotonically to $a_r(\pi^*)/\theta$, approval falls monotonically, and the elite is eventually ousted if and only if*

$$\beta + m_A \rho\, \pi^* \;<\; \varepsilon \sum_r \mu_r (1-\alpha_r)(1-\ell_{r,0})\, \frac{a_r(\pi^*)}{\theta}. \tag{7}$$

Proof. With literacy and taste fixed the share is fixed by Proposition 1. Then $S_{r,t+1} - S_{r,t} = a_r(\pi^*) - \theta S_{r,t}$, which is positive while $S_{r,t} < a_r(\pi^*)/\theta$ and contracts towards it. Approval is decreasing in S and converges to $\frac{1}{2}$ plus the left side of (7) minus the right. It crosses one half at a finite date if and only if the limit is below one half. □

Condition (7) compares two things. On the left is what the elite has in hand: its baseline majority β plus what its own milieu gives it for a share π^* of own code, $m_A \rho \pi^*$. On the right is what the excluded will eventually hold against it: the illiterate mass of each region, times the rate at which an illiterate citizen has a bad season at that share, times the number of bad seasons she remembers, $1/\theta$, times what each of them costs, ε. The stock on the right is not there on day one. It builds from zero at the rate θ, so that half of the eventual grievance is in place after about $\ln 2/\theta$ seasons, nine seasons at the calibration, and nine tenths of it after about thirty. That is what makes this path slow. An elite that satisfies (7) has satisfied it from the first season, while its approval is still comfortably above one half. Nothing changes along the way; the grievances simply pile up at the rate the memory allows, and the elite finds out a season at a time what was decided at the start.

The same condition says when the elite never falls. It fails when the excluded are few, when the memory is short or the weight of a grievance light, when the margin is wide, when the elite's taste for its own code is modest enough to keep π^* low, or when the code is open enough, with φ close to λ, that an illiterate citizen seldom has a bad season. On this path nothing in the audience changes, and nothing in the elite. The only thing that happens is that the stock of grievances fills up to its level, and approval sinks to where condition (7) said it would.

4.2 Faster, through too much infatuation

Proposition 3 (Infatuation)**.** *Let $\kappa = s = 0$ and $\eta > 0$, and let the initial share $\pi_0 = \pi^*(\ell_0, \psi_0)$ be interior. The share rises monotonically to $\bar{\pi} > \pi_0$, the smallest solution of*

$$\pi = \pi^*\big(\ell_0,\ \psi_0 + \eta\pi/\zeta\big)$$

that is at least π_0. Approval is nonincreasing along the path, and the elite is eventually ousted if and only if (7) *holds with $\bar{\pi}$ in place of π^*. Let β_0 and $\bar{\beta}$ be the margins at which* (7) *holds with equality at π_0 and at $\bar{\pi}$; then $\bar{\beta} > \beta_0$. The maintained voting bounds admit margins in $[\max\{0, \varepsilon/\theta - \frac{1}{2}\},\ \frac{1}{2} - \rho]$. Every admissible margin in $[\beta_0, \bar{\beta})$ is one at which the elite is safe under its initial taste and ousted under the taste it acquires; the set is nonempty whenever $\bar{\beta}$ exceeds the lower end of the admissible range and β_0 falls short of its upper end.*

Proof. By Lemma 1, π_t and h_t are nondecreasing and bounded, so they converge, to $\bar{\pi}$ and $\bar{h} = \eta\bar{\pi}/\zeta$, and continuity of π^* in ψ gives $\bar{\pi} = \pi^*(\ell_0, \psi_0 + \eta\bar{\pi}/\zeta)$. A monotone sequence starting at π_0 cannot pass a fixed point of a monotone map, so $\bar{\pi}$ is the smallest fixed point at or above π_0. For monotonicity of approval, write $M = m_A\rho$ and $C(\pi) = \sum_r c_r a_r(\pi)$, both fixed since literacy is. Approval in season t is $V_t = \frac{1}{2} + \beta + M\pi_t - C(\pi_t) - (1-\theta)\sum_r c_r S_{r,t}$. The initial optimum is interior, so $C'(\pi_0) = \psi_0/R + M \geq M$, and C is convex, so $M\pi - C(\pi)$ is nonincreasing on $[\pi_0, 1]$. The share is nondecreasing by Lemma 1, and so is each stock, from zero. Both moving terms therefore weakly lower approval season by season, and V_t is nonincreasing. It converges to $\frac{1}{2}$ plus the margin minus the limit grievance term at $\bar{\pi}$, and a nonincreasing sequence falls below one half at a finite date if and only if its limit is below one half, which is (7) at $\bar{\pi}$. For the interval: $\bar{\beta} > \beta_0$ if and only if the right side of (7) rises by more than $m_A\rho(\bar{\pi} - \pi_0)$ between π_0 and $\bar{\pi}$. Since a_r is convex, $a_r(\bar{\pi}) - a_r(\pi_0) \geq a_r'(\pi_0)(\bar{\pi} - \pi_0)$, and at an interior π_0 the first-order condition gives $\sum_r c_r a_r'(\pi_0) = \psi_0/R + m_A\rho$, so the right side rises by at least $(\psi_0/R + m_A\rho)(\bar{\pi} - \pi_0)/\theta$, which exceeds $m_A\rho(\bar{\pi} - \pi_0)$ whenever $\theta < 1$ and $\bar{\pi} > \pi_0$. The latter holds because π_0 is interior and $\eta > 0$, so that $h_1 > 0$ and π^* is strictly increasing in ψ at π_0. The intersection with the admissible margins is needed because the interval can lie outside them. □

This is what happens if Lena's love for her football jokes grows too fast, and it is what the *Tatort* story is about. The elite shows a little more each season because Lena likes the code a little more each season. Approval drifts down a little each season, and for a while the infatuation is worth it to her. But, finally, the line is crossed, by an elite that would have been safe under the taste it started with. The proof also says why the interval of margins is never empty, that is, why there are always margins at which the initial taste is safe and the acquired taste is not. When the taste rises, the elite shows more. Each extra unit of own code costs her, in this season's grievances, what she set it equal to at the start, her marginal benefit. But the audience remembers each grievance for $1/\theta$ seasons, so what

the acquired taste eats out of the margin is that marginal cost multiplied by the length of memory. A long memory turns a small infatuation into a lost majority.

4.3 Even faster, in the valley of hope

The third path is different in kind. On the first two, approval falls and stays down. On this one it falls and would come back, if the elite were still there to see it. The reason is leakage. When the code leaks, the illiterate become literate, and a literate citizen carries no grievance: the mass that weighs the stock in (3) shrinks towards zero, and with it the whole grievance term, whatever the stock per illiterate citizen has become. In the long run everyone reads the code, nobody counts a bad season as exclusion, and approval settles at $\frac{1}{2} + \beta + m_A \rho \pi_\infty$, above one half. Proposition 4 says so. The danger is on the way there. Leakage reaches the two regions at different speeds, the elite raises the share on the region that learns first, and the region that learns last collects grievances at the higher share before its own learning arrives. Approval dips, bottoms out, and recovers: a valley. Whether the bottom lies below one half depends on the parameters, and Section 4.4 shows which do and which do not. This is also why the figures there continue every path past the season of ousting. The recovery that follows the valley is what an elite that survives the crisis, because the contest is not held every season or is not decided by approval alone, would live to enjoy, and it is what the elite loses when it does not.

Proposition 4 (Leakage: the long run)**.** *Let $\pi_0 > 0$, and let each region have a learning seed: $s > 0$, or $\kappa > 0$ with $\alpha_r > 0$. Then $\ell_{r,t} \to 1$ for both regions, the share converges to some $\pi_\infty \geq \pi_0$, and approval converges to $\frac{1}{2} + \beta + m_A \rho\, \pi_\infty$, which exceeds one half. The elite, if it survives, ends safe.*

Proof. By Lemma 1, $\pi_t \geq \pi_0 > 0$, so (6) gives $\ell_{r,t+1} - \ell_{r,t} \geq (1 - \ell_{r,t})\big(\kappa \pi_0 \alpha_r + s\big)$, and the bracket is positive by the seed condition; so $1 - \ell_{r,t}$ decays at least geometrically to zero. The weight $\mu_r(1 - \alpha_r)(1 - \ell_{r,t})$ of the illiterate in (3) goes to zero, and $S_{r,t}$ is bounded by $1/\theta$, so the grievance term vanishes. The share is monotone and bounded, so it converges. Approval converges to $\frac{1}{2} + \beta + m_A \rho \pi_\infty > \frac{1}{2}$. □

So with leakage the long run is safe. The question is the way there.

Claim 1 (The valley). Let $\kappa > 0$, $s = 0$, and $\alpha_W > \alpha_E > 0$. Then $S_{E,t} \geq S_{W,t}$ at every date: the East carries at least as many unforgotten bad seasons per illiterate citizen as the West, always.

This follows from Lemma 1: $\ell_W \geq \ell_E$ gives $\nu_E \geq \nu_W$, hence $a_E(\pi_t) \geq a_W(\pi_t)$ at every date, and the stocks start equal. It is the only part of the valley that is proved. That the East's stock peaks later and higher than the West's, and higher than it would if the West did not learn, and that approval passes through a minimum, is shown in Section 4.4, and

it is not general: if the share is at the corner from the start, the West's learning changes neither the share nor the East's path. The ordering of the peaks is a race that I have not closed.

The mechanism is Proposition 1 applied season after season, with Lemma 1 keeping everything moving in one direction. The West has many citizens who know the code, so its illiterate citizens learn fast, and each citizen who learns stops counting in the elite's cost of showing her code. The elite sees a cheaper country and raises the share to what the citizens who are still illiterate can bear. The East has few who know the code. It learns little for a long time, and it faces a rising share that was set for a West which is learning. The East's bad seasons become more frequent as the share rises, its stock of grievances climbs, and only when the East's own learning catches up does the stock come down again. The East does not face a different screen from the West. It faces the screen that a largely literate West can bear, before it is literate itself.

4.4 The three paths in two pictures

The propositions are general, but the numbers that accompany them come from one calibration, which I fix here and use throughout: $\lambda = 0.2$, $\varphi = 0.9$, $m = 2$, $\mu_W = 0.7$, $\mu_E = 0.3$, $\alpha_W = 0.6$, $\alpha_E = 0.05$, $\omega = 0$, $\beta = 0.05$, $\rho = 0.012$, $\varepsilon = 0.044$, $\psi_0 = 0.003$, $R = 1$, $\zeta = 0.03$, $\theta = 0.08$ and $\ell_{r,0} = 0$; so $\varepsilon/\theta = 0.55 = \frac{1}{2} + \beta$, the longest memory the voting bound allows at this weight, and $m_A = 0.435$. The two rates that the three paths turn on, infatuation η and leakage κ, are set path by path.

Figure 1 runs the model for three hundred seasons under four settings and shows, for each, four things. The top left panel is the share of own code the elite chooses each season, π_t. The top right is approval, V_t, with the line at one half. The bottom left is literacy, the fraction of citizens outside milieu A who know code a, with the West solid and the East dashed. The bottom right is the stock of grievances, the number of unforgotten bad seasons per illiterate citizen, East solid and West dotted. Every path is continued past the season in which approval first falls below one half, as if the elite were still there, for the reason given in Section 4.3.

The four settings are the three paths and their combination. With no infatuation and no leakage, $\eta = \kappa = 0$, the share is constant at 0.41 and approval falls monotonically to 0.528, which is Proposition 2 with condition (7) not met: the stock builds, and the elite is never in danger. With infatuation alone, $\eta = 0.0004$, the share rises from 0.41 to one within a hundred seasons and approval falls to 0.43 for good, with the East at more than seven unforgotten bad seasons per illiterate citizen; the majority is lost in season 39. This is Proposition 3, and it is the *Tatort* path. With leakage alone, $\kappa = 0.03$, literacy rises, the West far ahead of the East, the share rises with it, and approval digs a valley that bottoms at 0.51 in season 108 and comes back to 0.56. The elite survives

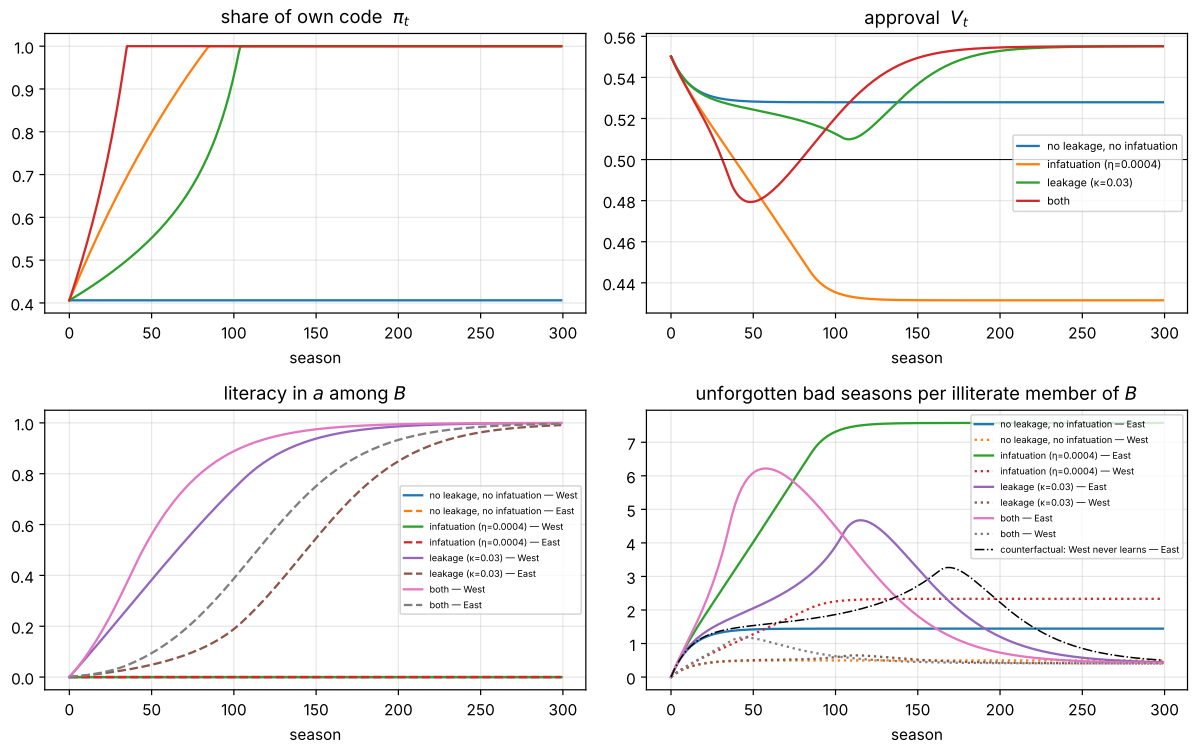


Figure 1: Four paths at the calibration of Section 4.4, continued past the season of ousting as if the elite had survived. Baseline: $\eta = \kappa = 0$. Infatuation: $\eta = 0.0004$. Leakage: $\kappa = 0.03$. Both. Bottom right also shows the East when the West's literacy is frozen at zero.

it. The East's stock peaks at 4.7 bad seasons in season 115, long after the West's, and the dash-dotted line in the bottom right panel is the counterfactual in which the West never learns: then the East peaks at 3.3. The difference between the two is the valley mechanism of Section 4.3, the share rising on the West's learning and the East paying for it. With both together, $\eta = 0.0004$ and $\kappa = 0.03$, the taste pushes the share up ahead of the audience's learning, and the valley that leakage alone would have survived becomes one it does not: the majority is lost in season 32, approval bottoms at 0.479 in season 48, and then recovers, to 0.555, for an elite that is no longer there.

The interval of margins in Proposition 3 can be checked against the calibration, and its intersection with the admissible margins is not a formality. With $\eta = 0.0001$ and no learning, $\pi_0 = 0.406$, $\bar{\pi} = 0.502$, $\beta_0 = 0.022$ and $\bar{\beta} = 0.032$, while admissible margins start at $\varepsilon/\theta - \frac{1}{2} = 0.05$: the share rises, but every admissible elite is safe. At $\eta = 0.0004$ the share goes to one, $\bar{\beta} = 0.119$, and every margin in $[0.05, 0.119)$ is an elite that is safe under its initial taste and ousted under its acquired one. The calibration's $\beta = 0.05$ is one of them, which is why the infatuation path of the figure crosses the line.

Figure 2 is the same exercise over the whole plane of infatuation and leakage. Each cell is a run; grey cells are elites that never lose their majority, and the colour of the others is the season in which they do. The safe region is the corner of small infatuation. The band

along the bottom, no leakage and some infatuation, is Proposition 3: there the elite is ousted for good, after many seasons near the edge of the safe region and after few further inside. Everything above the band is the valley: an elite that leakage would have made safe in the long run, but that did not last long enough to get there. And the frontier between safe and ousted bends. The safe region widens as leakage rises, and at small infatuation the elite is safe with no leakage, ousted with a little, and safe again with a lot. That is the subject of the next subsection.

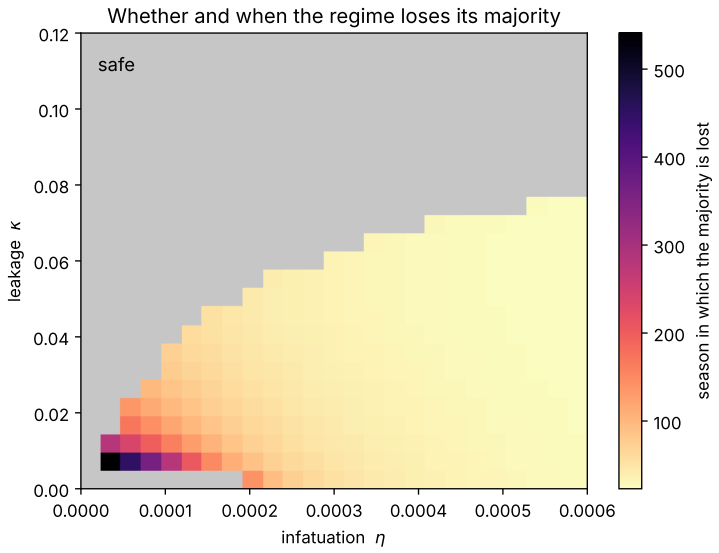


Figure 2: Whether and when the elite loses its majority, over infatuation η and leakage κ; other parameters at the calibration. Grey: approval never falls below one half, certified by a lower bound on all later approval, $L_t = \frac{1}{2} + \beta + m_A \rho \pi_t - \varepsilon \sum_r \mu_r (1 - \alpha_r)(1 - \ell_{r,t}) \max\{S_{r,t+1}, a_r(\pi^{\sup})/\theta\}$, where $\pi^{\sup}$ is the share at the limiting literacy and taste and bounds all later shares by Lemma 1; or, with no learning, by the exact steady state. Colour: the season in which the majority is lost.

4.5 Comparative statics

The three paths depend on the parameters, and this subsection examines how. The first part is the comparative static the paper is most about, learning, where the answer is a surprise, and in one limit a theorem. The second takes the remaining parameters in turn, memory first, since it governs the other two paths as well.

4.5.1 Learning, and what the elite does with it

Leakage, κ, is how often a kid at a mixed table learns what a nutmeg is; schooling, s, is how often the teacher explains it to the whole class. At a fixed share of the own code on the screen, learning of either kind is unambiguously good for the elite. Every term in a

region's bad-season rate falls as its literacy rises, so fewer citizens have bad seasons, the stock is smaller, and approval is higher. The dashed line in the left panel of Figure 3 is this benchmark. With the share held at its initial value, the minimum of approval rises monotonically in the rate of leakage, from 0.528 to 0.537.

The solid line is what happens when the elite responds. Learning makes it cheaper to use the own code, by Proposition 1, and the commander of the screen can afford to use more of it. There are more who enjoy it. In the limit of slow learning this leads to a theorem. Let $\kappa > 0$ and write literacy time $\sigma = \kappa t$. Without habit or schooling the share depends on the state only through literacy, so as $\kappa \to 0$ literacy follows

$$\frac{d\ell_r}{d\sigma} = (1-\ell_r)\,\pi(\ell)\,k_r, \qquad \pi(\ell) = \pi^*(\ell,\psi_0), \qquad \ell(0) = \ell_0, \tag{8}$$

and the stocks, which settle in a few multiples of $1/\theta$ seasons, are always fully built for the current share. The approval that goes with that is

$$\widetilde{V}(\sigma) = \tfrac{1}{2} + \beta + m_A\rho\,\pi(\ell(\sigma)) - \frac{\varepsilon}{\theta}\sum_r \mu_r(1-\alpha_r)\big(1-\ell_r(\sigma)\big)\,a_r\big(\pi(\ell(\sigma))\big), \tag{9}$$

with a_r evaluated at the literacy $\ell_r(\sigma)$.

Proposition 5 (A little learning)**.** *Let $\eta = s = 0$, let $\pi_0 = \pi^*(\ell_0,\psi_0) > 0$, let each region have a learning seed ($\alpha_r > 0$ or $\ell_{r,0} > 0$), and let $\kappa > 0$ be the rate of leakage. (i) As $\kappa \to 0$, the lowest approval along the path converges to $\inf_{\sigma\ge 0}\widetilde{V}(\sigma)$. (ii) $\widetilde{V}(0)$ is the limit of approval without leakage, so that*

$$\lim_{\kappa\to 0^+}\ \min_t V_t \;=\; \inf_{\sigma\ge 0}\widetilde{V}(\sigma) \;\le\; \widetilde{V}(0) \;=\; \inf_t V_t\big|_{\kappa=0}\,:$$

in the limit of slow learning a little leakage can never do the elite more good than none. The inequality is strict whenever $\widetilde{V}(\sigma) < \widetilde{V}(0)$ for some σ; and $\widetilde{V}'(0) < 0$, with the share interior at $\sigma = 0$, if and only if the elite's response to the first citizens who learn the code outweighs the relief that their learning brings,

$$\pi'(0)\left[\frac{\psi_0}{R} + (1-\theta)\,m_A\rho\right] \;>\; \varepsilon\sum_r \mu_r(1-\alpha_r)\,\ell_r'(0)\left[a_r + (1-\ell_r)\left|\frac{\partial a_r}{\partial \ell_r}\right|\right], \tag{10}$$

where a_r, $1-\ell_r$ and $\partial a_r/\partial\ell_r$ are evaluated at $(\pi_0,\ell_{r,0})$, $\ell_r'(0) = (1-\ell_{r,0})\pi_0 k_r$ and $\pi'(0) = \sum_r(\partial\pi^*/\partial\ell_r)\,\ell_r'(0)$. (iii) The elite is safe without leakage if and only if $\widetilde{V}(0) \ge \frac{1}{2}$. It is ousted at every sufficiently small positive rate of leakage if $\inf_\sigma \widetilde{V} < \frac{1}{2}$, and safe at every sufficiently small positive rate if $\inf_\sigma \widetilde{V} > \frac{1}{2}$.

Proof. Write $\Phi(\ell)$ for the right side of (8). The share $\pi^*(\ell,\psi_0)$ is Lipschitz in ℓ: it is the root of $G = \sum_r c_r a_r'(\pi)$, whose π-derivative is bounded away from zero wherever $\sum_r c_r$ is, and it equals one on a neighbourhood of $\ell = (1,1)$, where the cost vanishes. So Φ is Lipschitz, and (6) with $s = 0$ is Euler's method for (8) with step κ: for every T, $\sup_{\kappa t\le T}|\ell_t - \ell(\kappa t)| \to 0$

as $\kappa \to 0$, and with it the share and the rates a_r. Write $A_r(\sigma)$ for a_r along the literacy-time path, a Lipschitz function. The stock (4) is $S_{r,t+1} = \sum_{j \le t} (1-\theta)^{t-j} a_r(\pi_j)$, so

$$\begin{aligned}\left| S_{r,t+1} - \frac{1-(1-\theta)^{t+1}}{\theta} A_r(\kappa t) \right| &\le \sum_{i=0}^{t} (1-\theta)^i \left| A_r(\kappa(t-i)) - A_r(\kappa t) \right| + o(1) \\ &\le \frac{\kappa(1-\theta)}{\theta^2} \operatorname{Lip}(A_r) + o(1)\end{aligned}$$

uniformly on $\kappa t \le T$. Hence

$$V_t = \widetilde{V}(\kappa t) + (1-\theta)^{t+1} \frac{1}{\theta} \sum_r c_{r,t} A_r(\kappa t) + o(1)$$

uniformly on $\kappa t \le T$, where $c_{r,t}$ is c_r at $\ell_{r,t}$: up to the lag of the stocks, which vanishes with κ, the discrete path sits above the literacy-time path by the part of the stock that has not yet built, and that part vanishes after a few multiples of $1/\theta$ seasons. For the upper bound, take $t = \lfloor \sigma^*/\kappa \rfloor$ for any $\sigma^* > 0$; then $(1-\theta)^{t+1} \to 0$ and $V_t \to \widetilde{V}(\sigma^*)$. For $\sigma^* = 0$ take $t = \lceil \kappa^{-1/2} \rceil$; then $\kappa t \to 0$ and $(1-\theta)^t \to 0$, so $V_t \to \widetilde{V}(0)$. Thus $\limsup_\kappa \min_t V_t \le \inf_\sigma \widetilde{V}$. For the lower bound, on $\kappa t \le T$ the displayed term is nonnegative, so $V_t \ge \widetilde{V}(\kappa t) - o(1) \ge \inf \widetilde{V} - o(1)$. On $\kappa t > T$, literacy is nondecreasing, $\pi_t \ge \pi_0$ by Lemma 1, and $S \le 1/\theta$, so $V_t \ge \frac{1}{2} + \beta + m_A \rho \pi_0 - (\varepsilon/\theta) \sum_r \mu_r (1-\alpha_r)(1-\ell_r(T)) - o(1)$, and $\ell_r(T) \to 1$ as $T \to \infty$ by the seed condition and $\pi(\ell(\sigma)) \ge \pi_0 > 0$, while $\inf \widetilde{V} \le \widetilde{V}(0) \le \frac{1}{2} + \beta + m_A \rho \pi_0$. Choosing T and then κ gives $\liminf_\kappa \min_t V_t \ge \inf_\sigma \widetilde{V}$, and (i) follows. For (ii), $\widetilde{V}(0)$ is $\frac{1}{2}$ plus the left side of (7) minus its right, which is the limit of Proposition 2, and without leakage approval falls to it, so $\inf_t V_t = \widetilde{V}(0)$ at $\kappa = 0$. Differentiating (9) at an interior share and using the first-order condition $\sum_r c_r \, \partial a_r / \partial \pi = G$ to replace the coefficient of π' by $m_A \rho - G/\theta = -[\psi_0/R + (1-\theta) m_A \rho]/\theta$ gives $\theta \widetilde{V}'(0)$ as the right side of (10) minus its left side. (iii) is (i) together with Proposition 2, which gives safety without leakage if and only if its limit is at least one half. □

Three observations. The first is the inequality (10), which is the paper's mechanism in one line. The first citizens who learn the code bring relief: they take their own bad seasons with them, and the illiterate who remain have better Mondays, because the person they ask is now more likely to have got the episode. The elite responds to them by raising the share, by $\pi'(0)$ per unit of literacy time. She values a unit of her own code at G and pays G for it in this season's grievances, which is what the first-order condition says; but the audience keeps the grievance for $1/\theta$ seasons, so the true cost to approval of her response is $G/\theta - m_A \rho$ per unit of share. The second is that the limit does not depend on κ at all. In the limit of slow learning the speed of leakage only sets how many seasons each step takes. The sequence of literacy levels the country passes through, the share the elite sets at each of them, and the approval that goes with it, are the same at every speed; so is the depth of the valley. What a positive speed changes is two things. The stock of grievances lags behind the share while the share is moving, by an amount of order κ/θ^2; and at the start

the stock has not yet built up, which takes about κ/θ in literacy time. Both work in the elite's favour, and both grow with the speed of learning. They are what makes the second arm of the U: when learning is fast enough, the stock never catches up with the share before the East has learned. The third is that nothing in the proof is particular to leakage. With schooling in place of it, $\eta = \kappa = 0$ and $\sigma = st$, literacy follows $d\ell_r/d\sigma = 1 - \ell_r$, no seed is needed, and (i) to (iii) hold line for line.

With the habit the taste also settles at each literacy level in the slow limit, since it runs in calendar time, and the settled share of Proposition 3 takes the place of $\pi^*(\ell, \psi_0)$. The argument goes through for every $m \geq 2$. The settled share is the root of $F(\pi) = \sum_r c_r a_r'(\pi) - \eta\pi/(\zeta R) - \psi_0/R - m_A\rho$, which is convex in π because $a_r''' \geq 0$, and negative at zero whenever the initial share is positive; so it has at most one root, crossed upward, or none, in which case the settled share is one. The root moves continuously with literacy and joins the corner continuously. The numbers below use that version.

Remark 2 (Literacy time). The literacy-time path of Proposition 5 depends on θ only through ε/θ, the settled weight of a grievance, and on ε and the recognition weights through $(\psi_0/R + m_A\rho)/\varepsilon$, which sets the share, and $m_A\rho$, which enters approval directly. Away from the limit the discrete path differs from $\widetilde{V}$ by the lag of the stocks, so θ/κ enters; with the habit, η/κ and ζ/κ; with schooling, s/κ. That is why the frontier in Figure 2 bends.

The numbers. At the calibration with $\eta = 0$, approval without leakage settles at $\widetilde{V}(0) = 0.528$. The literacy-time path falls from the start, since (10) holds there, and bottoms at 0.5003, with the West three quarters literate, the East one fifth, and the share at its corner. A little learning takes the elite from a comfortable majority to within a thousandth of losing it, and it survives. The mild habit of the cuts below, $\eta = 0.0001$, takes $\widetilde{V}(0)$ to 0.518 and the bottom to 0.478, and the elite over the line. The discrete paths converge to the limit, from above at these parameters: their minima are 0.478, 0.479, 0.479 and 0.481 at $\kappa = 0.0001$, 0.0003, 0.001 and 0.003. The open circle in the left panel of Figure 3 is the limit, and the solid line is the cut. With no leakage the share rises only with the habit, and settles low enough that even the fully built stock leaves the elite safe, at 0.518. With a little leakage the share rises slowly enough that the East's stock is fully built when the share reaches the corner, and the East is still illiterate; the majority is lost, late, in season 439 at $\kappa = 0.003$. With a lot of leakage the share reaches the corner before the stock has had time to build, and the East begins learning before it finishes building; the valley is shallow, 0.520 at $\kappa = 0.12$, and the elite survives. Without the habit, the share reaches the corner at the same literacy levels whatever the speed of leakage; with the habit the route drifts, since the taste grows in calendar time, and in the cut the West's literacy when the share reaches one is 0.60 at $\kappa = 0.003$ and 0.75 at $\kappa = 0.12$. The dotted line is the same cut with the two regions given the same composition. Minimum approval is still non-monotone in leakage, so the U comes from the elite's response and not from the two regions learning at different speeds. The geography adds something else. At these parameters a country

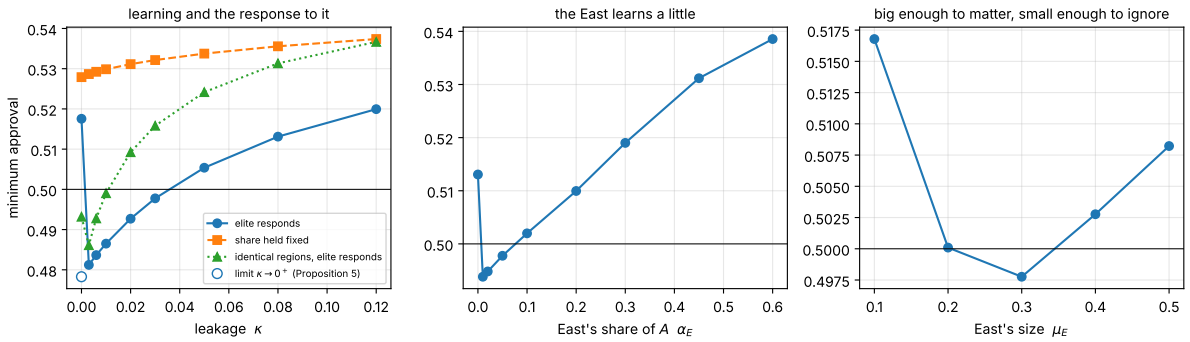


Figure 3: Minimum approval along the path. Baseline $\eta = 0.0001$, $\kappa = 0.03$, other parameters at the calibration; at the baseline the majority is lost in season 77. Left: leakage, with the elite responding (solid), with the share held at its initial value (dashed), and with the two regions given the same composition (dotted); the open circle is the limit $\kappa \to 0^+$ of Proposition 5 with the taste settled. Middle: the East's share of the elite's milieu. Right: the East's population share.

with identical regions has already lost its majority before any leakage, so the pattern of safe, then ousted, then safe again needs the two regions to differ. So the worst amount of leakage is a little. A little is enough to make the elite raise the share, because the West is learning, and not enough to make the East learn before its grievances have built.

The same result shows up in three more places, in the other panels of the figure and in the text, and these are shown, not proved. Schooling, which teaches both regions at the same rate per illiterate citizen, shows the same U when nothing else leaks: minimum approval is 0.518 with no schools, 0.483 at a schooling rate of 0.001 with the majority lost in season 144, and 0.534 at 0.05 with the elite safe. With leakage already present at the baseline the elite is already responding to the West's learning, and schooling then mainly speeds up the East; in that cut schools only help. So half-done schooling is dangerous where nothing else leaks. The East's share of the elite's milieu, α_E, shows it too. With $\alpha_E = 0$ the East never learns and its illiterate mass keeps the elite's marginal cost high. With α_E between 0.01 and 0.05 the East begins to learn, the share goes to one, and the East's grievances build before its learning arrives, so the majority is lost. From 0.1 the East learns fast enough to be safe. A region protects itself by being unteachable. And the East's size, μ_E: a small East costs little in the share and weighs little in the vote; a large East forces restraint; an East of three tenths of the country is big enough to matter and small enough to be ignored in the share, and that is where the elite falls.

4.5.2 Memory, closedness, and the rest

Memory plays an important role. It matters because a bad season is not cancelled by a good one: an illiterate citizen carries her bad seasons with her, each fading at the rate θ, so a share that the elite raises today is paid for over as many seasons as the audience

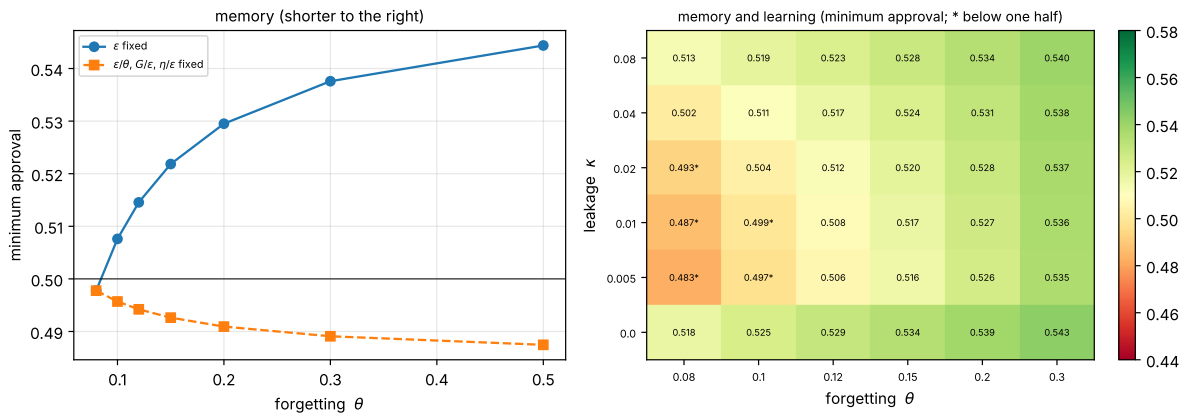


Figure 4: Memory. Left: minimum approval against forgetting θ at the baseline of Figure 3, with ε fixed (circles) and with ε/θ, G/ε and η/ε fixed (squares); the voting bound $\varepsilon/\theta \leq \frac{1}{2} + \beta$ allows no longer memory than $\theta = 0.08$ at the baseline weight. Right: minimum approval over forgetting and leakage; starred cells are below one half.

remembers. In the fable, θ is how fast a horse kid forgets a bad term. (Whether a grievance survives learning the code is a different question, and Section 7 checks it.) In the slow limit θ enters the path only through ε/θ: how long a grievance is remembered matters as a multiplier of what it weighs, and not as a speed. The left panel of Figure 4 shows two cuts. With ε fixed, minimum approval rises from 0.498 at $\theta = 0.08$, the longest memory the voting bound allows at this weight, to 0.544 at $\theta = 0.5$. With ε/θ held fixed, and G/ε and η/ε with it, so that only the speed moves, it falls from 0.498 to 0.487. At a fixed settled weight, faster forgetting is worse for the elite: the stock then tracks the share more closely, and the lag that fills the valley at a positive speed of learning is lost. So all of the improvement in the first cut, and a little more, comes from the weight of a grievance falling, not from the speed at which grievances are forgotten. Memory is also what makes the other two paths dangerous. In Proposition 3 the margin that the acquired taste eats is multiplied by $1/\theta$; in Proposition 5 the cost of the elite's response to learning is multiplied by $1/\theta$. In both cases the reason is the same: the elite pays for a unit of her own code once, in this season's grievances, and the audience charges her for it $1/\theta$ times. The grid in the right panel shows memory against leakage: longer memory lowers approval at every rate of leakage, and the U of the previous subsection is visible along every row, deepest at the longest memory. Long memory makes both of the other paths more dangerous, and it does so through the same stock of grievances in both.

Closedness and infatuation reinforce each other too. Closedness, φ, is the probability that someone who does not know the code misses an episode in it; in the fable it is how much of a football joke a horse kid misses, and at $\varphi = 1$ the joke is simply a blank to her. Infatuation, η, is how much fonder Lena gets of football jokes each time she puts one up. Figure 5 shows minimum approval over the two. A code with $\varphi \leq 0.7$ is safe at any infatuation in the grid; at $\varphi = 0.9$ the majority is lost at any infatuation of 0.0001 or more. Along the φ axis alone the effect is monotone up to 0.9, and then turns: a code

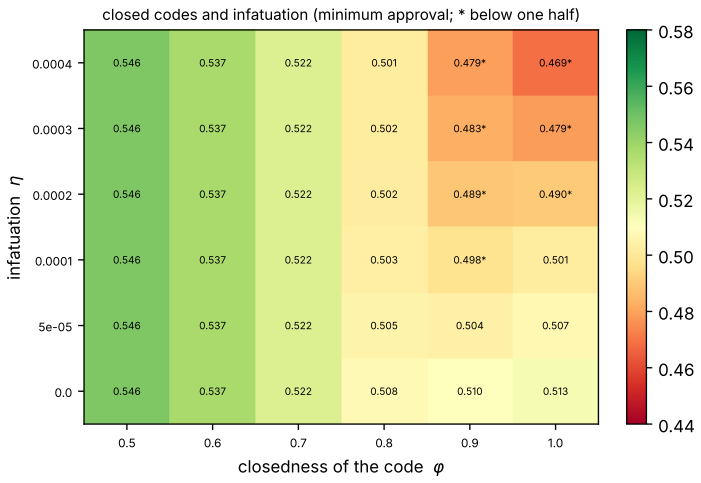


Figure 5: Closedness and infatuation: minimum approval at the baseline of Figure 3; starred cells are below one half.

closed enough is shown so sparingly that at $\varphi = 1$ and mild infatuation the elite is safe again, at 0.501. A very closed code is dangerous mainly in the hands of an elite that has grown fond of it.

The remaining parameters are monotone in the cuts, and one sentence each will do; these are illustrated patterns, not theorems. Tolerance protects. In the fable, m is how many bad Mondays make a bad term; when it takes three rather than two, $m = 3$, the baseline elite is safe. Asking n people instead of one puts the exponent n on F_r and λ in (1) and turns k_r in (6) into $1 - (1 - k_r)^n$; more acquaintances protect, through both channels, since all of them must have missed the episode for a Monday to be bad and any of them can teach the code. A habit that fades fast is harmless and one that persists is not: ζ is how quickly Lena's fondness wears off when she stops putting football jokes up, and the slower it wears off, the more dangerous the infatuation. A stronger habit is worse at any rate of fading. National contact fills the valley.

5 The Screen Changes Hands

The challenger has no move that reaches citizens. Only what is on the screen is got or missed, so what she shows at her rallies is irrelevant to votes. That is all the freedom she has. She can show code b to her own milieu at the rally, for recognition, and nobody who misses it is thereby excluded, because a rally does not come with the power to exclude. The asymmetry of the game is not between the milieus. It is that one of them holds the screen and the other does not.

When the elite is ousted, the leader of B holds the screen and chooses each season the share of code b, with the rest in code c. Code a leaves the screen by construction, since a leader

shows her own code or the common one. So what looks like hostility to the predecessor's code is, in the model, only the shape of the action set; Section 7 names the extension in which keeping or dropping the predecessor's code becomes a choice. Everything in Sections 3 and 4 applies with the labels exchanged. The excluded are now the citizens of milieu A and of C together, with regional mass $\mu_r(\alpha_r + \gamma_r)$; literacy is literacy in code b among them, starting at zero everywhere, since literacy in code a is inherited but enters nothing; the new elite's habit starts at zero; and the three paths to the downfall are open to her.

Three things follow. The first two are the exchange of labels and nothing more. The direction of resentment follows the screen: the citizens who hold grievances are those who do not hold the screen and have not yet learned its code. And what is shown follows the screen: if modernism is in code a it is shown while A holds the screen, provided the occupant's optimum is positive, and not after; if it is in code b, after and not before. The third is arithmetic. The mass exposed to exclusion under A was $1 - m_A$ at the start; under the successor it is $m_A + m_C$, which is larger if and only if $m_C > 1 - 2m_A$, a threshold of 0.13 at the calibration. When C is large the successor inherits a bigger problem than the one that undid her predecessor. The citizens of C are the audience of both elites: neither shows them their own code, since they have none, and under each of them they start counting bad seasons afresh. A large common-code group can therefore carry successive changes of elite, provided each incumbent builds enough grievances before its code has diffused. It does not make alternation permanent. With learning that persists across changeovers, each incumbency teaches its code to some of the audience, and after enough seasons in office the exposed mass is small enough that an elite with any positive margin is safe; so repeated turnover needs something the model does not have, such as codes that are renewed or literacy that is forgotten. Those are named in Section 7.

One thing follows from the dynamics and not from the labels. What the successor inherits is the geography, and the geography can flip. The region where the incumbent milieu is thin, and the excluded thick, is now the West, provided B is thicker in the East than in the West, $\alpha_W + \gamma_W > \alpha_E + \gamma_E$; the calibration fixes only α, and the condition holds whenever the common-code group is spread evenly. Under that condition it is the West that has few to learn from and the East that learns fast, and the region with the larger stock of grievances per illiterate citizen changes sides with the screen, by Lemma 1. Whether the successor's valley is dug the same way as the predecessor's, by the share rising on the other region's learning, the relabelling does not say. Under the new screen the excluded, before C is added, are 0.42 of the West and only 0.015 of the East, so the East's learning takes far less off the elite's cost than the West's learning did under the old screen, and the mechanism of Section 4.3 has less to work with. What is established is a conditional prediction, and it is testable: after a cultural changeover between milieus concentrated in different regions, the geography of grievance should invert.

This is the model's answer to Italy and to Dessau. In Weimar Germany modernism was in the code of the milieu that held the screen; the Bauhaus in Dessau was a municipal

institution of that milieu. The challenger who won in 1932 did what the model says a challenger does: she showed her own code, and modernism left the screen. In Italy modernism was in the challenger's code; futurism was part of what the fascists were before they held anything, and when the screen changed hands in 1922, modernism went onto it. The Dessau closure of 1932 was a decision of the city council, on a National Socialist motion; the Berlin successor closed in 1933. If an AfD government takes office in Sachsen-Anhalt, it is a leader of B who has taken a regional screen, and the model says the Bauhaus leaves it, because a leader shows her own code or the common one. In the model the Bauhaus simply stops being shown, because the action set has no room for it; the model derives no wish to suppress it, and it does not explain why an elite would attack an institution rather than stop showing it. Nor does it say which code modernism was in at any of these dates. That assignment is historical, and the Italian case in particular was not homogeneous. What the model says is that once the assignment is made, what is shown and who is aggrieved follow from who holds the screen.

What is proved and what is shown. Proved: the season's unique solution and its comparative statics (Proposition 1); the affine cost of the knowing Bayesian (Remark 1); the orderings along the path (Lemma 1); the slow ousting condition (Proposition 2); the infatuation fixed point and the interval of margins (Proposition 3); the safe long run under leakage (Proposition 4); that in the limit of slow learning a little leakage never raises the elite's lowest approval above what none gives, with the condition under which it lowers it (Proposition 5); and the role exchange, which is a relabelling. Shown, with figures and the mechanism: the valley (Claim 1), the second arm of the U, safety at fast learning, the other three faces of the learning result, and the complementarities of memory with infatuation and with slow learning. Assumed, and defended in the next section: that candidates come from milieus. Computed, for one calibration: the forward-looking occupant's path (Section 7). Noted: the rest of Section 7.

6 Discussion

In this section I offer some interpretation of, and thoughts about, the results above, clustered around seven themes.

Quality. The model says nothing about whether what is shown is good in any objective sense, indeed, it does not even say anything about quality in a subjective sense—other than that those who do not understand a used code will presumably find a show poor. Standing the test of time might be seen an indicator for quality: Wagner has survived a hundred and fifty years. Whether that is due to an innate quality or simply a sufficient supply of people who were able to read the code generation after generation is outside the

scope of this paper. The citizen's question in this paper is not whether the code is good. It is about understanding it.

The centrist. The model has no candidate who knows the common code only, although she would be a formidable one. With no own code to show she would never exclude, and if citizens knew that of her, so that a bad season under her screen could only be bad luck, she would carry no grievance term and sit at $\frac{1}{2} + \beta$ for ever. Such a leader is, however, absent by assumption, and the assumption is about the supply of candidates, not about potential popularity: a code is what makes a group a group. The football kids have a table, their jokes, and they have Lena; the horse kids have Paul. Those who only understand the common code have no joke of their own and are not perceived as a group. The proverb "a friend to all is a friend to none", commonly traced back to Aristotle's *Nicomachean Ethics*, captures some of this. In Bourdieu's (1985) language they can be classified but will not mobilise. So I defend the assumption for the present paper, where, quite clearly, both the current elite and the right-wing challengers have codes of their own.

One role or two. The model has one role where the world has two: the editor who gave the inspector an Eames chair and the minister who sanctions it. My assumption to have them as one is one about attribution. A disgruntled citizen may not be aware of the editor's existence, and she does not blame the minister for the chair; she simply concludes that the people who chose the chair and the people in office are the same kind of people, and votes accordingly. This paper's elite is a milieu that fills both roles. One final remark: the splitting of the roles in real life might be seen as a justification for the myopia assumption, with politicians in power for longer terms and editors who change.

Nobody is more cultivated than others. Yes, in the model there are people who can read two codes and there are others who can only read one. But there is no "pure gaze", and the opposition has its own code that will not be read by those who are fluent in Bauhaus and mid-century design. They have their dress codes, number codes and brand codes that the elite cannot read (and is not meant to). The professor's bafflement at a sweatshirt is the *Schrankwand* owner's bafflement at an Eames chair, in the other direction.

Infatuation. The habit stock is the one assumption in the model about the elite that is not about incentives, and it should be defended rather than hidden. It is Bourdieu's own claim about the habitus: use makes the code one's own. It is also what may become, more and more, a house style. The myopic elite who does not see that she is feeding her taste is a projection-bias assumption in the sense of Loewenstein, O'Donoghue and Rabin (2003), and whether a forward-looking elite would restrain herself is a question to be tackled in Section 7.

Majorities. Both milieus can be large or small, and those with only the common code might be the majority. The analysis shows that the downfall of the elite can even happen when it starts from a relatively safe position, though. A milieu of four fifths of the country loses its majority on a thin margin when the excluded fifth is concentrated, remembers long, and weighs its grievances heavily, which is the condition of Proposition 2 with $1 - m_A$ small and a_r large.

Schools. It is ever so tempting to think that cultural education—teaching codes at school—is in the best interest of everybody, in particular in the best interest of the elite. The model provides a caveat, though. Figure 3 shows the peculiarity. A little schooling can be worse than none—as it may render the elite just a tad too comfortable with showing its own code ever more. The upshot is simple and stark: cultural education, yes, but not half-assed.

7 Extensions

A forward-looking occupant. The one in power in this paper sets this season's share of shows in the closed code. One who valued the role's future would pay close attention to three things: that her share feeds her own taste (or that of her successor, in an overlapping-generations reading), that it feeds the stock of grievances, and that it teaches her audience. The last is an investment, since a literate audience is cheaper to please. The second is a debt, since the grievances will be paid for later. The first is both: a fonder future self enjoys the code more, and also shows more of it and pays a higher bill.

The first of these is Becker and Murphy's (1988) point against Pollak (1970), for the case they call beneficial consumption capital: an addict who foresees the habit values it, since a fonder future self enjoys every unit more, and consumes more rather than less. It holds here. So the infatuation path of Section 4.2 does not depend on the elite failing to see her own infatuation. If a forward-looking elite restrains herself, it is the stock of grievances that restrains her, not the foresight of her taste.

To get some feeling for how forward-lookingness would change the picture, I have computed what a forward-looking commander of the screen does at the calibration with infatuation and no learning; the method and the numbers are in the appendix. Two things come out. When approval is weighed as heavily as in the paper's calibration, she shows no own code at all and is never in danger: the stock of grievances, seen whole, outweighs the habit and the teaching together. So the positive share of the paper rests on the myopia in one place only, that the elite does not count the grievances she leaves behind. When approval is weighed lightly, she rides anyway, and knowingly; what her foresight buys her is the pacing, a few dozen seasons in office instead of seven. One caution. The forward-looking

elite keeps collecting the utility of her own code after approval has fallen below one half, since the model has no contest; a model of office-seeking would stop her payoff at defeat, and that is a different problem.

Simmel's reverse. Simmel (1904) explains fashion as a cycle driven by two motives at once: the lower strata imitate the upper, because the style signals rank, and the upper abandon a style as soon as it has been imitated, because it then signals nothing; so fashions trickle down and are replaced from above, over and over. In this paper the elite does the opposite. Once its code has leaked it no longer distinguishes, and an elite whose milieu valued distinction would then move on to a new closed code, and the whole path would start again with a new code a. That is Simmel's cycle with a screen attached, and the model has the ingredients for it: literacy as a state, and a role that chooses what to show.[2] What it would add is that each new code brings a new wave of exclusion, and a new stock of grievances.

National contact. The analysis sealed the regions, $\omega = 0$. With $\omega > 0$ a citizen's acquaintances are drawn from the whole country with probability ω, so the East borrows readers from the West, its learning starts earlier, and the valley is shallower. At the baseline of Figure 3 the majority is lost at $\omega = 0$ with a minimum of 0.498; at $\omega = 0.1$ the minimum is already 0.502 and the elite is safe, and it rises to 0.516 at $\omega = 1$, which is the limit in which the two regions are one. The valley is a product of the regions being sealed, and a little contact goes a long way towards filling it.

A different rule for a bad season. The rule of Section 3 makes a season bad when every Monday was. The convexity of the season's cost matters for how abruptly the share responds, since with $m = 2$ the interior window is narrow. The U-shape of Section 4.5 survives under a rule in which two bad Mondays out of three suffice, whose cost $3g^2 - 2g^3$ is not globally convex: minimum approval is 0.540 with no leakage, 0.461 at $\kappa = 0.003$ with the majority lost in season 2598, and 0.517 at 0.12 with the elite safe. So the U comes from the elite's response to learning, and the convexity of the cost only shapes how abrupt that response is.

Grievances that outlive illiteracy. A citizen who learns the code stops minding, in the model. If instead she kept her old grievances and only stopped acquiring new ones, the stock would be kept per original citizen outside milieu A rather than per illiterate

[2] Händler and Huck (2026) draw a line between fashion, which lives in the daringly unapproved, and luxury, which displays the approved to a public that can read it. In the terms of this paper, fashion is a sect with a closed code and luxury is lineage, a code that anyone can learn given time. The elite of this paper holds a screen, which is luxury's instrument, and shows a sect's code on it. Simmel's reverse would be an elite that behaved like fashion and moved on.

one, with the vote weighting it by $\mu_r(1-\alpha_r)$ alone. The elite's choice is unchanged, since past grievances are additive in her objective either way, so the share, habit and literacy paths are the same and only approval moves. At the calibration of Figure 3 the minima of approval become 0.518, 0.481 and 0.514 at $\kappa = 0$, 0.003 and 0.12, against 0.518, 0.481 and 0.520. The safe-unsafe-safe pattern survives this variant at these parameters.

The contest. The model has no election; it has approval, and the season in which approval falls below one half. Under a majority rule with no noise that season is the season of defeat. Under a probabilistic contest held every K seasons with success probability equal to approval, survival is a product over contest dates and the calendar K matters: with $K = 1$ even a safe elite is gone within a few seasons, since approval is far from one, so a contest every season is not the reading the paper intends. Specifying K would turn the dates in the figures into probabilities of defeat. The elite's choice would not change, since her objective is linear in approval.

Leaders who know all three codes. In the baseline a leader shows her own code or the common one. A leader who could show the other milieu's code would have a further move: slumming it, a leader of A showing code b. That might not be such a bad strategy, but I have not explored it. The move in the other direction, a leader of B who could show code a and refuses to, is the Bauhaus-fighting prime minister; in the baseline that refusal is forced on her by the action set, and here it would be her choice.

8 Coda

This paper began with a fable and continued with my reminiscences of the mid-noughties, when I finally had a satellite dish in London and was able to catch up with *Tatort*. The core idea for this paper formed then, so, in many ways, this is a project that has been twenty years in the making. But little did I know then where we would find ourselves two decades later with the AfD now in power in Sachsen-Anhalt and an assault on the Bauhaus in Dessau imminent.

This is depressing, and I am also a little sad that I did not take the time to write these thoughts up earlier.

Over the years I had my own grievances about politics, and there was a point in time when I asked myself why I was not joining a radical opposition. Two things held me back: a stance against racism, and the feeling of some kind of belonging. I compounded them in one sentence that could have given this paper another title: "Nazis don't buy Eames chairs."

A The forward-looking elite

The forward-looking commander of the screen of Section 7 maximises

$$\sum_{t \geq 0} 0.97^{t} \left[(\psi_0 + h_t)\, \pi_t + R\, V_t \right]$$

over the whole path of shares, with the stocks moving as in Section 3.2 and approval still a preference rather than a survival rule. The calibration is that of Section 4.4 with $\eta = 0.0004$ and no learning, $\kappa = s = 0$. The share is held constant over blocks of fifty seasons; the horizon is two thousand seasons, continued at the last share for a further three thousand so that nothing depends on the endpoint; and the optimisation is started from several paths, including the myopic one. In the no-learning case the problem has a closed form, since the stock of grievances can be summed forward and the objective is a discounted quadratic in the shares with a habit term; the block solution was checked against it. The table gives, for several values of R, the weight on approval, the forward-looking elite's initial share, her share after a hundred seasons, the lowest approval along her path, and the season in which she loses her majority, if she does, together with the same for the myopic elite of the paper.

			forward-looking		myopic	
R	π_0	π_{100}	lowest approval	majority lost	π_0	majority lost
1.00	0.00	0.00	0.550	never	0.41	season 39
0.30	0.04	0.05	0.546	never	0.81	season 7
0.20	0.10	0.14	0.546	never	1.00	season 7
0.15	0.19	0.34	0.518	never	1.00	season 7
0.10	0.65	1.00	0.431	season 36	1.00	season 7
0.08	1.00	1.00	0.431	season 7	1.00	season 7

At $R = 0.1$ the forward-looking elite's initial share depends on how finely the path is allowed to move: with shorter blocks it starts lower, around 0.4, and reaches the corner later, but it loses the majority in the same season to within a few seasons. The thresholds in R are illustrative and the discount factor is a choice; the pattern, nothing when approval is weighed heavily and a paced ride when it is weighed lightly, is not sensitive to either.

References

Adena, M. and S. Huck (2026). Non-familiarity with modernism and support for right-wing populists. Mimeo, WZB Berlin.

Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *Quarterly Journal of Economics* 115(3), 715–753.

Alesina, A. and N. Fuchs-Schündeln (2007). Good-bye Lenin (or not?): The effect of communism on people's preferences. *American Economic Review* 97(4), 1507–1528.

Bass, F. M. (1969). A new product growth model for consumer durables. *Management Science* 15(5), 215–227.

Becker, G. S. and K. M. Murphy (1988). A theory of rational addiction. *Journal of Political Economy* 96(4), 675–700.

Bisin, A. and T. Verdier (2000). "Beyond the melting pot": Cultural transmission, marriage, and the evolution of ethnic and religious traits. *Quarterly Journal of Economics* 115(3), 955–988.

Bisin, A. and T. Verdier (2001). The economics of cultural transmission and the dynamics of preferences. *Journal of Economic Theory* 97(2), 298–319.

Bonomi, G., N. Gennaioli and G. Tabellini (2021). Identity, beliefs, and political conflict. *Quarterly Journal of Economics* 136(4), 2371–2411.

Bourdieu, P. (1985). The social space and the genesis of groups. *Theory and Society* 14(6), 723–744.

Brierley, K. (2026). Wahlen in Sachsen-Anhalt: So radikal will die AfD die Kultur umbauen. *SRF Kultur*, 5 September 2026.

Bondar, M. and N. Fuchs-Schündeln (2023). Good bye Lenin revisited: East-West preferences three decades after German reunification. *German Economic Review* 24(1), 97–119.

Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste.* Translated by R. Nice. Harvard University Press.

Carmichael, C. M. (2025). Political representation in an era of income inequality and post-truth politics. *Constitutional Political Economy* 36, 171–199.

Corneo, G. and O. Jeanne (1997). Conspicuous consumption, snobbism and conformism. *Journal of Public Economics* 66(1), 55–71.

Downs, A. (1957). *An Economic Theory of Democracy.* Harper & Brothers.

Durante, R., P. Pinotti and A. Tesei (2019). The political legacy of entertainment TV. *American Economic Review* 109(7), 2497–2530.

Gennaioli, N. and G. Tabellini (2025). Presidential address: Identity politics. *Econometrica* 93(6), 1937–1967.

Gerbner, G. and L. Gross (1976). Living with television: The violence profile. *Journal of Communication* 26(2), 172–199.

Gethin, A., C. Martínez-Toledano and T. Piketty (2022). Brahmin left versus merchant right: Changing political cleavages in 21 Western democracies, 1948–2020. *Quarterly Journal of Economics* 137(1), 1–48.

Gidron, N. and P. A. Hall (2017). The politics of social status: Economic and cultural roots of the populist right. *British Journal of Sociology* 68(S1), S57–S84.

Guriev, S. and E. Papaioannou (2022). The political economy of populism. *Journal of Economic Literature* 60(3), 753–832.

Händler, E.-W. and S. Huck (2026). Fashion. In *Dialogues: Long-form conversations in English and German*. steffenhuck.github.io/fashion.

Karni, E. and D. Schmeidler (1990). Fixed preferences and changing tastes. *American Economic Review* 80(2), 262–267.

Landes, J. (2025). Modern denken. *Bauwelt* 1.2025.

Lazear, E. P. (1999). Culture and language. *Journal of Political Economy* 107(S6), S95–S126.

Lenski, G. E. (1954). Status crystallization: A non-vertical dimension of social status. *American Sociological Review* 19(4), 405–413.

Levy, G., R. Razin and A. Young (2022). Misspecified politics and the recurrence of populism. *American Economic Review* 112(3), 928–962.

Lévy-Garboua, L. and C. Montmarquette (1996). A microeconometric study of theatre demand. *Journal of Cultural Economics* 20(1), 25–50.

Loewenstein, G., T. O'Donoghue and M. Rabin (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics* 118(4), 1209–1248.

Mueller, J. E. (1970). Presidential popularity from Truman to Johnson. *American Political Science Review* 64(1), 18–34.

Noordzij, K., W. de Koster and J. van der Waal (2021). A revolt of the deplored? The role of perceived cultural distance in the educational gradient in anti-establishment politics. *British Journal of Sociology* 72(5), 1448–1463.

Norris, P. and R. Inglehart (2019). *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press.

Osborne, M. J. and A. Rubinstein (1998). Games with procedurally rational players. *American Economic Review* 88(4), 834–847.

Osborne, M. J. and A. Rubinstein (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior* 45(2), 434–441.

Passarelli, F. and G. Tabellini (2017). Emotions and political unrest. *Journal of Political Economy* 125(3), 903–946.

Pesendorfer, W. (1995). Design innovation and fashion cycles. *American Economic Review* 85(4), 771–792.

Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy* 78(4), 745–763.

Salant, Y. and J. Cherry (2020). Statistical inference in games. *Econometrica* 88(4), 1725–1752.

Shayo, M. (2009). A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political Science Review* 103(2), 147–174.

Simmel, G. (1904). Fashion. *International Quarterly* 10, 130–155.

Smaldino, P. E., T. J. Flamson and R. McElreath (2018). The evolution of covert signaling. *Scientific Reports* 8, 4905.

Spiegler, R. (2006). The market for quacks. *Review of Economic Studies* 73(4), 1113–1131.

Tuchman, G. (1978). Introduction: The symbolic annihilation of women by the mass media. In G. Tuchman, A. K. Daniels and J. Benét (eds.), *Hearth and Home: Images of Women in the Mass Media*. Oxford University Press, 3–38.

Vester, M., P. von Oertzen, H. Geiling, T. Hermann and D. Müller (2001). *Soziale Milieus im gesellschaftlichen Strukturwandel: Zwischen Integration und Ausgrenzung*. Suhrkamp.

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review* 99(5), 1899–1924.