Why Crackdowns Come in Waves—A Model of Authoritarian Repression

Steffen Huck*
WZB & UCL

October 17, 2026

Abstract

Authoritarian regimes often alternate crackdowns with quiet intervals. This paper offers an explanation for why repression occurs in waves. A regime that wants to survive can mount costly crackdowns accelerating the demolition of its opposition while structural consolidation erodes the opposition also during lulls. The paper shows how the timing of crackdowns depends on an opposition's innate strength and on the regime's surveillance capacity, and characterises three archetypal wave configurations.

1 Introduction

"Car, car over all, it has taken a terrible hold of us," Victor Klemperer noted in his diary in May 1936.[1] Klemperer was a Jewish philologist in Nazi Germany, dismissed from his Dresden chair the year before. His diary records from 1933 onward his fear and anguish over the Nazis' rise to power. But in May 1936 his attention is held by a motorist's ordinary worries—petrol prices, a failing starter. It was a comparatively quiet year of the persecution, and in later months his focus shifts again to the decrees that never paused, down to the one that would take his driving licence in 1938. His diary poses, from the inside, the question of this paper: why are there lulls in an authoritarian state—and how do they alternate with crackdowns?

Authoritarian regimes rarely repress at a constant intensity (Davenport and Appel, 2022;

*I worked on all parts of this paper in close collaboration with Claude (Fable)—to explore my initial idea, to fine-tune the model, to develop and verify the proofs, and to shape the draft. The numerical analysis is entirely Claude's, with me as mere supervisor. I also used ChatGPT (Astra) as an adversarial referee in successive rounds and once consulted Refine.ink. The usual disclaimer applies.

[1] Klemperer (1998, p. 201), in the Chalmers translation.

Truex, 2019). Periods of visible crackdown alternate with calmer intervals in which a weakened but surviving opposition and civil society operate openly. Sometimes the sequence closes with a decisive crackdown that destroys what remains; sometimes crackdowns simply stop, and the opposition fades without any final event. The puzzle is not why an authoritarian regime engages in crackdowns. It is why the regime sometimes prefers quiet spells—and what determines whether the story ends with a final blow or with quiet erosion.

The regime only ever does one of two things: it cracks down, or it does not. A *crackdown* is overt coercive mobilisation—arrests, prosecutions, raids, bans; a *lull* is the quiet between crackdowns. Beneath both runs a programme the regime set in motion from the outset: *structural consolidation*—association and foreign-agent laws, court capture, media consolidation, university purges, the closure of funding channels. Consolidation never pauses. That is an assumption, it does important work. I will discuss its role in detail and will show in which situations lulls are really thaws with regrowth of oppositional strength and when they can be final with quiet erosion doing the regime's work.

I build the simplest model in which these distinctions have teeth. A regime that extracts value from being in power faces an opposition that may topple it. The opposition regrows, when left alone, from a base of people, institutions, and organisational knowledge. It only collapses when it is eroded to a minimal level at which stage the regime is no longer threatened. The regime chooses when and for how long its coercive apparatus is mobilised for crackdowns. Mobilisation has a flow cost while it lasts and a fixed cost each time it begins. Surveillance has decreasing returns in scarcity—large visible oppositions are easy to disrupt, small clandestine remnants are not. As long as it holds power, the regime also drives oppositional capacity down through structural consolidation proceeding at a constant pace. The costs of consolidation are subsumed in the flow payoff the regime receives while in power.

Three main results are worth summarising. *First, repression comes in waves.* The fixed cost makes the regime buy repression in batches, much as a firm facing an ordering cost lets inventory run down before reordering. It cracks down when opposition reaches an upper trigger, stops at a lower one and tolerates regrowth, then repeats. Comparative statics on these fixed costs will illustrate how wave lengths are shaped by them.

Second, most crackdowns stall before finishing the opposition. This essentially stems from surveillance technology. As oppositions shrink, effective surveillance gets ever harder. As a consequence, all crackdowns (with the possible exception of one last hard one) will stop when oppositional strength reaches a lower threshold at which it is still viable. Surveillance technology, thus, becomes a key parameter, with mostly obvious comparative statics but a more subtle role for how campaigns with crackdowns and lulls unfold.

Third, the waves come in three archetypes. (a) the *sudden ending*. The regime alternates between relatively short phases with crackdowns and often longer phases of lulls. During these waves, consolidation does its more subtle work until the regime suddenly decides to engage one final longer crackdown finishing the opposition. (b) the *surprise ending*. Here the picture is initially almost identical, only that the regime decides at some stage for a much longer lull (that observers might misinterpret as a thaw). During that last long lull opposition initially regrows but then begins to erode thanks to accumulated erosion. When it becomes sufficiently weak, the regime strikes like in (a) with one prolonged and final crackdown. (c) the *silent ending*. This archetype looks almost identical to (b) only that the final crackdown is never necessary as erosion on its own does the job. These different configurations will be shown below in Figure 1, which is very much at the intellectual heart of this paper. One aspect is worth noting for all these archetypes: they all begin with an initial lull. It's rare that a new authoritarian regime begins with an immediate crackdown. This is again related to surveillance becoming more effective with a stronger opposition. Perhaps one of the ironies of this paper.

I hold society—the opposition and civil society—passive. Its capacity to resist and, perhaps, topple the regime follow simple stock dynamics which helps sharpen the findings on waves. But do notice that opposition is by no means naive or toothless. It carefully watches the regime, worries and strikes and I carefully analyse when it succeeds and the regime fails. In many parameter regions, the regime will fail more often than not. The paper remains deliberately abstract throughout: no history is claimed as an instance of the model, and Klemperer's diary is cited for what a lull looks like from inside one, not as a case the model explains.

The model is largely analysed numerically but also offers some low-hanging analytical fruit. Appendix A shows the proofs, Appendix B the details on the computations.

The paper proceeds as follows. Section 2 places the model in the literature and Section 3 sets it out. Section 4 contains the analysis; Section 5 discusses what the results do and do not claim, and Section 6 concludes.

2 Related literature

This section maps the neighbourhoods the paper borders. It has three.

Fast and slow. One literature studies repression and dissent in motion: the puzzle that repression sometimes deters and sometimes provokes (Lichbach, 1987; Davenport, 2007), and models of tipping and cascades (Granovetter, 1978; Kuran, 1991). Another studies

backsliding and autocratic legalism: the slow, statutory hollowing-out of the institutions that sustain organised opposition (Bermeo, 2016; Levitsky and Ziblatt, 2018), and the finding that organisational capacity, not grievance, is what makes opposition effective (Chenoweth and Stephan, 2011). My two state variables are these two literatures' objects—the fast literature tracks x, the slow literature documents K—and what this paper adds is the optimal timing of the fast instrument, conditional on the slow process.

The strategic state. An early antecedent is Spagat (2001): a survival-seeking dictator chooses repression period by period, and the state—fear and hate, both fed by past repression—drives alternation between maximal and minimal repression. There the cycling comes from society's state responding to repression; here society is passive, the cycling comes from a fixed mobilisation cost against a regenerating stock, and the separate reproductive base supplies the terminal lull and the two endings. The closest mechanical relative is Di Lonardo et al. (2026), a dynamic model in which repression addresses acute threats while consolidation reduces future mobilisation. There, repression interacts with consolidation directly—it can hasten or hinder it; here, consolidation runs independently of the crackdown decision, and its interaction with a startup cost and scarcity-sensitive detection is what produces the terminal lull—an object that does not appear in their analysis. Among informational theories, the closest relative is Lorentzen (2013), where an authoritarian government tolerates a regulated flow of protest because protest is informative; there, quiet is an information policy calibrated to a signal, while here it is cost-timing pinned to hitting times of a threshold, and the two are empirically separable. On the choice of annihilation over repression, Esteban et al. (2015) model mass killings as an investment in future power, and Gregory et al. (2011) document cost–benefit repression in Stalin's archives. This paper adds the timing margin—and, on part of the parameter space, its absence. The survival objective is the canonical one (Wintrobe, 1998; Svolik, 2012; Guriev and Treisman, 2019, 2022); throughout, regime survival means survival of the institutional order, not of an individual ruler.

Temporal and preventive repression. A quantitative literature studies when repression starts and stops: Truex (2019) documents calendar-driven, preemptive cycles; Davenport and Appel (2022) study the onset, termination, and recurrence of repression spells; Ritter and Conrad (2016) show that preventive repression creates a deep observational problem, since successful prevention leaves no dissent to observe; and Sullivan (2016a) shows, in the archives of the Guatemalan National Police, that repression aimed at clandestine organisational capacity, rather than at overt dissent, is what reduces subsequent mobilisation—the empirical cousin of K; his companion study of overt enforcement is Sullivan (2016b). These papers supply mechanisms—calendars, spells, anticipation—that are alternatives to the cost-timing mechanism here, and the hysteresis prediction below is where they separate.

Technique. The wave band is two-mode optimal switching with a fixed cost, the control cousin of the (s, S) inventory problem (Scarf, 1960; Stokey, 2009); logistic recruitment with removal, and refuges below which harvesting cannot push, are familiar from renewable-resource economics (Clark, 1990). What is new is the slow state: the refuge is moved by a process the controller rides rather than steers, which converts a feasibility constraint into a theory of timing.

3 The model

The regime and its objective. The regime holds power, worth a flow payoff $v > 0$—read throughout as net extraction, a reading made precise below; ousting ends the game with continuation normalised to zero: the regime keeps all payoffs received up to that point, and what it loses is the future. Ousting arrives at hazard $\lambda_0 + \lambda_1(x)$, where $\lambda_0 \geq 0$ is the exogenous hazard and x is the strength of organised opposition:

$$\lambda_1(x) = \chi + D(x) \ \text{ while organised opposition survives,} \qquad \lambda_1 = 0 \ \text{ after its collapse,}$$

with D increasing and convex, $D(\varepsilon) \approx 0$ at the collapse threshold ε introduced below, and $\chi > 0$ the residual hazard of any surviving autonomous organisation—the standing possibility that dormant networks reconnect. The regime wants to survive and to spend little doing so.

Opposition. Left alone, opposition regrows from its base:

$$\dot{x} = a_K(x) = g\,x\Big(1 - \frac{x}{K}\Big), \qquad g > 0, \tag{1}$$

growing below capacity and declining above it; K is the rest point of net recruitment, not a hard ceiling, and $x > K$ is transiently possible when the base erodes beneath the current stock. Organised opposition requires minimal scale, and I impose a single collapse threshold $\varepsilon > 0$ covering both margins: when the stock falls to ε, the remnant is too fragmented to coordinate or reproduce as an organised actor; and when the *base* falls to ε, no organisation can be sustained on the remaining substrate—a movement whose funding, legal personality, premises, and carriers of know-how are gone dissolves even with members still nominally present. Either event is *organisational collapse*, and it is absorbing: collapse eliminates organisational capacity, not persons. The Appendix shows that the computed switching topologies—waves and ending alike—survive substantially different base-side collapse rules, with dates, margins, and values shifting as expected.

Consolidation. As long as it holds power, the regime also erodes the base itself: structural consolidation—statutes, court appointments, media acquisitions—drives K down at a constant pace,

$$\dot{K} = -\theta, \qquad \theta > 0, \tag{2}$$

in crackdown years and quiet years alike. I treat this programme as given: its pace is a regime characteristic, and its flow cost is borne as long as the regime holds power—before and after the opposition's demise alike, for captured courts must be staffed and statutes maintained—so it nets out of the flow payoff. v is extraction minus the upkeep of the quiet machinery, and because the cost runs under every policy and every outcome, no margin in the model is distorted by it. The model studies only the timing of crackdowns on top of the programme. Because (2) runs identically under every crackdown policy, capacity will act as a clock. I hold K fixed through Sections 4.1–4.2 to isolate the fast dynamics, and let the clock run in Section 4.3.

Crackdowns. The coercive apparatus is either *off* or *mobilised*. Mobilised, it removes organisation at rate

$$m(x) = \bar{m}\,\frac{(s\,x)^2}{1+(s\,x)^2}, \qquad \bar{m} > 0,\ s > 0, \tag{3}$$

so $\dot{x} = a_K(x) - m(x)$ during crackdowns. The parameter s is *surveillance strength*: raising s raises removal at every opposition size. Removal saturates at $\bar{m}$—the apparatus can only process so many cases—and removal per unit of organisation, $m(x)/x$, vanishes as $x \to 0$: small, scattered remnants are hard to find, and how hard depends on s.[2] Mobilisation costs $c \in (0, v)$ per unit time; switching the apparatus on costs $\kappa \geq 0$; switching off is free. There is no intensity dial, deliberately: if a mobilised apparatus costs the same whether or not used to capacity, full application is optimal within any crackdown, and partial repression is delivered by cycling the apparatus on and off—which the fixed cost prices.

The problem. An admissible policy is a right-continuous mode process with finitely many switches on every compact interval; the final crackdown episode may have no finite end. Writing $\rho > 0$ for the effective rate combining time preference and the exogenous hazard, and $\Lambda(t) = \exp\big(-\rho t - \int_0^t \lambda_1(x(u))\,\mathrm{d}u\big)$ for survival discounting, the regime maximises

$$\int_0^\infty \Lambda(t)\Big[v - c\,\mathbf{1}\{\text{mobilised at } t\}\Big]\,\mathrm{d}t \;-\; \sum_i \Lambda(t_i)\,\kappa,$$

over admissible policies, where the t_i are the mobilisation dates and the continuation after collapse is the safe haven v/ρ. Two facts frame everything. The effective, survival-adjusted

[2] Surveillance is fixed technology here; endogenous investment in s is left to future work.

discount rate along any path is $\rho + \lambda_1(x)$: a threatened regime discounts as if impatient—the hazard mimics impatience without being a preference—while a nearly-safe one values the future almost at its pure rate ρ. And any policy under which the opposition survives forever is worth at most $v/(\rho + \chi) < v/\rho$: the gap is the guaranteed minimum value of ending coexistence for good, and it exists because the residual hazard vanishes at collapse. Time is measured in model periods, with no commitment to a calendar unit; I return to timescales in Section 4.3.

4 Analysis

4.1 Waves

Hold the base fixed. The first question is why repression comes in bursts rather than at a steady simmer.

Proposition 1 (No cycles under memoryless state rules)**.** *At any switching cost, including zero, no stationary rule that maps the opposition's strength alone into a mode generates cycles: under any such memoryless rule (with the regularity and solution convention of the Appendix), trajectories are monotone up to stagnation and never return to a state already visited.*

The qualifier matters. A rule that conditions on the current mode as well as the state—mobilise at an upper trigger, demobilise at a lower one—can cycle forever even at zero cost, because the mode itself supplies the memory. What Proposition 1 establishes is that cycling requires that memory. And at zero cost the relevant benchmark is not steadiness but *chattering*: holding the opposition at an interior point requires switching, and ideal containment is the limit of ever-faster, ever-smaller cycles—a sliding mode, cycling at infinite frequency with vanishing amplitude. The Appendix sketches, without a complete proof, why nothing beats this relaxed benchmark; weak dominance there would still leave costless cycles co-optimal, so the fixed cost is not what makes cycling possible. What the fixed cost does is set its scale. At $\kappa = 0$ the relaxed containment benchmark cycles invisibly fast; at $\kappa > 0$ the cycles acquire positive amplitude and finite frequency—waves wide enough to see, on the timescale of years rather than instants. Here the cost is concrete: the political, organisational, and reputational price of putting courts, security services, and propaganda on a campaign footing.

Proposition 2 (The switching cost opens the band)**.** *Conditional on the existence of an interior recurrent threshold optimum (within the class defined in the Appendix), every $\kappa > 0$ implies a strictly positive band: mobilise when opposition reaches an upper trigger $x_H(K)$, stand down at a lower threshold $x_L(K) > \varepsilon$. At $\kappa = 0$, collapsed bands attain a relaxed*

smooth-containment benchmark, computed in the Appendix; in the computations, optimal bands shrink as κ falls.

The economics is inventory management: the opposition is a renewable nuisance, and when every order carries a fixed charge, orders come in batches. A lull is not restraint; it is the regime amortising the startup cost of the next crackdown over the longest tolerable stretch of danger. Larger fixed costs mean longer lulls and deeper crackdowns—in all my computations, ceteris paribus, and provided crackdowns remain worth mounting at all—so polities where mounting a crackdown is politically expensive should exhibit longer cycles and harder individual crackdowns, not milder ones.

The fixed cost is paid at every launch, and one might count that as an argument against cycling: nine crackdowns, nine times κ. But the alternative to cycling is not one crackdown—while elimination is infeasible, there is no once. The alternatives are permanent mobilisation, which pays c forever to pin a remnant where removal works worst, and no mobilisation, which bears the hazard of an opposition at capacity. Cycling is the middle course: each κ buys a batch of removal taken where removal is efficient, and each lull is a stretch in which the regime pays neither κ nor c and the regrown danger is still tolerable. Section 4.4 computes the whole road as κ grows, and both of its ends.

These arguments can sound subtle, and they may get buried in the mathematics. So I asked Claude to explain the logic of lulls to an eight-year-old child. Here it goes:

> Why doesn't the gardener mow every single day? Because dragging the mower out of the shed is a lot of work, and the grass only grows a little each day—so he waits until it is long, cuts it all at once, and puts the mower away again. The quiet days are not the gardener being kind to the grass; they are just the cheapest way to keep it short.

4.2 Why elimination must wait

Near zero, regrowth (1) is linear in x while removal (3) is quadratic: a small remnant recruiting from a strong base outruns detection. A crackdown therefore pushes opposition down only to the point where removal and regrowth balance—call it the *floor*, $x_s(K)$—and no further, however long the crackdown persists. The floor rises with the base: against a strong base it lies well above the viability threshold, and the opposition survives beneath the apparatus's reach. Consolidation lowers regrowth at every size and so lowers the floor; only once capacity has eroded far enough does the floor slip below ε, and a crackdown pressed hard enough can, for the first time, carry the opposition below viability. The floor's descent over time is

what drives the life cycle of Section 4.3. (Formally I assume the balance point is unique and regular in the relevant range, verified at all parameters used; Appendix A.)

Proposition 3 (The feasibility frontier)**.** *The floor rises with the base. Suppose* $\bar{m}\, s^2\varepsilon < g(1+s^2\varepsilon^2)$*; then, within the verified descent basin, elimination at a fixed base is possible if and only if capacity has fallen below the closed-form frontier*

$$K^\star = \frac{\varepsilon}{1 - \dfrac{\bar{m}\, s^2 \varepsilon}{g\,(1+s^2\varepsilon^2)}}, \tag{4}$$

where the floor meets the viability threshold. When the feasibility inequality fails, set $K^\star = \infty$*: elimination is feasible against every base in the verified range.* $K^\star$ *rises with surveillance* ($\partial K^\star/\partial \bar{m} > 0$, $\partial K^\star/\partial s > 0$) *and falls with the opposition's regrowth rate; provided* $\bar{m} > g\varepsilon$*, it diverges as* s *rises to the critical surveillance strength* $s_c = \sqrt{g/\big(\varepsilon(\bar{m} - g\varepsilon)\big)}$.

Two readings. For the regime: against a healthy civil society, the question is not whether elimination is worth its price but whether it is possible at all—and for a wide range of surveillance technologies it is not. For the opposition: *a strong base is a shield.* While $x_s(K) > \varepsilon$, no crackdown can destroy organised opposition; what can be destroyed is the base, quietly, by statute—which is exactly what consolidation does.

One distinction matters for what follows and is easy to miss. $K^\star$ says when elimination becomes technologically possible at the current base. When an optimising regime actually *launches* is a separate question of value matching—the startup and flow costs must be worth paying—and its answer need not coincide with the frontier: at the benchmark economics the computed launches lie *above* it, in both final-crackdown types, with the crackdown completing as consolidation carries capacity down through the frontier; at leaner economics or slower erosion, launches descend to the frontier's neighbourhood. The Appendix develops the distinction precisely; the life cycle below shows the mechanism on the optimal path.

4.3 The life cycle

Now let consolidation run. As long as the regime holds power it drives the base down at a constant pace, $\dot{K} = -\theta$, the same in crackdown and calm—association laws do not pause when arrests do. Because capacity falls identically under every crackdown policy, $K(t)$ acts as a clock, and the regime's full dynamic problem can be computed as posed; every path and number below comes from one computed policy—the optimised hybrid of Appendix B, constructed, simulated, and tested against deviations as documented there. I hold one calibration throughout ($\theta = 0.006$ per period; parameters in the Appendix); Figure 1 varies

only the surveillance technology, and the diagrams that follow it vary the stakes of office as well.

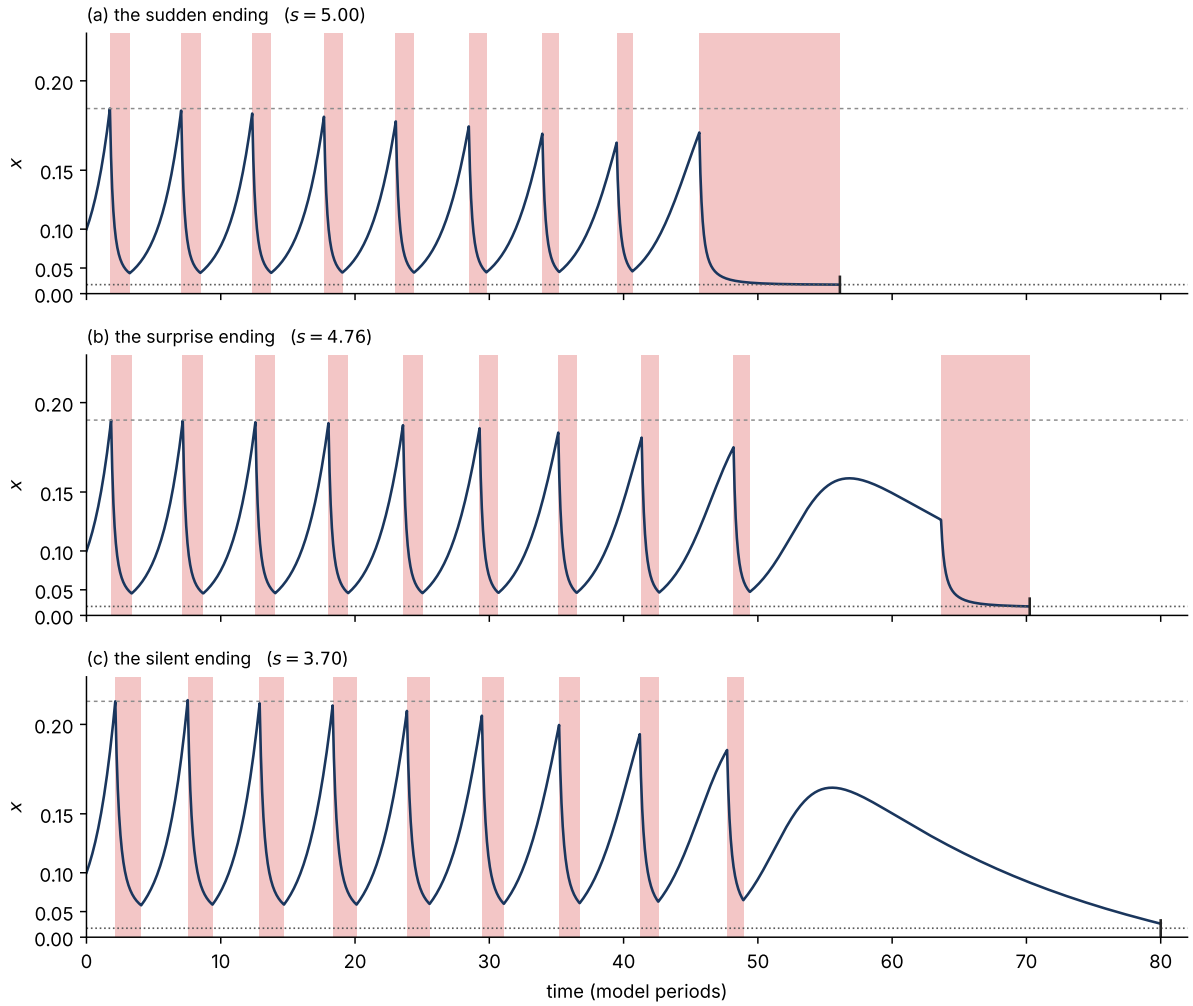


Figure 1: Three regime histories, identical except for surveillance strength s; the base erodes at the same pace in all three. Solid: opposition $x(t)$; shaded: crackdown on; dashed: the first peak, for reference; dotted: the viability threshold; tick: collapse. The vertical axis is mildly exponential, stretching the top of the panel so that the slow decline of the peaks is visible. (a) The sudden ending. (b) The surprise ending. (c) The silent ending.

Figure 1 shows three archetypes of repression. All three run the same opening—a rhythm of crackdowns and lulls, containment bought in batches—over the same silent clock, consolidation eroding the base at an identical pace. They differ in one thing: who delivers the ending, the apparatus or the clock. In panel (a) the apparatus delivers it on schedule: one ordinary crackdown turns out to be the last, because it never needs to stop. In panel (b) the apparatus delivers it late, after letting the clock cheapen the job, and the waiting is visible as a quiet several times longer than any before it. In panel (c) the clock delivers it alone: a final

crackdown is no longer necessary, the lull becomes indefinite, and the opposition declines with the base that can no longer carry it. I call these the *sudden ending*, the *surprise ending*, and the *silent ending*—names for how the ending looks in the crackdown record, not claims about anyone's beliefs. The first two cross over. In the sudden ending the crackdown is expected and the ending is not: the final crackdown begins at an ordinary trigger, looks like any other wave, and is revealed as final only by never standing down. In the surprise ending the ending has been in preparation for a decade and the crackdown is the surprise, arriving out of a calm that looked like peace. In the silent ending there is neither: no event at all, an ending visible only in retrospect. Formally, write 0 for off and 1 for on, and let $n \geq 0$ count completed ordinary crackdowns. A history before collapse is $0(10)^n1$ when a final crackdown occurs and $0(10)^n$ when it does not; the opening off-phase may have zero length. The count n is an attribute of a history, not a type. Its corners are the archetypes at birth: near initial feasibility the rhythm never forms and the first crackdown is the last, while at low stakes and high impatience crackdowns never start and the whole history is one lull. What separates the archetypes is one economic comparison and one labelling convention. The comparison is whether a final crackdown ever beats waiting for attrition; it divides the silent ending from the other two. The convention tells the sudden from the surprise ending, which share the sequence $0(10)^n1$: a final lull exceeding twice the ordinary average marks the surprise. The convention requires $n \geq 1$; a history with $n = 0$ has only its opening lull and is classified by the comparison alone. Moving the convention moves the sudden/surprise boundary without changing any policy. Both are made precise in Appendix B, and Section 4.4 shows the count stepping down through intervals, nowhere on a knife edge.

Each history begins with *waves*. In panel (b) the regime runs nine full cycles between its computed triggers—crackdowns of one and a half periods, lulls averaging four and a third—each launch a hitting time, the whole phase the batch-purchasing of Section 4.1 in motion. And each history opens the way it will continue, with a lull: born below its trigger, the regime waits exactly as it will wait after every stand-down. Letting the opposition grow before striking is not neglect; it is the batch logic of Section 4.1 from the first day, reinforced by the technology: a bigger opposition is easier to find, so waiting improves the apparatus's own efficiency—the irony noted in the introduction. The opening lull is short—about two periods against the ordinary four and a third—only because birth, at $x = 0.10$, sits closer to the trigger than the troughs do.

The line itself is worth reading slowly. Within a wave, the shape is asymmetric. The climb is gradual: regrowth is fed by the base, and it takes over four periods. The fall is fast: removal runs at full force, and it takes one and a half. The sawtooth is the cost structure made visible—patience on the way up, violence on the way down. Across the waves, everything

that drifts is the clock's doing. Without erosion the problem would be stationary and the figure a metronome: identical waves forever, constant triggers, no ending of any kind. With it, two things drift. The peaks decline: the regime moves at slightly lower stocks, because the falling base keeps shortening the threat's remaining future. The rebounds slow: regrowth is fed by the shrinking base itself. The troughs stay put: the regime always cuts to roughly the same depth. Peaks fall and floors hold, so the sawtooth narrows from above—slowly, and the slowness matters. To an observer the wave record looks stationary while the clock runs underneath it.

Then the waves stop, and they stop early. At a base still roughly twice the frontier, the last ordinary crackdown ends and no new one begins: at any recovered stock the instantaneous hazard is what it always was, but its *continuation* has shrunk with the base—the danger now has a shorter remaining life and a nearer collapse—and crackdowns no longer repay their fixed cost. What follows in panel (b) is the paper's central object. The *last lull* runs just over fourteen periods—more than three times the average ordinary lull. The opposition regrows, unmolested, toward stocks at which the regime used to crack down, peaks, and then begins to sink as the base erodes beneath it; the old triggers are never regained, because the shrinking base now caps the recovery below them—the lull peaks at 0.159, under every earlier launch. Appendix B's counterfactual separates the mechanisms: a regime mechanically running its old containment rule would have stayed quiet through the entire lull, because the shrinking base never delivers another trigger hit. The quiet itself is therefore no evidence of a changed policy; the genuine decision is the timing of the final launch into it. Arrests are at zero throughout; every crackdown series reads as liberalisation; a resistance that measures danger by prosecutions concludes the storm has passed. Meanwhile the floor descends toward the viability threshold. Nothing in the crackdown record marks the difference between this lull and a genuine thaw—the difference lives entirely in statute books, court rosters, and funding channels.

The surprise ending (panel b). With mid-range surveillance, the terminal crackdown eventually pays: launched at capacity 0.118, *above* the frontier $K^\star = 0.105$, against an opposition of stock 0.129—the lowest stock at which it ever launched a crackdown—and lasting six and a half periods, nearly five times the average ordinary crackdown: the crackdown holds the opposition pressed to the floor while consolidation carries the floor below viability. The final blow is launched from the lowest stock of its sequence and runs the slowest: a grind, not a blitz, ending in collapse.

The silent ending (panel c). With weaker surveillance, the frontier sits so low that the terminal crackdown never repays its costs—and it is never launched. After the last stand-down the

regime simply waits. The opposition tracks its collapsing base downward and is extinguished when the base itself reaches viability: collapse by *attrition*, at a finite, deterministic date, without a single further arrest. There is no final event to observe at all. A literature searching for terminal crackdowns will code this history as the dog that did not bark.

The sudden ending (panel a). Panel (a) shows what strong surveillance does to the boundary between the phases: the excess delay disappears. The final crackdown is launched from an ordinary wave trigger—*above* the frontier—and simply never stands down, completing as consolidation carries capacity down through the frontier; it is the longest crackdown of its sequence, an ordinary crackdown that never ends. Its launch even sits a touch above the last ordinary trigger—the final launch comes at 0.174 against its predecessor's 0.168—and the kink is informative. Ordinary triggers decline because each crackdown protects a future that consolidation keeps shrinking. The final crackdown protects no further cycle—what it buys is the safe haven itself—so its launch answers a different question: how much regrowth to tolerate against a deadline the clock has already set. The staircase of triggers falls; the last step is placed by the endgame. That two histories so alike produce endings so different has a precise source. The two technologies are nearly the same instrument where the waves live: removal at ordinary stocks differs by a few percent between the panels, and the wave phases come out nearly identical. But removal near viability is quadratic in $s\,x$, and the frontier amplifies the difference hyperbolically—five percent more surveillance moves $K^\star$ from 0.105 to 0.184, three-quarters higher. The wave phase uses the technologies where they agree; the ending uses them where they differ. At $s = 5.00$ the ordinary trigger already sits a few periods of capacity above the frontier, and the prize is large: staying mobilised until consolidation delivers feasibility is worth its cost, and finishing beats another cycle. At $s = 4.76$ the same hold does not pay—pressing on from the last ordinary trigger costs 2.2 in value against standing down—so the regime lets the clock work unpaid instead, and launches only two periods above the frontier. Push surveillance a little further and the rhythm itself goes: close to initial feasibility the wave count falls to one and then to none, with the computed optimal delay before the single crackdown reaching exactly zero a sliver inside the feasibility boundary—the regime launches at birth and grinds through. Across the three archetypes, surveillance deletes the quiet from the life cycle: first the surprise ending's long lull, then the waves themselves, one at a time. Holding the consolidation pace and costs fixed, improved surveillance moves the regime through this ordered sequence at every pace computed; slower consolidation widens the final-crackdown region—when free lethality recedes, almost every regime eventually strikes—while faster consolidation shrinks it, since waiting delivers the same death without the bill.

4.4 The rhythm

Before the maps, one piece of formal comparative statics—the piece the problem yields without touching the free boundary.

Proposition 4 (Comparative statics in the large)**.** *For any admissible policy with finite discounted switching cost, the regime's objective is affine in (v, c, κ), with nonnegative coefficients on v, $-c$, and $-\kappa$. The value $V(v, c, \kappa)$ at any state is therefore convex, nondecreasing in v and nonincreasing in c and κ, and wherever it is differentiable,*

$$\frac{\partial V}{\partial v} = \int_0^\infty \Lambda(t)\ \mathrm{d}t, \qquad -\frac{\partial V}{\partial c} = \int_0^\infty \Lambda(t)\,\mathbf{1}\{\textit{mobilised at } t\}\ \mathrm{d}t, \qquad -\frac{\partial V}{\partial \kappa} = \sum_i \Lambda(t_i),$$

evaluated along the optimal policy: the survival-discounted time in power, the survival-discounted time on campaign footing, and the survival-discounted number of crackdowns.

Corollary 1. *The attained share of the ceiling, $V(v, c, \kappa)/(v/\rho)$, is nondecreasing in v.*

The proposition is deliberately thin—a supremum of affine functions, Appendix A—but its slopes are precisely the objects the figures map: $-\partial V/\partial\kappa$ is the discounted cousin of the count below, $-\partial V/\partial c$ aggregates the shaded bars of Figure 1, and $\partial V/\partial v$ is discounted time in power, whose hazard-adjusted terminal piece is the survival factor of Section 4.5. The value whose slopes these are is itself the map of Section 4.6. Convexity means the marginal value of cheaper repression rises as repression gets cheaper. The corollary is one line: a policy's share of the ceiling is $\rho\,(A_\pi - B_\pi/v)$, with A_π its discounted time in power and $B_\pi \geq 0$ its discounted repression bill, and a supremum of functions nondecreasing in v is nondecreasing (Appendix A). It is the theory behind the third map. What the same logic does not deliver is the second map: discounted time in power rising with v does not order an undiscounted survival probability at a policy-dependent collapse date, so the survival monotonicity reported below is a computed finding. What theory does not deliver cheaply is monotonicity of the policy itself: that bands widen in κ or that triggers fall with the base is supermodularity analysis in a nonstationary free-boundary problem, and the computations carry those claims instead. Now the maps.

How many cycles does a history contain? The count is a wave-phase object, and at the benchmark the model separates it from the ending: forbidding terminal launches below $K = 0.15$ moves the wave triggers by at most 0.004 (Appendix B), so there the rhythm does not know about the ending until the ending reaches back into it—a benchmark diagnosis, not a theorem across the map. The count itself decomposes. It is the stretch of capacity the wave phase covers, divided by the clock speed times the cycle length, and everything in that

expression is containment economics: the band from κ, c, and the hazard at ordinary stocks; the clock from θ. Figure 2 maps the count over surveillance and the flow value of power—the stakes of office: raising v at fixed crackdown costs raises what incumbency is worth relative to what repression costs—on the same plane and axis as the survival map of the next subsection. Three features carry it. The count rises with the stakes: at mid surveillance, from four cycles at the lowest stakes to eleven at the highest—higher stakes tighten the containment band. The stakes also pick the ending. At benchmark surveillance the final crackdown is not worth its costs at $v \leq 1.2$, and those histories end in attrition, by margins of 0.01 to 0.11 (margins, here and throughout, are value differences at the last stand-down); at $v = 1.5$ the strike pays (+0.24) after a fourteen-period lull; by $v = 1.9$ the lull has compressed to nine periods (+0.87), and by $v = 2.5$ to five, under the labelling cutoff (+2.1). Raising the stakes thus walks the ending through the archetypes in order: silent, surprise, sudden. The technology is unchanged; what rises is the value of ending coexistence for good. The count is single-peaked in surveillance, with a modest amplitude of one or two cycles: weak detection stretches each crackdown, so fewer cycles fit into the same span of erosion, while past the peak the ending begins to consume the rhythm. At the right edge the consumption accelerates: each deleted wave requires the terminal launch to climb one cycle's worth of capacity, and near feasibility that capacity climbs steeply in s. So the staircase narrows into the white band where the rhythm never forms, descending one count at a time—every count a band of positive, shrinking width (Appendix B). The colour bands are the level sets of an integer: each count occupies a region of positive measure, and the boundaries between them are the threshold curves. Slower consolidation lengthens the rhythm without changing it: two dozen cycles and more at $\theta = 0.002$. A higher combined discount rate thins everything. Repression is investment—κ and c now against reduced hazard later—so an impatient regime cracks down less and loses the final crackdown first; at benchmark stakes and technology it is gone by $\rho = 0.07$, and at low stakes besides, crackdowns never start.

The fixed cost itself is swept the same way. Raising κ from 0.1 to 1.6 at the benchmark thins the rhythm from seventeen cycles to four and stretches ordinary lulls from under three periods toward seven: bigger batches, bought less often. The lulls never disappear—they win. What disappears is first the final crackdown and then the waves. At $\kappa = 1.6$ the regime still cycles but no longer strikes, its ending gone silent; at $\kappa = 3.2$ a single wave remains; at $\kappa = 6.4$ the apparatus never leaves the shed. That is one exit from cycling, and there is another. Raise the stakes to $v = 2.5$, or halve the mobilisation cost, and the same $\kappa = 1.6$ produces an everlasting crackdown instead: launched from an ordinary trigger and never stood down, the regime paying κ once precisely to avoid paying it again. Expensive crackdowns end the rhythm either way. Whether a polity exits into quiet or into permanent emergency is selected by the flow cost of repression and the value of power.

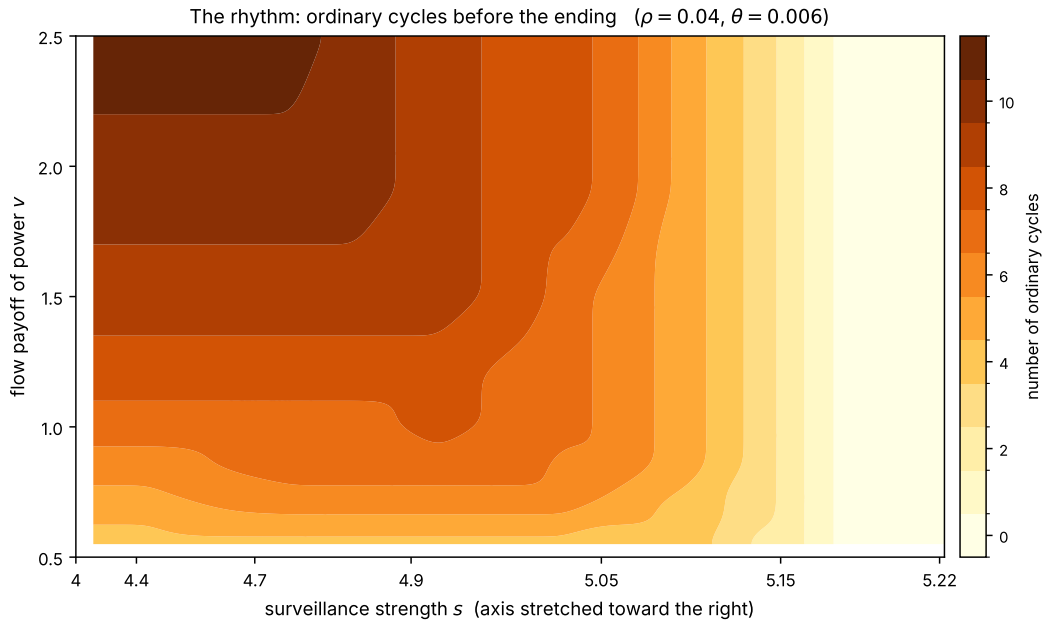


Figure 2: The rhythm: ordinary cycles before the ending, in surveillance and the stakes of office, on the axis of Figure 3. Bands are the level sets of the cycle count, computed on a fine grid, interpolated between grid points, and lightly smoothed; the thin bands' positions are verified by direct probes (Appendix B).

4.5 Survival

The survival landscape. Figure 3 colours a subset of the same parameter plane by the opposition-channel survival factor at collapse: the probability that the regime outlives its opposition, absent exogenous ousting. The landscape climbs elevenfold across the diagram, from 0.08 in the low-stakes, weakly surveilled corner to 0.88 at the right edge, where the single crackdown launches at once, and its gradient decomposes into three parts. Survival rises with the stakes of office everywhere, even deep inside the no-final-crackdown region where every history ends at the same attrition date: higher-stakes regimes run more and deeper crackdowns, hold opposition lower through the same eighty periods, and accumulate less hazard—containment pays in survival even when no ending ever comes. That is a computed fact, cell by cell, not an implication of optimisation alone: revealed preference disciplines discounted time in power (Proposition 4) and delivers Corollary 1, but neither orders survival at a policy-dependent collapse date. The boundary at which final crackdowns begin contributes a visible but modest step, the strike advancing a collapse whose remaining tail hazard is already low. And the steep part of the landscape is the approach to feasibility: between $s = 5.00$ and s_{init} survival more than doubles, as histories shorten from a full wave phase toward a single early crackdown. All hazard in the plotted opposition channel comes

from coexistence with a surviving opposition, so everything that shortens or thins it—a cheaper wave, an earlier ending—shows up in this map. The numbers are upper bounds with respect to the exogenous hazard, which multiplies in $e^{-\lambda_0 t}$ at each cell's collapse date and therefore compresses the slow histories hardest, steepening the gradient further.

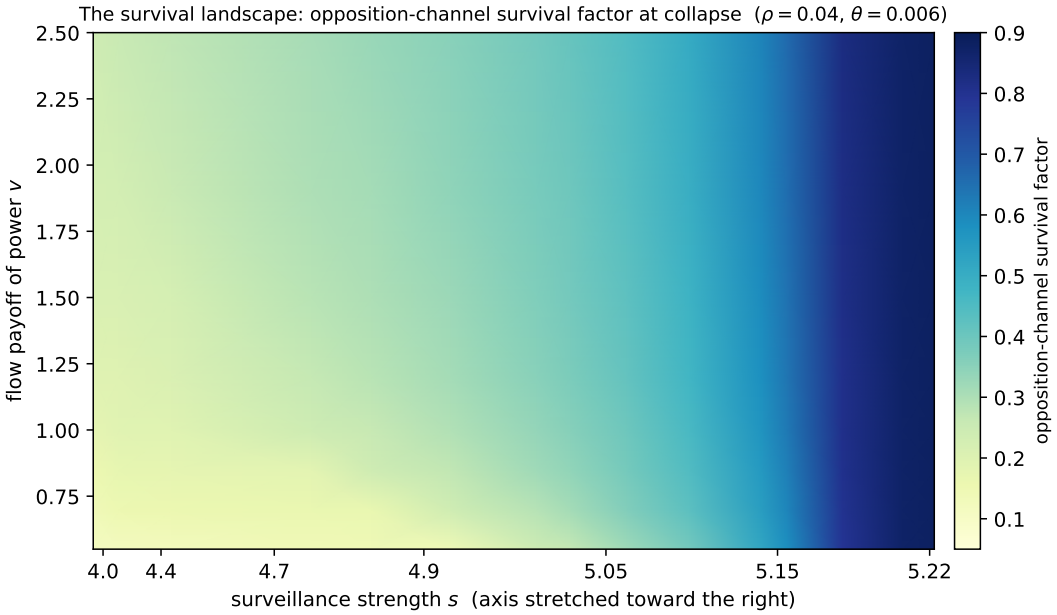


Figure 3: The survival landscape: the probability of outliving the opposition, absent exogenous ousting. The horizontal axis is stretched toward the right, where the landscape is steep; an exogenous hazard multiplies each cell by $e^{-\lambda_0 t_{\text{collapse}}}$ and steepens the gradient.

4.6 The price of coexistence

The third map prices what the first two describe. The regime's ceiling is the safe haven v/ρ: the value of office once coexistence is over. No policy attains it from birth—getting there costs money and runs risk—and the shortfall splits exactly in two. Writing V for the value of the computed policy at the common initial state and T for its collapse date,

$$\frac{v}{\rho} - V \;=\; \underbrace{\int_0^T \big(e^{-\rho t} - \Lambda(t)\big) v \, \mathrm{d}t + \Big(e^{-\rho T} - \Lambda(T)\Big)\frac{v}{\rho}}_{\text{the toll}} \;+\; \underbrace{\int_0^T \Lambda(t)\, c\, \mathbf{1}\{\text{mobilised at } t\} \, \mathrm{d}t + \sum_i \Lambda(t_i)\, \kappa}_{\text{the bill}}.$$

The toll is value lost to the risk of being ousted while the opposition still stands; the bill is what the regime pays its apparatus. Their sum is the price of coexistence: what the regime would give, at birth, to have the opposition vanish costlessly. Figure 4 maps the attained share $V/(v/\rho)$ on the plane of Figures 2 and 3.

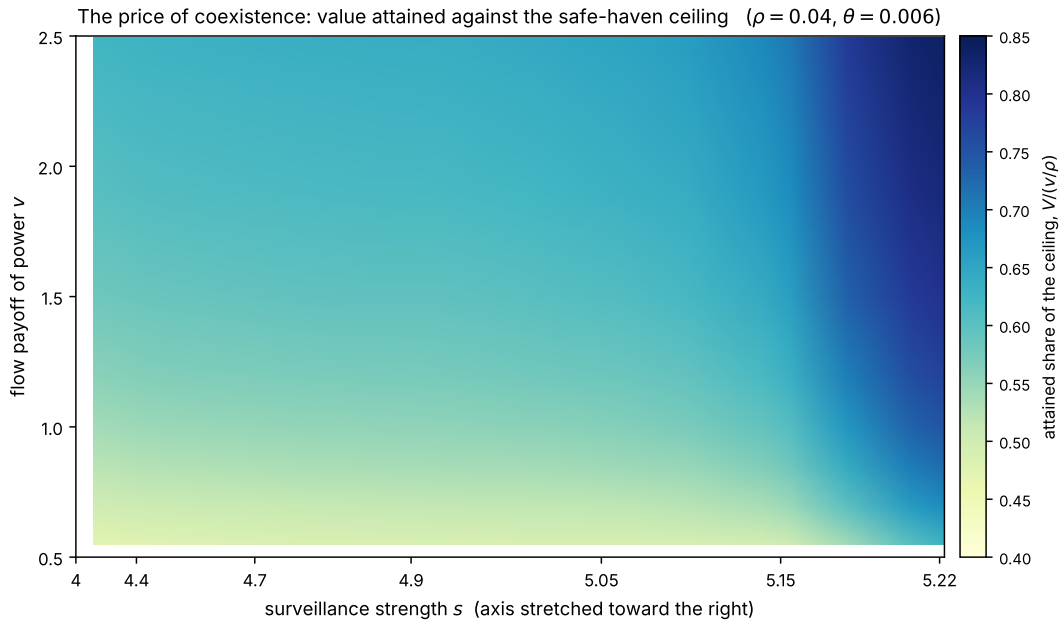


Figure 4: The price of coexistence: the share of the safe-haven ceiling v/ρ attained at birth, on the plane and axis of Figures 2–3.

Two findings are this map's own. First, stakes buy efficiency, not only survival: the attained share rises with v at every surveillance level—Corollary 1 proves the direction, the map supplies the magnitude—from 45 percent of the ceiling in the poor, weakly surveilled corner to 84 percent in the rich, nearly feasible one. At the benchmark the regime attains 59 percent—the opposition's existence costs it two fifths of what power would be worth without it. Second, the split. At the benchmark, 77 percent of the price is toll and 23 bill: the regime's main expense is not repression, it is risk. Near feasibility the split turns—a quick kill leaves little time to die in, and at low stakes the toll falls to under a third of the price. Surveillance does not just shrink the price of coexistence; it changes who is paid, from risk borne to wages paid.

4.7 Reading the maps together

Figures 2–4 share one plane, and an accounting identity runs through them: log-survival is minus the accumulated hazard, the hazard accumulates over the phases of Figure 1, and what the hazard does not take from the ceiling is either paid to the apparatus or kept. Where the risk lives is therefore a computable question, and for the three histories of Figure 1 the answer is stark. The waves carry most of the risk everywhere—75, 70, and 62 percent of accumulated hazard in panels (a), (b), and (c)—at a per-cycle cost that is nearly a technological constant: 0.09 to 0.11 in log-survival, so that at benchmark stakes each ordinary wave the regime

chooses to run costs it roughly a tenth of its chance of outliving the opposition. Up to that constant, the count of Figure 2 is a price list. The rest of the risk sits where the quiet is. The surprise ending's last lull contributes 22 percent of its history's hazard—as much exposure as three ordinary cycles—while its final crackdown, the most violent stretch of the whole history, contributes 8: the dangerous part of the surprise ending is the waiting, not the blow. The silent ending is costliest of all, its dying tail accumulating 38 percent of its total, and across the panels survival orders sudden, surprise, silent: 0.37, 0.27, 0.21. Speed is safety. The history with the fewest arrests is the one most likely to end with the regime gone first.

Read together, the maps also resolve what looks like a tension. The count rises with the stakes, and so does survival: the regime runs more crackdowns exactly where it lives longest. Choice and consequence, not contradiction. Higher stakes tighten the band, so each cycle runs at lower stocks and cheaper hazard; the regime buys more cycles and less exposure at once, and Figure 4 shows the purchase paying—more of the ceiling attained wherever more is at stake.

What the stakes axis measures deserves one careful paragraph. The flow value v is the value of office relative to the costs of repression—everything on the maps is measured in units of c and κ—and it is itself net: extraction minus the upkeep of consolidation. The dark corner of Figure 4 is therefore not wealth; it is rents with cheap repression—office worth a great deal to its holder, in a polity where putting courts and security services on a campaign footing costs that holder little. That is the long-lived rentier autocracy of the empirical record, not the rich institutionalised state, where raising κ and c is much of what constitutions are for. Read this way, the maps locate the contest: where the stakes of office are high, every unit of consolidation pace, every cheapening of the crackdown, and every improvement in surveillance pays its holder more than it would anywhere else. The high-stakes polity is where the parameters of this model matter most.

The mechanism also binds itself to data, jointly: distinct mobilisation and stand-down triggers at comparable capacity; band widths that move with the political cost of mobilisation; institutional erosion that continues while overt repression declines; and, where a final crackdown occurs, a launch below every earlier launch. An arrest series alone identifies none of this—the predictions need measures of opposition strength and of its base—and hysteresis or a long lull by itself would not separate the model from calendar-driven or informational alternatives.

Timescale and survivorship. One honest caveat belongs in the main text. Under hazards calibrated in Appendix B to observed autocratic failure rates, a regime following the computed plan reaches the last stand-down with probability 0.40, the terminal launch with 0.30, and completes the history with 0.27 through the opposition-linked channel alone; under the

illustrative decomposition of the discount rate, the exogenous hazard multiplies these by $e^{-\lambda_0 t}$, leaving a total completion probability near 0.07. Completion is thus rare and concentrated by construction among the durable; the majority is ousted mid-cycle, with the opposition still standing—so most observed autocracy deaths should occur inside the wave phase, a censoring prediction in its own right. The endgame carries about seven percent of the plan's time-zero value, and almost all of that—over ninety-nine percent—is already secured by the attrition path; the strike's acceleration of collapse adds less than one percent of it. This, together with the launch-prohibition experiment of Appendix B, is why the wave-phase policy is numerically insensitive to how the endgame is modelled. A model period is read as a year, consistent with the hazard calibration: ordinary lulls then last about four years, the last lull about fourteen, and the completed benchmark history roughly seven decades—timescales of institutional orders, not of individual rulers.

5 Discussion

The trivial insight from the analysis is that surveillance strength works unambiguously in the regime's favour. There is nothing counterintuitive here. Regimes with stronger surveillance capabilities do better on all counts: they are more likely to survive, and they extract more surplus in expectation. Less obvious is perhaps what surveillance does to the rhythm. The number of waves is single-peaked in surveillance strength, at every level of what the regime can extract: without knowing where a polity sits, no prediction links the number of waves to surveillance capability.

Surveillance also does other work. Its technology—more effective when oppositions are larger, simply because it is easier to find an opponent—shapes waves and lulls and predicts that regimes will rarely begin their rule with harsh crackdowns. Initially, they let erosion do its quiet work.

Flow payoffs capture a country's riches: what can be extracted, net of consolidation costs. The analysis shows that higher flow payoffs also work unambiguously in the regime's favour—higher survival, and a larger share of the attainable value, a true double whammy. At the same time they increase the number of crackdowns and lulls: when there is more at stake, the regime polices more tightly—it tolerates less regrowth, strikes more often, and relies less on erosion alone. High-stakes polities should thus oscillate more between calm and harshness, not less. Countries in which the value of holding power has reached unusual heights are, in this light, of particular interest; the contemporary United States comes to mind.

How costly is it to start a new crackdown? Intuitively, one might think that as κ grows, lulls

disappear. The computations of Section 4.4 show the opposite: lulls win. A larger fixed cost does imply fewer crackdowns though and, intuitively, they become more sustained. But aside from extreme parameter configurations large crackdown costs never eliminate waves.

Impatience works in one direction only: crackdowns are an investment, so an impatient regime cracks down less. It is also more vulnerable: the impatient regime represses less, forgoes the final blow, and is markedly less likely to outlive its opposition (Appendix B).

Erosion of opposition strength through sheer consolidation sets the pace. Slower consolidation lengthens the rhythm without changing its character while faster consolidation shrinks the final-crackdown region: the regime just lets erosion do its work and costly crackdown phases become less appealing.

The three archetypes deserve their own word. Which ending a regime reaches is not a type it is born with but a selection by the parameters: surveillance orders the endings—silent for weak detection, surprise in the middle, sudden for strong—and the stakes of office walk the same path, from attrition at low stakes to the sudden ending at high. This also answers the question the introduction promised to settle: when is a lull a thaw? During the waves, every lull is a genuine recovery—the opposition regains its strength, and only the base beneath it has thinned. The last lull of the surprise ending looks the same from the outside and is the opposite: the recovery caps below the old peaks and then turns down, erosion quietly doing what crackdowns no longer need to do. What separates the two is not whether erosion continues—in the model it always does—but what the recovery does: in a thaw it holds; in the last lull it caps below the old peaks and turns down.

There are a few extensions that follow-up work might tackle. Right now the consolidation pace is given exogenously and its costs are folded into the regime's flow payoff. What if the regime had to choose the rate of consolidation efforts at the outset? The relative costs of consolidation and crackdowns would then play a key role, presumably with obvious comparative statics (which is why I reserve this for later work).

The other obvious parameter that is exogenous in the present analysis and would warrant an endogenous choice is surveillance strength, in particular as it does so much work in the present draft. But there are two fundamental reasons why I decided to keep it exogenous for now. First, whatever would be chosen, the choice would not affect the three wave archetypes I identified. Second, a new regime investing in a surveillance state would send an important signal to its opposition and that would require a rather different model of how the opposition behaves.

Which leads naturally to the third and, perhaps, most fundamental extension: an opposition

that acts strategically. Whether a strategic opponent would compel the regime to avoid waves and resort to constant crackdowns seems questionable though.

Does this paper have a core that could be taken to empirics? It started with Klemperer and a stylised fact and tried to make sense of it. Whether anybody would be able to turn it into something that could be structurally estimated is a tantalising thought. As the model also allows—in its most extreme parameter configurations—for immediate crackdowns that only cease when the opposition is fully obliterated, it may lend itself to a structural approach rather than to attempts at falsification.

6 Conclusion

An authoritarian regime that wants to extract maximal surplus while staying in power may alternate crackdowns and lulls for a long time, then fall quiet altogether, and either strike once more or never strike again—the three archetypes of Figure 1.

It is tempting to ask about historical antecedents for these archetypes.

I started with Klemperer's diary—the elephant in the room on these final stretches. Does the Nazis' way to absolute power map into one of the archetypes? This needs to be handled with care; the gap between my model and history is huge. So, with extreme caution:

(i) 1933 saw an extremely harsh crackdown almost immediately, without much time for the opposition to grow—in the spirit of this paper, closer to a corner than to the generic picture. (ii) Yet 1936 was a real lull: the Olympic Games, far fewer inmates in the concentration camps than in 1933—though quietly rising again—and many international observers surprisingly positive about Germany.[3] (iii) The biggest wave of Jewish emigration from Germany came not in 1933, when Thomas Mann left, but only in 1938/9, after the apparent thaw of 1936 had truly ended. (iv) In many ways, Germany's going to war can be read as a final crackdown internally.[4] So, if a label is wanted, the shape is closest to the surprise ending which, presumably, reflects some of the surprise, also internationally, when appeasement had failed.

Whether other historical examples map into the three archetypes or not is a question left for the reader. And so is the question whether some of the (pseudo)authoritarian regimes of the second half of the 2020s will fall into any of them.

The sliver of hope that the paper offers is this: they can be toppled.

[3]Another great diary of the twentieth century serves as a chilling example—Henry "Chips" Channon's (Channon, 2021). A British MP at the time, he went to the Berlin Games and was thrilled by what he saw, particularly enjoying Göring's do.

[4]At the outbreak of war in September 1939, the Gestapo re-arrested thousands of former party and union functionaries from lists prepared in advance—a strike launched from calm.

Appendix A: Proofs

Formal statements are given with the domains under which they are proved; the objective and the admissible policy class are those of Section 3.

Proposition 1, and the zero-cost dominance sketch. Consider a stationary rule assigning to each x a single drift $F(x) \in \{a_K(x),\, a_K(x) - m(x)\}$, with finitely many sign changes. We adopt the Filippov convention throughout, which describes the closure of the admissible finite-switching class: at a discontinuity where the drift keeps its sign, Filippov solutions pass through monotonically; where the drift opposes (positive below, negative above), they slide, i.e. remain constant—the relaxed limit of ever-faster admissible switching; where it repels (negative below, positive above), solutions leave the point monotonically. Finitely many sign changes partition the state space into maximal intervals of a.e. constant drift sign, and on each interval every Filippov solution is monotone. Non-Lipschitz zeros of F may be crossed, but only monotonically, since F is sign-preserving in a neighbourhood of any such crossing. Concatenating across boundaries, every solution is globally monotone in time up to stagnation at a rest or sliding point, so no non-constant trajectory revisits a value. □

Why a memoryless policy should do at least as well (remark). A mode-dependent hysteresis rule is stationary in (x, j) and cycles at $\kappa = 0$; but with free switching $V^0 \equiv V^1$, the mode carries no value-relevant information, and any hysteresis cycle uses, on part of each leg, the drift that does not maximise the Hamiltonian of the common value function—so we expect it weakly dominated by memoryless feedback. Completing the verification argument is left aside; that an at-least-as-good memoryless policy exists is a supported expectation, not a theorem; weak dominance would not by itself exclude a costless cycle from being co-optimal under Hamiltonian ties. The memoryless hover is itself a Filippov sliding mode—the relaxed limit of admissible cycles of vanishing amplitude and diverging frequency, not an admissible finite-switching policy—so what $\kappa > 0$ excludes within the admissible class is not switching but vanishing amplitude.

Proposition 2. Threshold policies (ℓ, h) are compared through the recurrent-cycle value $V_K(\ell, h) = [P^0 + e^{-B^0}(-\kappa + P^1)]/[1 - e^{-(B^0+B^1)}]$, with P^j, B^j the recurrent legs' payoffs and survival exponents; a common initial state x_0 is connected to the band by an initial transient leg whose payoff and exponent are evaluated separately and are continuous in (ℓ, h), so the transient plays no role in the band-collapse argument that follows. Restrict to the closure of $\{(\ell, h) : \max\{\varepsilon, x_s(K) + \delta\} \le \ell \le h \le \bar{x} < K\}$ for fixed $\delta > 0$ and suppose an optimum exists satisfying standard verification. On this set descent and lull speeds are bounded below, so along any band-collapsing sequence the cycle length $\tau \to 0$ and the amortized charge $\kappa/[(\rho + \lambda_1)\tau] \to \infty$ while payoffs are $O(\tau)$: for $\kappa > 0$ no optimum is collapsed. For $\kappa = 0$,

collapsed sequences have the relaxed-hover limit $[v - c\,a_K(x)/m(x)]/[\rho + \lambda_1(x)]$, the time-average cost of holding that constant state; we claim no more than this limit (in particular, no dominance over the permanently mobilised path, which follows a different, lower-hazard trajectory). The comparative static (larger κ, longer cycles, deeper crackdowns) holds ceteris paribus in all computations and is not claimed as a theorem. □

Single crossing, the floor, and the frontier. *Assumption (regular single crossing).* For each $K \le K_0$, $m - a_K$ has exactly one sign change on $(0, \bar{x}]$, from negative to positive, with $(m - a_K)' > 0$ at the root. Verified by direct root counts at all (s, K) used; for large K outside the verified range three crossings are possible and all statements are restricted to the current descent basin. *Lemma (floor monotone).* Raising K lowers $m - a_K$ pointwise on $(0, K)$, so the unique regular root moves strictly right; as $K \to \infty$ it converges to the smaller root of $gz^2 - \bar{m}z + g/s^2$ when $\bar{m}s \ge 2g$, which is exactly when real roots exist; when $\bar{m}s < 2g$, removal loses to recruitment at every stock and the balance point diverges with K. *Frontier.* $K^\star$ solves $a_K(\varepsilon) = m(\varepsilon)$, requiring $\bar{m}s^2\varepsilon < g(1 + s^2\varepsilon^2)$; above it the mobilised drift is positive below the floor so no policy reaches ε in finite time at fixed K (and no mode process descends faster than full mobilisation); strictly below it the descent time $\int \mathrm{d}z/(m - a_K)$ is finite; at $K^\star$ the threshold is a rest point. $s_c = \sqrt{g/(\varepsilon(\bar{m} - g\varepsilon))} = 5.30$; s_{init}, which solves $K^\star(s) = K_0$, is 5.19.

The menu at fixed capacity. (i) Coexistence ceiling $v/(\rho + \chi)$. (ii) *For $K \ge K^\star$* (non-eliminating case): permanent mobilisation converges to the floor with value at most $(v - c)/(\rho + \chi)$; for $K < K^\star$ it eliminates and falls under (iv). (iii) For $x_0 \le K$, stand-down values lie between $v/(\rho + \lambda_1(K))$ and $v/(\rho + \lambda_1(x_0))$. (iv) Write P^j for the flow payoff accumulated over one leg of a cycle in mode j, $P^j = \int_0^{T_j} e^{-\int_0^t (\rho + \lambda_1)\,\mathrm{d}u} \pi_j \,\mathrm{d}t$, and $B^j = \int_0^{T_j} (\rho + \lambda_1)\,\mathrm{d}t$ for the leg's survival exponent, so that e^{-B^j} is its discount factor. With $v \ge c$, $V_E \ge -\kappa + e^{-(\rho + \bar{\lambda})T^E} v/\rho$, where T^E is the fixed-base elimination time from x_0 and $\bar{\lambda} = \max_{x \le x_0} \lambda_1(x)$, so elimination dominates all never-eliminating policies whenever $T^E < (\rho + \bar{\lambda})^{-1} \ln[(v/\rho)/(\kappa + v/(\rho + \chi))]$, nonvacuous iff $\kappa < v\chi/(\rho(\rho + \chi))$ (0.4 against 10.23 at benchmark). Sufficient only: when it fails, elimination may still pay, since actual coexistence can sit far below its ceiling; this bound is a minimum for the prize, not the whole of it.

The diagonal value match $K^\dagger$. The frozen launch decision is state- and mode-contingent, and the scalar below compares only two policies from the off-state diagonal $x = K$—a third policy, such as a further ordinary wave, may dominate both, so $K^\dagger$ is a named value match rather than the launch threshold of the full frozen switching problem. Its existence is assumed; for large κ the comparison set can be empty. The definition: $K^\dagger$ is the largest K at which, with the apparatus off and $x = K$, immediate mobilisation pressed to collapse at fixed K

weakly beats indefinite stand-down at that K. At benchmark parameters: $(s, K^\star, K^\dagger) =$ (3.70, 0.0393, 0.0392), (4.17, 0.0527, 0.0526), (4.55, 0.0763, 0.0760), (4.76, 0.1047, 0.1043): the diagonal comparison favours elimination essentially up to the frontier.

Comparative statics in the large. Fix an admissible policy. Along its path $\Lambda(t)$ does not depend on (v, c, κ), so the objective equals $v\int_0^\infty \Lambda\, \mathrm{d}t - c\int_0^\infty \Lambda\, \mathbf{1}\{\text{mobilised}\}\, \mathrm{d}t - \kappa\sum_i \Lambda(t_i)$: affine, with nonnegative coefficients on v, $-c$, and $-\kappa$. The value is the pointwise supremum of these affine functions over admissible policies, hence convex, nondecreasing in v, nonincreasing in c and κ. At any point of differentiability, Danskin's theorem identifies the gradient with the coefficient vector of an optimal policy. For the corollary, write a policy's payoff as $vA_\pi - B_\pi$ with $A_\pi, B_\pi \geq 0$ and divide by v/ρ: each policy's share $\rho(A_\pi - B_\pi/v)$ is nondecreasing in v, and the supremum inherits the monotonicity. □

Appendix B: The numerical approach

In the diagrams of Section 4, the benchmark occupies the cell $(s, v) = (4.76, 1.5)$. Surveillance values are sampled on a grid in $\sigma = 1/s$ and reported to two decimals in s; all computations use the exact grid values. Benchmark parameters: $v = 1.5$, $g = 0.5$, $\bar{m} = 0.9$, $s = 4.76$, $\rho = 0.04$, $c = 0.5$, $\kappa = 0.4$, $\chi = 0.015$, $D(x) = 0.3x^2$, $\varepsilon = 0.02$ (both margins), $K_0 = 0.5$, $x_0 = 0.10$, $\theta = 0.006$. The hazard scale is calibrated to the survival evidence: with ρ read as pure time preference plus an exogenous hazard (illustratively $0.02 + 0.02$), total annual failure hazards along the path run between 3.5 and 5 percent, in line with observed autocratic failure rates (Svolik, 2012).

Method. The two value functions solve the coupled quasi-variational inequalities

$$\max\big\{v + a_K(x)V^0_x - \theta V^0_K - (\rho + \lambda_1(x))V^0,\ V^1 - \kappa - V^0\big\} = 0,$$

$$\max\big\{v - c + (a_K(x) - m(x))V^1_x - \theta V^1_K - (\rho + \lambda_1(x))V^1,\ V^0 - V^1\big\} = 0,$$

with $V^j = v/\rho$ on the base-collapse boundary $K \leq \varepsilon$ and on the portion of the stock boundary that trajectories can actually cross; for $K > K^\star$ the stock boundary is unreachable and carries interior values, a distinction the Verification paragraph develops. Because K falls at rate θ under every policy, it acts as a clock, and the problem is solved by semi-Lagrangian backward induction in K from that boundary: characteristics advance by explicit Euler within each capacity step, hazards enter through survival discounting, continuation values by linear interpolation, and the same-slice switching update is the coupled obstacle form $V^0 = \max\{C^0, C^1 - \kappa\}$, $V^1 = \max\{C^1, C^0\}$, computed from continuation values first and hence order-independent. Steps crossing ε receive their within-step crossing time. The grid

carries the wave phase down to an exact terminal slice just above the frontier, described under Verification; paths are simulated forward on the resulting hybrid policy.

Verification. The stock-collapse boundary requires care. Above the frontier both drifts at $x = \varepsilon$ point into the interior: no trajectory can reach stock collapse before capacity has eroded to $K^\star$, so the value function is *discontinuous* there—its interior limit sits strictly below the safe-haven ceiling v/ρ that a Dirichlet condition imposes at the boundary node. A scheme that interpolates linearly across that jump diffuses the collapse prize into interior states, and the effect is not hypothetical: in an uncorrected solve it biased terminal launches early, producing a grid launch at $K = 0.134$ that is dominated within its own admissible family by simply waiting. The reported solution therefore removes the boundary from the grid's reach. Backward induction starts from an *exact terminal slice* at $K_{\mathrm{sw}} = K^\star + 0.025$: at every stock node the off-mode value is the better of attrition and the optimal single-launch family (off for a waiting time, then mobilise and hold to collapse), the mobilised value the better of standing down and holding to collapse, all integrated in continuous time (DOP853 at tolerance 10^{-10}, cross-checked against RK45) with each cell's own stakes v. Above the slice the mobilised floor exceeds viability, so grid characteristics never meet the stock boundary; the solver counts boundary contacts, and every reported run logs zero. The slice values are family-restricted, and the restriction is tested explicitly: a three-parameter deviation family—wait, run an intervening crackdown, stand down, wait again, then launch and hold—nests the single launch and covers an extra crackdown during the final lull and a stop-and-restart of the grind; at the benchmark, at the other two panels, and at seven further boundary cells spanning the thin verdicts, it never improves on the better of the single launch and attrition, and at the benchmark its optimum collapses onto the single launch itself. What the construction delivers is an *optimised hybrid policy*, not a certified global optimum; the deviation evidence is what stands behind the slice. The corrected solver reproduces the benchmark to displayed precision—nine waves, a final lull of 14.2 periods, launch capacity 0.1182, a grind of 6.6, a value margin of +0.24 over attrition—and re-solving every audited (v, s) cell under it changes no classification: every attrition verdict returns with its margin unchanged to within 0.01, and every final-crackdown verdict stands. The wave phase is numerically independent of the endgame at the benchmark: forbidding terminal launches below $K = 0.15$ moves wave triggers by at most 0.004 and $V^0(x_0, K_0)$ by less than 0.005 on a scale of 22. Relative to the coarse grid, two assignments change, both toward attrition—the strike does not survive the exact comparison at benchmark stakes for $s = 4.55$, nor at $(v, s) = (0.7,\ 5.00)$—and the rest stand. The deciding margins run from 0.012 to 3.1, and the thinnest verdicts—no final crackdown at $v = 0.55$ near $s = 5$ (margins 0.012–0.013) and at benchmark stakes for $s \leq 4.55$ (margins 0.05–0.09)—are decided narrowly. At $s = 5.00$ the terminal optimum is shallow in the launch date: corrected re-solves with different handoff placements launch between

$K = 0.21$ and 0.24, always after at most an ordinary-length pause and always well above the frontier 0.1836; the figure draws the single-launch optimum from the last stand-down, $K = 0.2262$. Table 1 collects these checks in one place. At $s = 5.21$ the computed optimal delay before the single crackdown is exactly zero: near feasibility the rhythm never forms. The surprise/sudden distinction uses a fixed cutoff, the final lull exceeding twice the average ordinary lull, defined for $n \geq 1$; at $n = 0$ the single crackdown is classified by the economic comparison alone. Finally, the trigger counterfactual behind the lull discussions of Sections 4.3 and 4.7: along the benchmark's final lull the off-mode trajectory never reaches the wave-phase trigger $x_H(K)$—its closest approach is 0.035 below, and the lull's peak stock of 0.159 sits under the last ordinary trigger of 0.178—so an unchanged containment rule would have remained quiet through the whole lull.

cell	construction	waves	final lull	launch K	collapse	margin
benchmark (1.5, 4.76)	grid, $\Delta K = 5 \times 10^{-4}$	9	14.24	0.1182	70.2	+0.238
	corrected hybrid, $\Delta K = 2.5 \times 10^{-4}$	9	14.17	0.1182	70.3	+0.239
(a) (1.5, 5.00)	corrected, handoff $K^\star + 0.025$	8	4.98	0.2262	56.1	+2.33
	corrected, handoff $K^\star + 0.05$	7	4.21	0.2086	57.4	—
	corrected, handoff $K^\star + 0.10$	7	7.08	0.2415	55.3	—
(c) (1.5, 3.70)	both constructions	8	—	no launch	80.0	−0.05

Table 1: Audit of the three histories and the thin claims across constructions, grids, and handoff placements. Margins are value differences at the last stand-down. The $s = 5.00$ launch date is a shallow optimum; its handoff rows quantify the resulting timing uncertainty.

Robustness. Varying the residual hazard, with a full re-solve at each value: the final crackdown persists down to $\chi \approx 0.011$ and the ending is attrition below it, including at $\chi = 0$—the residual hazard thereby joins surveillance, stakes, and the consolidation pace as a selector of the ending. Smoothing the disappearance of χ near viability—replacing the constant by $\chi(x) = \chi_0\big(1 - e^{-(x-\varepsilon)/w}\big)$, which leaves hazards at ordinary stocks unchanged while phasing the collapse prize in continuously on the stock margin (the base margin keeps its jump)—is run as a sweep in the width w, from 2×10^{-4} to 2×10^{-2}. The governing scale is the gap between the floor and viability, about 0.003 at the relevant capacities. For widths below that gap the benchmark's surprise ending persists essentially unchanged: at $w = 2 \times 10^{-4}$, nine waves, a final lull of 13.7, launch at 0.123, margin +0.29. As the width passes the gap the lull compresses continuously and the ending turns sudden—launch 0.166 at $w = 3.5 \times 10^{-3}$, launch 0.335 with a forty-period grind at $w = 10^{-2}$—and at the thin cell $s = 4.55$ widths beyond the gap flip the silent ending into a late strike. Collapse need not be discontinuous; the extinction reward must be concentrated below the floor. The calibration, moreover, pins the total failure hazard, not its split into removable and permanent parts—and the split prices the elimination prize. A hazard-preserving sweep makes the point

exact: with pure discounting $r = 0.02$, the family $(\lambda_0, \chi, \rho) = (0.035 - \chi,\ \chi,\ 0.055 - \chi)$ leaves every pre-collapse hazard and its survival discounting unchanged while moving the removable share. At $\chi = 0$, 0.005, and 0.010 the benchmark keeps its nine waves but ends in attrition (margins -0.06), with total survival to collapse 0.045; at the benchmark split $\chi = 0.015$ the strike returns ($+0.24$), collapse advances from 80 to 70, and total survival rises to 0.066. Identical observed hazards, different endings: the division between removable and permanent risk selects the ending, and the benchmark split is illustrative rather than identified by aggregate failure rates. Re-solving the three panel histories under five collapse rules—the baseline; base-side thresholds moved to 0.03 and 0.01; no base threshold, with K eroding toward zero and collapse through the stock margin alone; and bounded decline, which caps the *autonomous* drift at $\dot{x} \geq -0.05$ while leaving crackdown removal unrestricted and both thresholds at baseline—and carrying each through the terminal stage gives Table 2: the ending is invariant—a final crackdown in panels (a) and (b), with launch capacity and collapse date unchanged to displayed precision, attrition in panel (c)—while the margins move with the rule as the value of waiting moves with the deadline. The bounded-decline cap never binds along the computed paths, whose natural decline tracks the consolidation rate; its row is an implementation check—evidence that the cap never binds, not robustness to a binding bound on autonomous decline. The stock-side threshold is varied separately, with the base threshold fixed: at $\varepsilon_x = 0.015$ the frontier drops to $K^\star = 0.038$ and the benchmark ending flips to attrition (margin -0.065), while at $\varepsilon_x = 0.025$ the floor never clears viability and elimination is feasible from the outset. The viability threshold moves the frontier exactly as Proposition 3 prescribes, and the ending moves with the frontier; the benchmark sits where the floor clears viability narrowly, the same gap the smoothing sweep exploits. The detection exponent is varied in the same spirit: replacing the quadratic by $m(x) = \bar{m}\,(sx)^q/\big(1 + (sx)^q\big)$ at $q = 1.5$ and $q = 3$ leaves the wave machinery intact—at the panel technologies, $q = 3$ still produces nine to ten cycles with ordinary lulls near three and a half periods—while moving the frontier sharply, as the frontier condition prescribes: stronger small-stock detection ($q = 1.5$) makes elimination feasible from the outset at the panel technologies, weaker ($q = 3$) drops the frontier to $K^\star \approx 0.02$ and sends all three histories to attrition. The taxonomy is relocated, not lost: rescaling surveillance to restore a benchmark-sized frontier recovers all three endings in the usual order—at $q = 1.5$, silent at $s = 2.26$, surprise at $s = 2.27$ with an eleven-period last lull, sudden at $s = 2.28$; at $q = 3$, silent at $s = 5$, surprise at $s = 9.9$ with a sixteen-period lull, sudden at $s = 10.4$. The exponent chooses which surveillance levels are interesting; the architecture—waves from the fixed cost, endings from the frontier—survives it. The fixed-cost sweep reported in Section 4.4, κ from 0.1 to 6.4 plus the high-stakes and cheap-mobilisation cells at $\kappa = 1.6$, runs on the same boundary-corrected solver at the standard grid, as does the hazard decomposition of Section 4.7. The value map of Section 4.6 is computed cell by cell as

the value of the simulated optimal path, with the toll and the bill accumulated along it so the decomposition closes by construction; in the near-feasibility columns, where the terminal-slice construction truncates waves that still pay, each cell takes the better of the corrected and the full-grid construction—both values of feasible, deviation-tested policies—and every row of the map is then monotone in surveillance, as it must be. Figures 2 and 3 are generated from the fine-grid full-grid construction; in the value map's spliced columns the winning policy is that same construction, and each cell's reported count, survival, and value come from one simulated path per cell. Across the 136 cells the value map's counts agree with Figure 2 cell for cell, except in five cells on band boundaries, where the two constructions differ by one cycle; the values there agree to 0.002. Opposition rebuilding is checked the same way: replacing the clock by $\dot{K} = -\theta + \beta x$, with $\beta = 0.005$ and 0.01—capacity still falls along every computed path, with erosion slowed by up to a third at the opposition's peaks—leaves all three panel archetypes in place. Histories lengthen. The final crackdown strengthens, its benchmark margin rising from +0.24 to +0.42 and +0.58 while the final lull shortens toward ten periods, and the silent ending's attrition date moves out from 80 to 89 and 99 periods: rebuilding makes the quiet expensive, so the regime waits less and gains more from finishing. The solver change is a state-dependent time step, $\mathrm{d}t = \mathrm{d}K/(\theta - \beta x)$, with the terminal stage integrating K as a state. Crackdown-driven erosion, $\dot{K} = -\theta - \eta$ while mobilised, is checked the same way at $\eta = 0.003$ and 0.006: the faster clock during crackdowns raises the opportunity cost of quiet, so the benchmark's surprise ending turns sudden—the final lull falls from 14 to 6.6 and 4.9 periods—wave counts fall, margins rise, and the weak-surveillance panel remains attrition. The map of Figure 2 is computed at the fine grid ($\Delta K = 2.5 \times 10^{-4}$, 1400 stock nodes); the standard grid agrees at every cell to within one cycle, and the features carrying the text survive the refinement: each row of the map is single-peaked in s, and the staircase's final steps—six, five, three, two, none on the benchmark row—are unchanged. The map of Figure 2 draws the count's level sets by linear interpolation between the computed grid points, lightly smoothed. The interpolation only draws; it does not by itself establish that the intermediate counts occur. The direct probes establish it where the bands are thinnest: along the benchmark row, four waves on $s \in [5.118, 5.133]$ and one on $s \in [5.165, 5.171]$; in the stakes direction, five waves at $v \in \{0.60, 0.625, 0.65\}$ for $s = 5.025$. The drawn band edges match these intervals to within 0.006, and every count down to zero occupies a band of positive width—the staircase has no knife edges. The unsmoothed cell-level map accompanies the code. Impatience is swept at the benchmark cell: at $\rho = 0.055$ and 0.07 the wave count falls to eight and seven, the final crackdown loses its margin (-0.06, -0.01), and the opposition-channel survival factor falls from 0.268 to 0.204 and 0.186—the impatient regime represses less and dies more. The corrected terminal timing moves the survival factors of Figure 3 by at most 0.03. Code accompanies the paper.

	(a) $s = 5.00$			(b) $s = 4.76$			(c) $s = 3.70$	
Collapse rule	launch K	collapse	margin	launch K	collapse	margin	collapse	margin
baseline (both margins 0.02)	0.226	56.1	+2.33	0.118	70.2	+0.24	80.0	−0.05
base threshold 0.03	0.226	56.1	+2.24	0.118	70.2	+0.07	78.3	−0.07
base threshold 0.01	0.226	56.1	+2.40	0.118	70.2	+0.36	81.4	−0.05
no base threshold	0.226	56.1	+2.40	0.118	70.2	+0.36	81.4	−0.05
bounded decline ($\dot{x} \geq -0.05$)	0.226	56.1	+2.33	0.118	70.2	+0.24	80.0	−0.05

Table 2: The ending under five collapse rules (defined in the text), carried through the terminal stage. Margins are the final crackdown's value advantage over attrition; in panel (c) no launch occurs and the best feasible launch loses.

References

Bermeo, N. (2016). On democratic backsliding. *Journal of Democracy*, 27(1):5–19.

Chenoweth, E. and Stephan, M. J. (2011). *Why Civil Resistance Works: The Strategic Logic of Nonviolent Conflict*. Columbia University Press.

Channon, H. (2021). *Henry 'Chips' Channon: The Diaries, Volume 1: 1918–38* (S. Heffer, Ed.). London: Hutchinson.

Clark, C. W. (1990). *Mathematical Bioeconomics: The Optimal Management of Renewable Resources*. Wiley, 2nd edition.

Davenport, C. (2007). State repression and political order. *Annual Review of Political Science*, 10:1–23.

Davenport, C. and Appel, B. J. (2022). *The Death and Life of State Repression: Understanding Onset, Escalation, Termination, and Recurrence*. Oxford University Press.

Di Lonardo, L., Sun, J. S., and Tyson, S. A. (2026). Repression and autocratic consolidation. *Games and Economic Behavior*, 159:393–424.

Esteban, J., Morelli, M., and Rohner, D. (2015). Strategic mass killings. *Journal of Political Economy*, 123(5):1087–1132.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443.

Gregory, P. R., Schröder, P. J., and Sonin, K. (2011). Rational dictators and the killing of innocents: Data from Stalin's archives. *Journal of Comparative Economics*, 39(1):34–42.

Guriev, S. and Treisman, D. (2019). Informational autocrats. *Journal of Economic Perspec-*

tives, 33(4):100–127.

Guriev, S. and Treisman, D. (2022). *Spin Dictators: The Changing Face of Tyranny in the 21st Century*. Princeton University Press.

Klemperer, V. (1998). *I Shall Bear Witness: The Diaries 1933–1941*. Translated by M. Chalmers. Weidenfeld and Nicolson, London.

Kuran, T. (1991). Now out of never: The element of surprise in the East European revolution of 1989. *World Politics*, 44(1):7–48.

Levitsky, S. and Ziblatt, D. (2018). *How Democracies Die*. Crown.

Lichbach, M. I. (1987). Deterrence or escalation? The puzzle of aggregate studies of repression and dissent. *Journal of Conflict Resolution*, 31(2):266–297.

Lorentzen, P. L. (2013). Regularizing rioting: Permitting public protest in an authoritarian regime. *Quarterly Journal of Political Science*, 8(2):127–158.

Ritter, E. H. and Conrad, C. R. (2016). Preventing and responding to dissent: The observational challenges of explaining strategic repression. *American Political Science Review*, 110(1):85–99.

Scarf, H. (1960). The optimality of (S, s) policies in the dynamic inventory problem. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Sciences*. Stanford University Press.

Spagat, M. (2001). The dynamics of repressive dictatorships. Working paper, Royal Holloway, University of London.

Stokey, N. L. (2009). *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Sullivan, C. M. (2016a). Political repression and the destruction of dissident organizations: Evidence from the archives of the Guatemalan National Police. *World Politics*, 68(4):645–676.

Sullivan, C. M. (2016b). Undermining resistance: Mobilization, repression, and the enforcement of political order. *Journal of Conflict Resolution*, 60(7):1163–1190.

Svolik, M. W. (2012). *The Politics of Authoritarian Rule*. Cambridge University Press.

Truex, R. (2019). Focal points, dissident calendars, and preemptive repression. *Journal of*

Conflict Resolution, 63(4):1032–1052.

Wintrobe, R. (1998). *The Political Economy of Dictatorship*. Cambridge University Press.