THE EXACT MINIMAX PEAK ON THE SUB-DISTURBANCE WINDOW: THE VALUE OF SCALAR ADVERSARIAL ADAPTIVE CONTROL FROM A START BELOW THE DISTURBANCE LEVEL

KOREA SUPERINTELLIGENCE LABS

Abstract. For the scalar adaptive-control game $x_{t+1} = ax_t + u_t + w_t$ with $\|w\|_\infty \le 1$, an unknown constant pole $a \in [-\Delta, \Delta]$, causal deterministic control and the worst-case peak $\|x\|_\infty$ as cost, Ho has recently determined the minimax value from rest, $\gamma^*(\Delta) = 1 + \Delta$, and from every start with $|x_0| \ge 1$, and leaves one region open: on the *sub-disturbance window* $0 < |x_0| < 1$ his results give only the bracket $[\,1 + \Delta|x_0|,\ 1 + \Delta\,]$, and he writes that "what is open is the exact value between them, which we defer." We determine the value on two explicit sub-intervals of that window, one at each of its ends. On the plateau $|x_0| \le \min(1, 1/(2\Delta - 1))$ the value is the *upper* end of the bracket, $\gamma^*(x_0, \Delta) = 1 + \Delta$; for every $\Delta \le 1$ that plateau is the whole window, so there the deferred question is closed outright. On the upper regime $|x_0| \ge (1 + \sqrt{1 + 8\Delta})/(4\Delta)$ the value is $1 + \max\big(\Delta s,\ M(\Delta s)/(2s),\ 1/s\big)$, where $s = |x_0|$ and M is Ho's own one-step exchange maximum; in general this is neither end of the bracket. Two mechanisms carry the matching lower bounds. On the plateau it is a *free move*: below the disturbance level one disturbance value repositions the state and is admissible for every parameter at once, so the adversary reveals nothing and commits to nothing, and the attack is a finite chain of such moves closed by one identification spike. On the upper regime it is three explicit finite attacks, the last of which drives the state back to rest and cashes Ho's own value theorem at the reduced and recentred parameter interval that survives. Three exact values follow: $\gamma^*(1/2, 3) = 3$ and $\gamma^*(1/2, 10) = 9$ lie strictly between the two ends of the bracket, while $\gamma^*(9/10, 3) = 37/10$ is its lower end exactly — so the window carries all three possibilities. A residual band between the two regimes remains open. Every theorem, proposition, lemma and corollary below has been formally verified in Lean 4 against *Mathlib*; no theorem here rests on a finite computation, and the computations we do report support no claim and are labelled as such.

1. Introduction

1.1. **The game.** Fix $\Delta > 0$. A scalar plant evolves by

$$x_{t+1} \;=\; a\,x_t + u_t + w_t, \qquad t \in \mathbb{N} = \{0, 1, 2, \dots\}, \tag{1}$$

where the state $x_t \in \mathbb{R}$ is measured exactly, the disturbance obeys $\|w\|_\infty = \sup_t |w_t| \le 1$, and the pole $a \in \mathbb{R}$ is *constant, unknown, and known only to lie in* $[-\Delta, \Delta]$. The controller is any causal deterministic law

$$u_t \;=\; K_t(x_t, x_{t-1}, \dots, x_0), \qquad K = (K_t)_{t\in\mathbb{N}} \in \mathcal{K}, \tag{2}$$

and the cost is the worst-case peak of the state — not an ℓ_2 gain, not a regret, not a discounted or averaged cost:

$$\gamma^*(x_0, \Delta) \;=\; \inf_{K\in\mathcal{K}} \sup_{|a|\le\Delta} \sup_{\|w\|_\infty\le 1} \|x(K, a, w)\|_\infty. \tag{3}$$

The interest of the problem is that Δ may be arbitrarily large: the parametric uncertainty is not a perturbation, and the controller must identify and regulate at once, under a criterion — the peak — that refuses to amortise one large excursion against many small ones.

Ho [1] determines this value from rest. His Theorem 1 is $\gamma^*(0, \Delta) = 1 + \Delta$, with the infimum and both suprema attained; the optimal law is set-membership certainty-equivalent deadbeat control at the *midpoint* of the consistent parameter interval,

$$u_t = -\operatorname{mid}(A_t)\, x_t, \qquad A_0 = [-\Delta, \Delta], \qquad A_{t+1} = \begin{cases} A_t \cap W_t, & x_t \neq 0, \\ A_t, & x_t = 0, \end{cases} \tag{4}$$

where $W_t = \{\alpha : |x_{t+1} - \alpha x_t - u_t| \leq 1\}$ is the set of parameters consistent with the transition just observed; we write K_Δ^{mid} for this law. From rest the second line gives $A_1 = A_0$, which is how [1] writes the initialisation; from a nonzero start the update runs from $t = 0$, as [1] prescribes when it treats such a start. From a nonzero start of magnitude at least the disturbance level, his Corollary 2 gives $\gamma^*(x_0, \Delta) = \max(|x_0|, 1 + \Delta|x_0|)$ for $|x_0| \geq 1$.

1.2. **What was deferred.** Between those two regimes lies a window that [1] leaves open. We quote its closing section verbatim.

> "What remains open is the sub-disturbance window $0 < |x_0| < 1$, where the lower bound $1 + \Delta|x_0|$ parts ways with the law's guarantee of $1 + \Delta$. Safety is not in question there, and the window's endpoints $|x_0| = 0$ and $|x_0| = 1$ both carry the value $1 + \Delta$; what is open is the exact value between them, which we defer."

So on $0 < |x_0| < 1$ all that was known is the bracket

$$1 + \Delta|x_0| \;\leq\; \gamma^*(x_0, \Delta) \;\leq\; 1 + \Delta, \tag{5}$$

whose two ends differ by a factor approaching $1 + \Delta$ as $x_0 \to 0$. The lower end is the one-step *identification spike*: whatever the control, an extreme parameter can be aligned with it and then the disturbance with the result. The upper end is the guarantee of the law (4), which survives a nonzero sub-disturbance start unchanged. Neither end is obviously the truth, and the source's two boundary values do not decide between them: at $|x_0| = 0$ the two ends are as far apart as they ever get and the value is the upper one, while at $|x_0| = 1$ the two ends coincide and the bracket says nothing about which of them the interior follows.

1.3. **Results.** Write $s = |x_0|$ throughout, and set

$$b^*(\Delta) \;=\; \begin{cases} 1, & 0 < \Delta \leq 1, \\ \dfrac{1}{2\Delta - 1}, & \Delta > 1, \end{cases} \qquad \text{so that} \qquad s \leq b^*(\Delta) \iff \big(s \leq 1 \text{ and } (2\Delta - 1)s \leq 1\big). \tag{6}$$

The product form on the right is the one we use: it is a single inequality valid for every $\Delta > 0$, with no case split at $\Delta = 1/2$ where $1/(2\Delta - 1)$ ceases to exist.

Our main result is that on $[0, b^*(\Delta)]$ the bracket (5) is tight at its upper end.

Theorem 1.1 (Theorem A)**.** *Let $\Delta > 0$ and let $x_0 \in \mathbb{R}$ satisfy $|x_0| \leq 1$ and $(2\Delta - 1)|x_0| \leq 1$. Then*

$$\gamma^*(x_0, \Delta) \;=\; 1 + \Delta.$$

Equivalently: for $\Delta > 1$ the value is $1 + \Delta$ at every start with $|x_0| \leq 1/(2\Delta - 1)$, and for $\Delta \leq 1$ at every start with $|x_0| \leq 1$.

Corollary 1.2. *For every $\Delta \in (0, 1]$ the deferred window is closed outright: $\gamma^*(x_0, \Delta) = 1 + \Delta$ for every x_0 with $|x_0| \leq 1$, endpoints included.*

Two features are worth isolating. First, the value on the plateau is *constant* in x_0: it does not interpolate between the two ends of (5), and the source's lower bound $1 + \Delta|x_0|$ is maximally loose there, understating the truth by a factor that tends to $1 + \Delta$ as $x_0 \to 0$. Second, the plateau is a genuine sub-interval and not the whole window. The following shows that the upper end of the bracket is not the answer throughout either.

Proposition 1.3 (Proposition C)**.** *Let $\Delta \geq 0$ and let $0 < s = |x_0| \leq 1$. Against every admissible (a, w) the midpoint law* (4) *keeps*

$$\sup_t |x_t| \;\leq\; 1 + U(s, \Delta), \qquad U(s, \Delta) \;:=\; \max\Big(\Delta s,\; \frac{(\Delta s + 1)^2}{8s},\; \min\big(\tfrac{1}{s}, \Delta\big)\Big),$$

and hence also $\sup_t|x_t| \le 1+\min\bigl(\Delta, U(s,\Delta)\bigr)$. *Moreover, for* $\Delta>0$, $U(s,\Delta)<\Delta$ *— so that this cap is strictly better than the guarantee* $1+\Delta$ *of* [1] *— exactly when*

$$\frac{1}{\Delta} < s < 1 \qquad \textit{and} \qquad (\Delta s)^2 - 6\,\Delta s + 1 < 0, \quad \textit{i.e.} \quad 3-2\sqrt{2} < \Delta s < 3+2\sqrt{2}.$$

The set of s *so described is nonempty exactly for* $\Delta > 1$.

Proposition 1.3 is an upper bound on the value, obtained from one specific controller; it does not compute the value where it applies. What it does establish is qualitative and, we think, the main structural news about the window:

Corollary 1.4. $1+\Delta$ *is not the value throughout the deferred window. At* $\Delta = 3$ *one has* $\gamma^*(1/5,3) = 4 = 1+\Delta$ *by Theorem 1.1, while* $\gamma^*(1/2,3) < 4$ *by Proposition 1.3, which caps the peak of the midpoint law at* 3 *from that start.*

Remark 1.5. Combining this with the source's own endpoint value $\gamma^*(x_0,\Delta) = 1+\Delta$ at $|x_0| = 1$ [1, Corollary 2]: at $\Delta = 3$ the function $s \mapsto \gamma^*(s,3)$ equals 4 at $s = 1/5$, is strictly below 4 at $s = 1/2$, and equals 4 again at $s = 1$. So the value is not monotone on the window, and the window carries at least two regimes — a shape that the two endpoint values alone could not reveal.

1.4. **The upper regime.** The second regime settled here sits at the other end of the window. Let

$$M(R) \;=\; \max_{y\in\mathbb{R}} \bigl|[-R,R]\cap[y-1,y+1]\bigr|\cdot|y| \;=\; \begin{cases} 2R(1-R), & 0 \le R \le \frac{1}{3}, \\ (R+1)^2/4, & \frac{1}{3} \le R \le 3, \\ 2(R-1), & R \ge 3, \end{cases}$$

the *one-step exchange maximum* at prior half-width R: the largest product of a surviving parameter length with the state magnitude bought by it. The quantity is the source's own; [1] computes it — from the constraints of its Lemma 1 — in a footnote to its Proposition 3, at the single half-width $R = \Delta$ that a maximal-exposure history reaches. Put

$$U^\sharp(s,\Delta) \;=\; \max\Bigl(\Delta s,\; \frac{M(\Delta s)}{2s},\; \min\bigl(\tfrac{1}{s},\Delta\bigr)\Bigr),$$

which is the cap U of Proposition 1.3 with its middle branch $(\Delta s+1)^2/(8s)$ replaced by $M(\Delta s)/(2s)$; since $M(R) \le (R+1)^2/4$ for every $R \ge 0$, we have $U^\sharp \le U$ pointwise (Section 7).

Theorem 1.6 (Theorem B)**.** *Let* $\Delta > 0$ *and let* $x_0 \ne 0$ *have* $s = |x_0| \le 1$ *and*

$$s+1 \;\le\; 2\Delta s^2, \qquad \textit{equivalently} \qquad s \;\ge\; \frac{1+\sqrt{1+8\Delta}}{4\Delta}.$$

Then

$$\gamma^*(x_0,\Delta) \;=\; 1+U^\sharp(s,\Delta).$$

The hypothesis forces $\Delta s \ge 1$, *so the third branch of* $U^\sharp$ *is* $1/s$ *there.*

Corollary 1.7. $\gamma^*(1/2,3) = 3$, $\qquad \gamma^*(1/2,10) = 9$, $\qquad \gamma^*(9/10,3) = 37/10$. *The first two lie strictly inside the bracket* (5), *which is* $[5/2,4]$ *and* $[6,11]$ *at those two starts; the third is its lower end* $1+\Delta|x_0|$ *exactly.*

Corollary 1.7 is the qualitative payoff. Before it, the only values known anywhere on the deferred window were the plateau's, all of them at the bracket's upper end; now the window is known to carry all three possibilities — the upper end, the lower end, and points strictly between — and which one occurs is decided by where $|x_0|$ sits, not by Δ alone.

The two theorems cannot conflict, because they never overlap: Theorem 1.6's hypothesis forces $(2\Delta-1)|x_0| \ge 1$, so its region misses the interior of the plateau entirely, and the two hypotheses hold together only at the single point $\Delta = |x_0| = 1$, where both give $\gamma^* = 2$ (Proposition 9.2).

What remains open after this paper is the band between the two regimes, $\min(1, 1/(2\Delta-1)) < s < (1+\sqrt{1+8\Delta})/(4\Delta)$, which is empty for $\Delta \le 1$. Section 6 reports what is computed about it and Section 10 what closing it would take.

1.5. **The mechanism.** The upper half of Theorem 1.1 is the source's own guarantee, re-run from a nonzero start (Theorem 3.3 below). The content is the lower half, and it rests on a two-line inequality whose interest is entirely in its quantifier order.

> *Free move.* If the current state x satisfies $\Delta|x| \le 1$ and the controller plays u, then for *every* target y with $|y-u| \le 1-\Delta|x|$ and *every* admissible parameter a, the disturbance $w := y - ax - u$ is admissible and realises y as the next state.

Because the same successor y is realised for every $a \in [-\Delta, \Delta]$ at once, a prefix of free moves produces the *same observed state history* whatever the parameter is: the controller learns nothing, and the adversary has committed to nothing. This is what makes the lower bound elementary. Attacks in adaptive control usually have to keep a set of parameters alive along the play and argue at the end that some parameter is consistent with everything the controller saw — a lazy-adversary or deferred-commitment reduction. Here no such device is needed: the adversary chooses the parameter at the *leaf* of the recursion, and each free move in the prefix accommodates that choice after the fact, at zero cost.

The attack itself is then a finite descent. With $R = \Delta s$ and $\kappa = |u_0|$ the magnitudes reachable by a free move form the interval $[(\kappa - 1 + R)^+,\ \kappa + 1 - R]$, so either $\kappa \ge R$, and the adversary lands at magnitude $\kappa + 1 - R \ge 1$ and spikes for $1 + \Delta|y| \ge 1 + \Delta$; or $\kappa < R$, and it lands at magnitude 0 or strictly below $2R - 1 = 2\Delta s - 1$ and repeats. The recursion descends precisely when $2\Delta s - 1 < s$, which is the hypothesis $(2\Delta - 1)s < 1$; the closed right endpoint $(2\Delta - 1)s = 1$ needs one extra round run with a strict inductive hypothesis. The base of the recursion is the source's own two-move attack from rest. So the whole lower bound is an explicit adversary: on the open plateau a chain of at most $\lceil s/\delta \rceil + 1$ free moves, where $\delta = \min\bigl(1, 1 - (2\Delta - 1)s\bigr) > 0$, closed by a single identification spike, with the closed right endpoint costing one further round.

Above the plateau the free move is no longer available: the data window of width 2 clips the parameter interval, and the attack must continue against a smaller interval which is in general not centred at the origin. The mechanism of Theorem 1.6 is therefore a *landing* lemma — the adversary forces any successor the window still admits, and inherits everything forced at the interval that survives (Lemma 8.3) — together with the observation that the surviving interval is a legitimate starting position for the same game, because the value depends on the interval only through its length (Lemma 8.1). Three tunings of the landing point then suffice, and the third is the one that makes the theorem an equality: the adversary drives the state back to *rest*, where the whole width-2 window survives, and cashes [1]'s own value from rest at the reduced interval. Which tuning is available depends on how far the control has biased its move, and the hypothesis $s + 1 \le 2\Delta s^2$ is exactly what makes the case split close.

1.6. **What is not claimed.** The value on the band between the two regimes is not determined here, and nothing about it is asserted; see Section 6, where the numerical evidence is reported as such and kept outside every theorem. Outside the region of Theorem 1.6, Proposition 1.3 and its sharpening Theorem 7.2 are guarantees of one controller, not values: where they beat $1 + \Delta$ they say only that the value is smaller, not what the value is. That the plateau of Theorem 1.1 ends exactly at $b^*(\Delta)$ — rather than at some larger point our method fails to reach — is supported by computation only and is stated as Remark 6.2, not as a theorem. Neither is the left endpoint of Theorem 1.6's region claimed to be sharp: the computations of Section 6 put the transition strictly below $(1 + \sqrt{1 + 8\Delta})/(4\Delta)$, and closing that residue is discussed, and not attempted, in Section 10.

1.7. **Related work, and where this claim was looked for.** The nearest technical line is minimax adaptive control in the dynamic-game formulation of Rantzer and coauthors [2, 3, 4, 5, 6, 7]: the same information-state view of an unknown-parameter control problem, with explicit Bellman or Riccati solutions. That line certifies a different currency: where it certifies a closed-loop gain at all, the gain is the induced ℓ_2 (H_∞) one [2, 3, 4], or the ℓ_1 gain for positive systems [7]. An ℓ_2 certificate amortises one large excursion against many small ones, and the peak does not. We found no value statement for the induced ℓ_∞ criterion in that line, and none from a nonzero initial condition below the disturbance level.

The searches behind that sentence, all run on 3 September 2026, were: a single-pass phrase search over a local corpus of 369 shards (86 GB) of arXiv full texts, in which "minimax peak", "sub-disturbance" and "identification spike" each occur in exactly one shard, the one containing [1] itself; extraction and reading of every corpus paper containing "minimax adaptive control" — nine distinct e-prints, in none of whose full texts the word "peak" occurs more than once and none of whose abstracts states a peak value; live arXiv full-text queries, of which `abs:"worst-case peak" AND abs:"adaptive control"` and `all:"consistent model chasing"` each return exactly one paper, namely [1]; and a citing-work lookup for [1] in OpenAlex, which returns HTTP 404 because the three-week-old e-print is not indexed there — *no signal, not a negative*. Every one of these negatives should be read as "not found", never as "new". What makes the question a settled open one rather than folklore is not the searches but the source: [1] states it, in its own words, as deferred.

1.8. **Changes from version 1.** Version 1 of this paper proved Theorem 1.1 and Corollary 1.2, stated Proposition 1.3, and named the upper regime as its natural next step, blocked on one structural point: the return-to-rest attack needs the value from rest at a parameter interval that is neither centred at 0 nor of half-width Δ. Four things have changed.

(1) The upper regime is settled. Theorem 1.6, the three exact values of Corollary 1.7 and the controls of Proposition 9.2 are new here, and so are the tools they rest on: the recentring bijection of Lemma 8.1, which removes the structural obstruction, and the landing lemma and its three tunings (Lemmas 8.3 and 8.4).
(2) The cap of Proposition 1.3 is sharpened to $U^\sharp$ (Theorem 7.2) by restoring the two outer branches of M. Remark 5.1 called that restoration elementary and observed that it changes no statement of version 1; both halves were true then, but it is not cosmetic now — without it Theorem 1.6 would settle no start at which the two caps differ, and in particular not $\gamma^*(1/2, 10) = 9$.
(3) The threshold proved for the upper regime is $(1 + \sqrt{1 + 8\Delta})/(4\Delta)$ — $1/2$ at $\Delta = 3$ and $1/4$ at $\Delta = 10$. Only the return-to-rest branch of the attack constrains it; the spike and the echo survive any bias of the control. The band left open is correspondingly narrow.
(4) The strictness half of Proposition 1.3 is now formally verified (Proposition 5.3). It was the one statement version 1 had to except from the verification claim, and neither the abstract nor Section 11 carries an exception of that kind any longer. One correction to version 1 is recorded, as Remark 6.3.

2. The game, and the encoding of the value

2.1. **Plays and the two halves of the value.** Fix $\Delta > 0$ and $x_0 \in \mathbb{R}$. A *play* is a pair (a, w) with $|a| \le \Delta$ and $w : \mathbb{N} \to \mathbb{R}$, $|w_t| \le 1$ for all t; against a causal law K it generates the trajectory $x(K, a, w)$ by (1) with $u_t = K_t(x_t, \ldots, x_0)$. We keep the two halves of (3) apart and never form the inner supremum.

Definition 2.1. For $c \in \mathbb{R}$:

- K *guarantees* c from x_0 if $|x_t(K, a, w)| \le c$ for every play (a, w) and every $t \in \mathbb{N}$;
- c is *forced* from x_0 if for every causal law K there is a play (a, w) and a *finite* time t with $|x_t(K, a, w)| \ge c$;
- $\gamma^*(x_0, \Delta) = \inf\{c : \text{some causal } K \text{ guarantees } c \text{ from } x_0\}$.

The set of guaranteeable caps is the up-set of $\{\sup_{a,w} \|x(K, a, w)\|_\infty : K \in \mathcal{K}\}$, so its infimum is the $\inf_K \sup_{a,w}$ of (3); and it is bounded below by $|x_0|$, since the play $a = 0$, $w \equiv 0$ and the time $t = 0$ show every guaranteeable cap is at least $|x_0|$. So the infimum is a real number and not a degenerate one. The two halves meet in the obvious way, which is the only tool we use to convert bounds into values: if K guarantees c and c' is forced, then $c' \le c$; if K guarantees c and c is forced, then $\gamma^*(x_0, \Delta) = c$.

Note that "forced" asks for the bound to be *attained* at a finite time, which is stronger than a statement about the supremum, and matches the source's own remark that in this game the infimum and the suprema are all attained.

Remark 2.2 (the formal encoding). In the Lean development a causal law is a function of the state history newest-first, $K : [x_t, x_{t-1}, \ldots, x_0] \mapsto u_t$; the length of the list carries the time index, so causality is structural and there is no side condition to state or to get wrong, and the type of such functions is exactly the class $\mathcal{K}$ of (2). The two halves of Definition 2.1 are two separate predicates. `Guarantee` $K\,x_0\,\Delta\,c$ is a universal statement, $\forall a,\ |a| \le \Delta \to \forall w,\ (\forall t,\ |w_t| \le 1) \to \forall t,\ |x_t| \le c$; `Forced` $x_0\,\Delta\,c$ is $\forall K\,\exists a\,\exists w\,\exists t$ with the same side conditions and $c \le |x_t|$, so the adversary's parameter and disturbance are chosen *after* the law, and the witnessing time is exhibited. The value is the infimum of the guaranteeable caps. No supremum over plays is ever formed, and no measurability, selection or minimax-swap argument is needed anywhere.

2.2. **The consistent set.** For the upper bounds we carry the consistent parameter interval A_t of (4) as the pair

$$(m_t, D_t) \;=\; \big(\mathrm{mid}(A_t),\ \tfrac{1}{2}|A_t|\big), \qquad (m_0, D_0) = (0, \Delta),$$

and update it in the source's own regressor coordinate $\varphi_t(\alpha) = (\alpha - m_t)x_t$, in which the prior's image is $[-R_t, R_t]$ with $R_t = D_t|x_t|$ and, by the source's Corollary 1, the data window's image is $[x_{t+1} - 1, x_{t+1} + 1]$. Writing $\ell = \max(-R_t, x_{t+1} - 1)$ and $h = \min(R_t, x_{t+1} + 1)$, the update on $x_t \neq 0$ is

$$m_{t+1} = m_t + \frac{\ell + h}{2x_t}, \qquad D_{t+1} = \frac{h - \ell}{2|x_t|}, \tag{7}$$

and on $x_t = 0$ nothing changes, which reproduces $A_1 = A_0$ from rest automatically. We call $g_t = 2D_t|x_t|$ the *exposure* at time t, as in [1].

3. The two facts inherited from the source

Both halves of the source's bracket (5) are reproduced in the formal development, because Theorem 1.1 needs both: the upper one is its upper half, and the lower one is where its recursion bottoms out. The mathematics of this section is [1]'s; only the certificates are ours.

Proposition 3.1 (the identification spike; source Proposition 1 (Step 2), Proposition 3(i))**.** *Let $\Delta \ge 0$. For every x_0, the value $1 + \Delta|x_0|$ is forced from x_0, and it is attained at $t = 1$. From rest, $1 + \Delta$ is forced: $\gamma^*(0, \Delta) \ge 1 + \Delta$.*

Proof sketch. For the first claim, let the law play $u = K_0(x_0)$. Choosing the sign of the extreme parameter $a = \pm\Delta$ so that ax_0 and u do not cancel, and then $w_0 = \pm 1$ aligned with the result, gives $|x_1| = |ax_0 + u + w_0| = 1 + \Delta|x_0| + |u| \ge 1 + \Delta|x_0|$. The second claim is the source's two-move attack: from $x_0 = 0$ the disturbance $w_0 = \pm 1$ aligned with u_0 moves the state to $|x_1| = |u_0| + 1 \ge 1$ while leaking nothing about a (at $x_0 = 0$ the transition is parameter-independent), and the spike applied at x_1 then forces $1 + \Delta|x_1| \ge 1 + \Delta$. □

The upper half needs the source's window lemma. For $R \ge 0$ and $y \in \mathbb{R}$ put $I(R, y) = [-R, R] \cap [y - 1, y + 1]$ and $\mu = |I(R, y)|$ when the intersection is nonempty.

Lemma 3.2 (source Lemma 1(i),(ii), the third bound from its proof)**.** *$\mu + |y| \le R + 1, \quad \mu\,|y| \le 2R$, and $\mu \le \min(2, 2R)$.*

Proof sketch. The third is immediate, $I(R, y)$ being the intersection of an interval of length $2R$ with one of length 2. For the first, if $y \ge 0$ then $\mu \le R - (y - 1)$ and $|y| = y$, while if $y < 0$ then $\mu \le (y + 1) + R$ and $|y| = -y$; either way $\mu + |y| \le R + 1$. For the second, the only nontrivial branch is $|y| > \max(1, R)$, where $\mu \le R + 1 - |y|$ by the first part and $2R - (R + 1 - z)z = (z - 1)(z - R) + R > 0$ at $z = |y|$. □

Theorem 3.3 (source Theorem 2 in the form of its Remark 7)**.** *Let $\Delta \ge 0$ and $|x_0| \le 1$. Then the midpoint law $K^{\mathrm{mid}}_{\Delta}$ guarantees $1 + \Delta$ from x_0; in particular $\gamma^*(x_0, \Delta) \le 1 + \Delta$ on the whole window.*

Proof sketch. Four facts about one update (7) are checked directly, using Lemma 3.2 for the last: the true parameter survives, $|a - m_{t+1}| \le D_{t+1}$; the half-width stays nonnegative; $A_{t+1} \subseteq A_t$, i.e. $D_{t+1} \le D_t$; and the exposure drops, $D_{t+1}|x_{t+1}| \le D_t$, which is the source's invariant $g_{t+1} \le |A_t| \le 2\Delta$. Chaining them gives $D_t \le \Delta$ for every t by nestedness from $D_0 = \Delta$, and $D_t|x_t| \le \Delta$ for every t — at $t = 0$ because $|x_0| \le 1$, and for $t \ge 1$ by the invariant. The one-step reach bound $|x_{t+1}| \le D_t|x_t| + 1$, which is the closed loop $x_{t+1} = (a - m_t)x_t + w_t$ together with $|a - m_t| \le D_t$, then gives $|x_{t+1}| \le 1 + \Delta$ for every t, and $|x_0| \le 1 \le 1 + \Delta$ starts it. □

Two points about Theorem 3.3 deserve emphasis, since they are what make the plateau an equality rather than a bound. The law is the source's, unchanged: the single law K_Δ^{mid}, which depends on Δ alone, is used at every start and is not retuned for a nonzero one. And the hypothesis is only $|x_0| \le 1$; the proof never uses $x_0 = 0$, which is why the guarantee extends across the whole deferred window.

4. The free move and the plateau

Lemma 4.1 (free move)**.** *Let $x, u, y \in \mathbb{R}$ and $|a| \le \Delta$. If $|y-u| \le 1 - \Delta|x|$ then $|y - ax - u| \le 1$.*

Proof. $|y - ax - u| \le |y-u| + |a||x| \le (1 - \Delta|x|) + \Delta|x| = 1$. □

The inequality is a triangle inequality; the statement is about the order of its quantifiers. The disturbance $w := y - ax - u$ that realises the target y depends on a, but the realised *state* y does not: for every parameter in $[-\Delta, \Delta]$ the same successor is reachable at admissible cost. Consequently a play in which every move is free produces the same state history for every admissible parameter; a causal law sees the same data whatever a is, and the adversary may postpone its choice of a indefinitely. In the formal development this appears as a transport lemma: if c is forced against the residual law $L \mapsto K(L \mathbin{+\!\!+} [x_0])$ from a free-move successor y, then c is forced against K from x_0 — the parameter supplied by the recursion at y is simply carried back through the prefix, and the disturbance value that does it is admissible by Lemma 4.1. This is the step that would otherwise require a lazy-adversary argument, and it costs two lines.

Lemma 4.2 (landing)**.** *Let $R = \Delta|x| \le 1$ and let the law play u at x, $\kappa = |u|$. Then:*

(1) *if $\kappa \ge R$ there is y with $|y-u| \le 1 - R$ and $|y| \ge 1$;*
(2) *if $\kappa < R$ there is y with $|y-u| \le 1-R$ and either $y = 0$ or $|y| < 2R - 1$.*

Proof. The free move reaches every y in a segment of half-width $1 - R \ge 0$ centred at u, so the reachable magnitudes are exactly $[(\kappa - 1 + R)^+,\ \kappa + 1 - R]$. In case (1) take the far end, $|y| = \kappa + 1 - R \ge R + 1 - R = 1$. In case (2), if $\kappa \le 1 - R$ the origin is reachable; otherwise take the near end, $|y| = \kappa - 1 + R < R - 1 + R = 2R - 1$. □

Theorem 4.3 (Theorem A, lower half)**.** *Let $\Delta > 0$ and $|x_0| \le 1$ with $(2\Delta - 1)|x_0| \le 1$. Then $1 + \Delta$ is forced from x_0: no causal law keeps the peak below $1 + \Delta$.*

Proof sketch. One round. Suppose $\Delta|x_0| \le 1$ and $2\Delta|x_0| \le 1 + \beta$ for some $\beta \ge 0$ such that $1 + \Delta$ is already known to be forced from every state of magnitude at most β. Given a causal law K, apply Lemma 4.2 at x_0 with $u = K_0(x_0)$. In case (1) the adversary lands at $|y| \ge 1$ and then applies Proposition 3.1 at y, forcing $1 + \Delta|y| \ge 1 + \Delta$. In case (2) it lands at $|y| \le \max(0, 2\Delta|x_0| - 1) \le \beta$ and invokes the hypothesis against the residual law. Either way $1 + \Delta$ is forced from x_0.

The descent. Fix $\delta \in (0, 1]$ with $(2\Delta - 1)|x_0| \le 1 - \delta$. Then each round replaces the current magnitude s by at most $\max(0, 2\Delta s - 1) \le \max(0, s - \delta)$, and the hypotheses $s \le 1$, $(2\Delta - 1)s \le 1 - \delta$ are preserved. So by induction on n: if $|x_0| \le n\delta$ then $1 + \Delta$ is forced from x_0, the base case $n = 0$ being $x_0 = 0$, which is the second half of Proposition 3.1. For $|x_0|$ in the strict interior, $(2\Delta - 1)|x_0| < 1$, take $\delta = \min\bigl(1, 1 - (2\Delta - 1)|x_0|\bigr)$ and n Archimedean with $|x_0| \le n\delta$. At the closed right endpoint $(2\Delta - 1)|x_0| = 1$ the descent is circular, since there $2\Delta|x_0| - 1 = |x_0|$ exactly; one further round run with the strict hypothesis "$1 + \Delta$ is forced from every state of magnitude *strictly less than* $|x_0|$" closes it, and that hypothesis is the interior case already proved. □

The adversary produced by this proof is explicit and finite: at most $\lceil |x_0|/\delta \rceil + 2$ free moves, each of which is one disturbance value, closed by the single move of the spike. It uses one parameter value, chosen at the end.

Proof of Theorem 1.1 and Corollary 1.2. Theorem 3.3 gives a law guaranteeing $1 + \Delta$ from x_0, since $|x_0| \le 1$; Theorem 4.3 forces $1 + \Delta$; so the value is $1 + \Delta$. For the corollary, $\Delta \le 1$ makes $(2\Delta - 1)|x_0| \le |x_0| \le 1$ automatic — for $2\Delta \le 1$ the left side is nonpositive — so every $|x_0| \le 1$ lies in the plateau. □

Remark 4.4. Two instances, both verified formally. At $\Delta = 3$ the plateau's right endpoint is $b^*(3) = 1/5$ and $\gamma^*(1/5, 3) = 4$, where the source's bracket (5) is [1.6, 4]. At $\Delta = 1$ the plateau's right endpoint coincides with the window's and $\gamma^*(1, 1) = 2$.

5. The plateau ends

Proof sketch for Proposition 1.3. Write $s = |x_0| > 0$ and run the midpoint law from x_0. The four bounds are:

- $t = 0$: $|x_0| = s \le 1 \le 1 + U$.
- $t = 1$: the reach bound gives $|x_1| \le D_0 s + 1 = \Delta s + 1$, the first branch.
- $t = 2$: $|x_2| \le D_1|x_1| + 1$. Let $\mu = 2sD_1$ be the length of the surviving interval in the regressor coordinate at time 0. Lemma 3.2, applied with $R = D_0 s = \Delta s$ and $y = x_1$, gives $\mu + |x_1| \le \Delta s + 1$, whence by the arithmetic–geometric mean inequality $\mu|x_1| \le (\Delta s + 1)^2/4$ and $D_1|x_1| = \mu|x_1|/(2s) \le (\Delta s + 1)^2/(8s)$, the second branch.
- $t \ge 3$: $|x_t| \le D_{t-1}|x_{t-1}| + 1 \le D_{t-2} + 1 \le D_1 + 1$ by the exposure invariant and nestedness of Theorem 3.3's proof, and $\mu \le \min(2, 2\Delta s)$ from Lemma 3.2 gives $D_1 = \mu/(2s) \le \min(1/s, \Delta)$, the third branch.

The second assertion, $\sup_t|x_t| \le 1 + \min(\Delta, U(s, \Delta))$, is this together with Theorem 3.3. For the comparison, $U(s, \Delta) < \Delta$ holds if and only if each of the three branches is smaller, and for $\Delta > 0$ and $0 < s \le 1$: $\Delta s < \Delta$ is $s < 1$; $\min(1/s, \Delta) < \Delta$ is $s > 1/\Delta$; and $(\Delta s+1)^2/(8s) < \Delta$ is $(\Delta s+1)^2 < 8\Delta s$, i.e. $R^2 - 6R + 1 < 0$ at $R = \Delta s$, i.e. $3 - 2\sqrt{2} < R < 3 + 2\sqrt{2}$. All three clauses are needed: $s < 1$ fails at $s = 1$, where the first branch already equals Δ, and $s > 1/\Delta$ fails at every small s. Note that $s > 1/\Delta$ forces $R > 1 > 3 - 2\sqrt{2}$, so only the upper root of the quadratic binds; and since $1/\Delta < s < 1$ requires $\Delta > 1$, while for $\Delta > 1$ every s slightly above $1/\Delta$ satisfies all three, the region is nonempty exactly for $\Delta > 1$. □

Remark 5.1. The middle branch is the source's own one-step exchange maximum

$$M(R) \;=\; \max_y \, |[-R, R] \cap [y - 1, y + 1]| \cdot |y|,$$

restricted to its middle range $\frac{1}{3} \le R \le 3$, where $M(R) = (R+1)^2/4$; [1] computes the same three-branch maximum, from the constraints of its Lemma 1, in the footnote to its Proposition 3, at $R = \Delta$. Since $(R+1)^2/4 - 2R(1-R) = (3R-1)^2/4$ and $(R+1)^2/4 - 2(R-1) = (R-3)^2/4$, the branch $(R+1)^2/4$ dominates M everywhere, so the cap above is valid for all Δs and is only slightly weaker than the three-branch form at very small and very large Δs. Restoring the two outer branches is elementary piecewise real arithmetic and changes no statement of this paper; we have not done it.

Proposition 5.2 (controls)**.** *Let $\Delta > 0$.*

(1) (The bracket's lower end is not the value.) *If $x_0 \ne 0$, $|x_0| < 1$ and $(2\Delta - 1)|x_0| \le 1$, then $0 < |x_0|$ and $1 + \Delta|x_0| < \gamma^*(x_0, \Delta)$.*
(2) (The bracket's upper end is not the value.) $\gamma^*(1/2, 3) < 1 + 3$.
(3) (Non-vacuity.) *For every $\Delta > 0$ the plateau contains a nonzero start: $x_0 = \min\big(1, 1/(2\Delta)\big)$ satisfies $x_0 \ne 0$, $|x_0| \le 1$ and $(2\Delta - 1)|x_0| \le 1$.*
(4) (Nothing below is guaranteeable.) *On the plateau no causal law guarantees any $c < 1 + \Delta$.*
(5) (Nothing above is forced.) *For $|x_0| \le 1$ no $c > 1 + \Delta$ is forced.*

Proof. (1) is Theorem 1.1 together with $\Delta|x_0| < \Delta$. (2) is Proposition 1.3 at $s = 1/2$, $\Delta = 3$, where $U = \max(3/2,\, 25/16,\, \min(2, 3)) = 2$, so the midpoint law guarantees 3. (3) is arithmetic: $(2\Delta - 1)/(2\Delta) < 1$. (4) and (5) are Theorem 4.3 and Theorem 3.3 against the definition of guarantee. □

Items (1) and (2) are the two ends of the source's bracket, and they fall on opposite sides: the lower end is refuted as the value everywhere on the plateau, the upper end is refuted as the value at one explicit point of the window. Items (4) and (5) are the guard that the two predicates of Definition 2.1 are discriminating in both directions — a value statement in a game is only as good as the failure of the two one-sided statements that bracket it.

The strictness half of Proposition 1.3 was the one statement version 1 of this paper could not present as verified. It is verified now, in the following form, which also records two elementary controls on the region it describes.

Proposition 5.3. *Let $\Delta > 0$ and $s > 0$. Then*

$$U(s,\Delta) < \Delta \iff \big(\tfrac{1}{\Delta} < s < 1 \text{ and } (\Delta s)^2 - 6\Delta s + 1 < 0\big) \iff \big(\tfrac{1}{\Delta} < s < 1 \text{ and } \Delta s < 3 + 2\sqrt{2}\big),$$

the second equivalence because $s > 1/\Delta$ already forces $\Delta s > 1 > 3 - 2\sqrt{2}$, so that only the upper root of the quadratic binds. On that region $\gamma^*(x_0, \Delta) < 1 + \Delta$ for every x_0 with $|x_0| = s \le 1$. The region is nonempty for every $\Delta > 1$ and empty for every $\Delta \le 1$; and its clause $s < 1$ is not implied by the other two, since at $\Delta = 3$, $s = 1$ those two hold while $U(1, 3) = 3 = \Delta$.

Proof sketch. The equivalences are the branch-by-branch algebra already displayed in the proof of Proposition 1.3. The value statement is Proposition 1.3 itself, which exhibits a law guaranteeing $1 + U(s,\Delta) < 1 + \Delta$. Nonemptiness for $\Delta > 1$ is witnessed by $s = (\Delta+1)/(2\Delta)$ when $1 < \Delta \le 2$ and by $s = 2/\Delta$ when $\Delta > 2$. Emptiness for $\Delta \le 1$ holds for two reasons at once: $1/\Delta \ge 1$ makes $1/\Delta < s < 1$ vacuous, and directly $\min(1/s, \Delta) = \Delta$ for every $s \le 1$, so $U(s,\Delta) \ge \Delta$ there. The second reason is the algebraic shadow of Corollary 1.2: when $\Delta \le 1$ the whole window carries the value $1 + \Delta$, so no sharpening of that guarantee can exist. The last clause is the displayed evaluation at $(\Delta, s) = (3, 1)$. □

6. The numerical landscape of the window

Everything in this section is computation, obtained outside the formal development, and no statement of this paper depends on any of it. It is reported because it is what located the two regimes that are settled here, and because it is the only information anyone has about the band between them.

6.1. **The band.** Between the plateau and the regime of Theorem 1.6 there is an interval on which we assert nothing. The following is computational evidence, obtained outside the formal development, and is recorded as such.

Remark 6.1 (outside the formal development). A two-sided value iteration for the Bellman operator $\Phi(s, D) = \inf_c \sup_y \max\big(|y|, \Phi(|y|, \mu/(2s))\big)$ on a grid in the two-dimensional information state (s, D), run upward from $1 + Ds$ and downward from $1 + D$ — both monotone, and meeting to four decimals — with the exact boundary data $\Phi(0, D) = 1+D$ and $\Phi(\sigma, D) = 1+D\sigma$ for $\sigma \ge 1$, gives the following profile. At $\Delta = 3$: 4.000 at $s = 0.20$, then 3.865, 3.713, 3.606, 3.493, 3.406 at $s = 0.24, 0.28, 0.32, 0.36, 0.40$, and from $s \approx 0.45$ upward the values agree with $1 + U(s,\Delta)$. At $\Delta = 10$: 11.000 at $s = 0.05$, then 9.703, 9.080, 8.545 at $s = 0.075, 0.100, 0.125$, and from $s \approx 0.15$ upward agreement with $1+U(s,\Delta)$. So the computed transition point above which $1 + U$ is the value is $s_2 \approx 0.575, 0.45, 0.15$ at $\Delta = 2, 3, 10$; on $b^*(\Delta) < s < s_2(\Delta)$ the computed value lies strictly below $1 + \Delta$, and below $1 + U(s,\Delta)$ except immediately under s_2, where the two already agree to the grid's fourth decimal. A diagnostic at the converged fixed point ($\Delta = 3$) reports that inside the band the optimal implicit selector is biased *away* from the midpoint (a bias of roughly 0.55–0.63), whereas it is the midpoint on the plateau and above the band. None of this is certified: the one-step lower-bound column of that computation is an infimum over a finite grid of controller actions and is therefore not a valid lower bound, and no statement of this paper depends on any of it.

Remark 6.2 (outside the formal development). The same computation puts the right edge of the plateau at $1/(2\Delta - 1)$ exactly — the computed value drops below $1 + \Delta$ immediately above 0.3333, 0.2 and 0.05263 at $\Delta = 2, 3, 10$ — so Theorem 1.1 appears to be tight and not merely to cover a sub-interval. An independent adversary search against the midpoint law, over 60 pairs (Δ, s) with $\Delta \in \{0.5, 1, 2, 3, 10, 30\}$ and $s \in \{0.05, \ldots, 1.0\}$ — exhaustive over four parameter values times all 3^5 disturbance prefixes in $\{+1, -1, 0\}^5$, plus 150,000 random plays, plus the myopic greedy adversary from both parameter endpoints, all at horizon 10 — never exceeds the cap of Proposition 1.3 and attains it to four decimals in 24 of the 60 cells, among them all of $\Delta = 3$ with $s \in \{0.4, 0.5, 0.6, 0.75, 0.9\}$ and $\Delta = 10$ with $s \in \{0.2, 0.3\}$. The "never exceeds" half of that sentence is Proposition 1.3 and is proved; the attainment half is a finite search and certifies only what it explored.

Remark 6.3 (a correction to version 1). Remark 6.1 compares the computed profile with $1+U(s,\Delta)$, the cap of Proposition 1.3. The comparison function in the computation itself is the sharpened $1+U^\sharp(s,\Delta)$, and the two can differ wherever the middle branch of M is not the active one, i.e. wherever Δs lies outside $[1/3,3]$. At $\Delta=3$ and $s\geq 1/9$ the middle branch is active and the two agree, so the $\Delta=3$ profile is unaffected; at $\Delta=10$ they part company above $s=3/10$, and there the computed profile follows $1+U^\sharp$ and not $1+U$. The clearest case is $\Delta=10$, $s=1/2$: the computed value is 9.000, while $1+U^\sharp=9$ and $1+U=10$. Theorem 1.6 confirms the computation, $\gamma^*(1/2,10)=9$. Version 1 of this paper read the agreement at $\Delta=10$ from $s\approx 0.15$ upward as agreement with $1+U$; it is agreement with $1+U^\sharp$, and Remark 6.1 should be read that way.

6.2. **Above the band.** The computed profile agrees with $1+U^\sharp(s,\Delta)$ for every s above $s_2(\Delta)\approx 0.575,\ 0.45,\ 0.15$ at $\Delta=2,3,10$, whereas the threshold at which Theorem 1.6 takes over is $(1+\sqrt{1+8\Delta})/(4\Delta)\approx 0.640,\ 0.5,\ 0.25$ there. So the theorem's left endpoint is not claimed to be sharp, and the residue between $s_2(\Delta)$ and that threshold is part of what remains; see Section 10.

7. The exchange maximum and the sharpened cap

[1] needs the quantity M once, and in one place only: in a footnote in its Section 5.2, attached to the paragraph that follows its Proposition 3, to show that a maximal-exposure history is never followed by another one. We quote the computation.

> "Maximizing $\mu_{t+1}|x_{t+1}|$ under these constraints gives $2(\Delta-1)$ for $\Delta\geq 3$, $(\Delta+1)^2/4$ for $1/3\leq\Delta\leq 3$, and $2\Delta(1-\Delta)$ for $\Delta\leq 1/3$; in every case strictly below 2Δ: a maximal-exposure history is never followed by another one."

The constraints are those of Lemma 3.2, in the closed-loop form the source's Corollary 1 gives them — which is what the footnote invokes — and the half-width at which [1] writes them is Δ itself, because at maximal exposure its $g_t/2$ equals Δ. A start with $|x_0|<1$ never reaches maximal exposure, so what the sub-disturbance window needs is the same maximisation at a general half-width. That is the one extension of the source's own text made here.

Lemma 7.1. *Let $R\geq 0$ and let $\mu,z\geq 0$ satisfy $\mu\leq\min(2,2R)$ and $\mu+z\leq R+1$. Then $\mu z\leq M(R)$.*

Proof sketch. Three branches, each with one case split on z. For $R\leq 1/3$: if $z\leq 1-R$ use $\mu\leq 2R$ and $2Rz\leq 2R(1-R)$; otherwise $\mu\leq R+1-z$ and $2R(1-R)-(R+1-z)z=(z-2R)\big(z-(1-R)\big)\geq 0$, both factors being nonnegative because $1-R\geq 2R$. For $1/3\leq R\leq 3$ it is the arithmetic–geometric mean inequality, $4\mu z\leq(\mu+z)^2\leq(R+1)^2$. For $R\geq 3$: if $z\leq R-1$ use $\mu\leq 2$; otherwise $\mu\leq R+1-z$ and $2(R-1)-(R+1-z)z=(z-2)\big(z-(R-1)\big)\geq 0$, because $R-1\geq 2$. □

The bound is attained — the displayed M really is the maximum and not merely an upper bound for it — at $z=1-R$, $z=(R+1)/2$ and $z=R-1$ on the three branches, where the surviving length is $2R$, $(R+1)/2$ and 2 respectively. We use the maximiser in Section 9; here only the inequality is needed.

Theorem 7.2. *Let $\Delta\geq 0$ and $0<s=|x_0|\leq 1$. Against every admissible (a,w) the midpoint law (4) keeps*

$$\sup_t |x_t| \;\leq\; 1+U^\sharp(s,\Delta),$$

and $U^\sharp(s,\Delta)\leq U(s,\Delta)$, so this is at least as strong as Proposition 1.3.

Proof sketch. The skeleton is that of Proposition 1.3, with Lemma 7.1 in place of the arithmetic–geometric mean step at $t=2$. There, with $R=D_0s=\Delta s$, $\mu=2sD_1$ the length of the surviving interval in the regressor coordinate at time 0, and $z=|x_1|$, Lemma 3.2 supplies exactly the two hypotheses $\mu\leq\min(2,2R)$ and $\mu+z\leq R+1$; hence $D_1|x_1|=\mu z/(2s)\leq M(\Delta s)/(2s)$, the middle branch. The bounds at $t=0$, $t=1$ and $t\geq 3$ are those of Proposition 1.3 verbatim. For the comparison, $(R+1)^2/4-2R(1-R)=(3R-1)^2/4$ and $(R+1)^2/4-2(R-1)=(R-3)^2/4$ are both nonnegative, so $M(R)\leq(R+1)^2/4$ for every $R\geq 0$ and the middle branch of $U^\sharp$ never exceeds that of U; the other two branches are shared. □

Remark 7.3. The inequality $U^\sharp \le U$ can be strict, and where it is strict it can decide a start. At $(\Delta, s) = (10, 1/2)$ one has $U^\sharp = 8$ against $U = 9$, so Proposition 1.3 caps the peak of the midpoint law at 10 and Theorem 7.2 caps it at 9 — and 9 is the value, by Theorem 1.6. No general strictness is claimed: the two caps also coincide at many starts with Δs outside $[1/3, 3]$, namely wherever a branch other than the middle one attains both maxima, as at $\Delta = 5$, $s = 1$. Nor is it the sharpening that decides every start with $\Delta s > 3$: where the two caps coincide, the cruder one already closes the cell. What is true is the statement Remark 5.1 of version 1 did not anticipate: restoring the outer branches is elementary, and it is nonetheless load-bearing, because without it Theorem 1.6 settles no start at which the two caps differ — among them $(\Delta, s) = (10, 1/2)$, the cell of Corollary 1.7.

8. The prior as a parameter, and the landing lemma

Theorem 1.6's lower half needs one thing that Sections 2–4 do not provide: a move that hands the game on when only *part* of the parameter interval survives, and a way to treat the surviving part as a starting position of the same game. Here are the two.

8.1. **Recentring, and the two-dimensional information state.** For $m \in \mathbb{R}$ and $D \ge 0$ let $\gamma^*_{m,D}(x_0)$ denote the value (3) of the same game with the parameter constraint $|a| \le \Delta$ replaced by $|a - m| \le D$, and read the two predicates of Definition 2.1 with the same replacement. [1] records the reduction to $m = 0$ as a normalisation, in its problem formulation:

> "the causal substitution $\tilde{u}_t = u_t + cx_t$ recenters any prior interval at the origin, so the symmetry of $[-\Delta, \Delta]$ is a convention rather than an assumption (Remark 9(c) records the resulting value for an asymmetric prior)".

Lemma 8.1 (recentring and reflection)**.** *Let $m \in \mathbb{R}$, $D \ge 0$ and $x_0, c \in \mathbb{R}$.*

(1) *Write $\rho_m(K)$ for the law $L \mapsto m\,L_0 + K(L)$, where L_0 is the current state at the head of the history L. Then ρ_m is a bijection of $\mathcal{K}$, with inverse ρ_{-m}; a law K guarantees c from x_0 against the interval $[m - D, m + D]$ if and only if $\rho_m(K)$ guarantees c from x_0 against $[-D, D]$; the same equivalence holds for "forced". Consequently $\gamma^*_{m,D}(x_0) = \gamma^*(x_0, D)$.*

(2) $\gamma^*(-x_0, \Delta) = \gamma^*(x_0, \Delta)$.

So the value depends on the initial state and the initial parameter interval A_0 only through the pair $\big(|x_0|,\, |A_0|/2\big)$.

Proof sketch. (1) One induction on t shows that the trajectory of $\rho_m(K)$ under the parameter b is the trajectory of K under $m + b$, from the same x_0 and the same disturbance; the substitution $a = m + b$ matches the two parameter constraints, and both predicates transport because ρ_m is onto — the forcing predicate quantifies over all laws, so surjectivity is what is needed there. (2) is the reflection $K \mapsto \big(L \mapsto -K(-L)\big)$, under which the trajectory from $-x_0$ with disturbance $-w$ is the negative of the trajectory from x_0 with w. □

Corollary 8.2. *From rest, and for a parameter interval $[\underline{a}, \overline{a}]$ of diameter $\bar{D} = \overline{a} - \underline{a} > 0$, the value is $1 + \bar{D}/2$, whatever the centre.*

Proof. Lemma 8.1(1) with $D = \bar{D}/2$, and then Theorem 1.1 at $x_0 = 0$. □

Corollary 8.2 is Remark 9(c) of [1] — the remark the sentence quoted just above points to for "the resulting value for an asymmetric prior" — and the mathematics of both is the source's; the reason to isolate them is that the third attack of Theorem 1.6 lands the state at rest with a parameter interval that is neither symmetric nor of half-width Δ, and it is exactly here that it is cashed.

8.2. **Landing.** Fix a state $x \neq 0$, let the law play $u = K(x)$, and put $R = \Delta|x|$. In the coordinate $\varphi(\alpha) = \alpha x + u$ — the successor predicted by the parameter α — the image of the parameter interval is the segment of centre u and half-width R, and by Corollary 1 of [1] the image of the data window at a realised successor y is $[y - 1, y + 1]$. Write

$$I(y) \;=\; [u{-}R,\, u{+}R] \cap [y{-}1,\, y{+}1], \qquad \ell(y) \;=\; \min(u{+}R,\, y{+}1) - \max(u{-}R,\, y{-}1), \qquad D'(y) = \frac{\ell(y)}{2|x|}.$$

Thus $\ell(y)$ is the length of $I(y)$ when that intersection is nonempty and is negative exactly when it is empty, and $D'(y)$ is the half-width, in parameter units, of the interval that survives the transition to y.

Lemma 8.3 (landing)**.** *Let $x \neq 0$ and let y satisfy $\ell(y) \geq 0$. If c is forced from y against a parameter interval of half-width $D'(y)$, then c is forced from x against K.*

Proof sketch. Pull the continuation attack back through Lemma 8.1(1), at the centre $m' = \big(\operatorname{mid} I(y) - u\big)/x$ of the survivor measured in parameter units. A parameter a with $|a - m'| \leq D'(y)$ has $ax + u \in I(y)$; lying in the first factor gives $|a| \leq \Delta$, and lying in the second gives $|y - ax - u| \leq 1$, so the single disturbance $w := y - ax - u$ is admissible and realises y as the next state. Splicing that one move in front of the continuation play gives the claim. The free move of Lemma 4.1 is the special case $[u - R, u + R] \subseteq [y - 1, y + 1]$, in which nothing is clipped and $D'(y) = \Delta$. □

Three tunings of the landing point are what Theorem 1.6 uses. The first is the identification spike with the control's own move kept rather than discarded; the second is built from the two primitives [1] names in its Definition 4, a "Spike of height y" (its part (b)) followed by an "Echo" (its part (c)); the third is the one version 1 of this paper could not reach.

Lemma 8.4 (three attacks)**.** *Let $\Delta \geq 0$, let $x \neq 0$, write $u = K(x)$ and $R = \Delta|x|$, and put $\mu(R, z) = \min(R, z + 1) + \min(R, 1 - z)$. Then, against K from x:*

(1) (sharp spike) $1 + \Delta|x| + |u|$ *is forced;*
(2) (echo) *for every $z \geq 0$ with $\mu(R, z) \geq 0$, the value $1 + \mu(R, z)\, z/(2|x|)$ is forced;*
(3) (return to rest) *if $|u| + 1 \leq R$, then $1 + 1/|x|$ is forced.*

Proof sketch. (1) is the identification spike of Proposition 3.1 with the term $|u|$ kept rather than discarded — Proposition 3(i) of [1] in full. (2) Land at $y = u \pm z$, the sign chosen to be that of u, and then spike at the survivor. Two facts make the bound independent of what the control played: the landing magnitude is $|u| + z \geq z$, and the surviving length is $\mu(R, z)$ for either sign, because the image of the parameter interval is symmetric about u, so translating it does not change what a window of width 2 can cut off. Lemma 8.3 followed by Proposition 3.1 at the survivor forces $1 + D'(y)|y|$, and $D'(y)|y| = \mu(R, z)(|u| + z)/(2|x|) \geq \mu(R, z)z/(2|x|)$, so the displayed value is forced too. (3) The hypothesis $|u| + 1 \leq R$ gives $[-1, 1] \subseteq [u - R, u + R]$, so landing at $y = 0$ leaves $\ell(0) = 2$ exactly and $D'(0) = 1/|x|$: the whole width-2 window survives, and nothing about the parameter has been learnt beyond it. Cashing the second half of Proposition 3.1 — the two-move attack from rest — at the origin with that interval, through Lemma 8.1(1), forces $1 + 1/|x|$. □

9. Theorem B

Theorem 9.1 (Theorem B, lower half)**.** *Let $\Delta > 0$ and let $x_0 \neq 0$ satisfy $s = |x_0| \leq 1$ and $s + 1 \leq 2\Delta s^2$. Then $1 + U^\sharp(s, \Delta)$ is forced from x_0: no causal law keeps the peak below it.*

Proof sketch. Write $R = \Delta s$, fix a causal law K and let $u = K(x_0)$. The hypothesis reads $2R \geq 1 + 1/s$, whence $R \geq 1$ — since $s \leq 1$ makes $1/s \geq 1$ — and therefore $\min(1/s, \Delta) = 1/s$, so the target is $1 + \max\big(R, M(R)/(2s), 1/s\big)$. Forcing is closed under maxima: whichever of two forced values the maximum *is*, it is forced. So it is enough to force each of the three numbers separately.

$1 + R$ is the identification spike, Proposition 3.1.

$1 + M(R)/(2s)$ is the echo of Lemma 8.4(2) at a landing distance where the product $\mu(R, z)\, z$ is exactly $M(R)$: take $z = (R+1)/2$ for $R \leq 3$ and $z = R - 1$ for $R \geq 3$, at which $\mu(R, z)$ equals $(R+1)/2$ and 2 respectively. (The third maximiser, $z = 1 - R$ for $R \leq 1/3$, cannot occur here, as $R \geq 1$.) This bound is available whatever the control played, which is the point of the second half of Lemma 8.4(2).

$1 + 1/s$ is a case split on the control's own move. If $|u| + 1 \leq R$ the return to rest, Lemma 8.4(3), forces $1 + 1/s$ outright. Otherwise $|u| > R - 1$, and the sharp spike of Lemma 8.4(1) forces $1 + R + |u| > 2R \geq 1 + 1/s$. Either the control stays close enough to the midpoint for the width-2 window at the origin to swallow the whole parameter image, or it has already paid for the deviation. This split is the only place the hypothesis $s + 1 \leq 2\Delta s^2$ is used, and it is used exactly in the form $2R \geq 1 + 1/s$. □

Proof of Theorem 1.6. Theorem 7.2 exhibits a law guaranteeing $1 + U^\sharp(s, \Delta)$ from x_0, since $0 < s \le 1$; Theorem 9.1 forces the same number; so it is the value. $\square$

Proof of Corollary 1.7. Three evaluations, each with the hypothesis checked first. At $(\Delta, s) = (3, 1/2)$: $s + 1 = 3/2 = 2\Delta s^2$, so the hypothesis holds with equality; $R = 3/2$ lies in the middle branch, $M(R) = 25/16$, $M(R)/(2s) = 25/16$, $\Delta s = 3/2$ and $\min(1/s, \Delta) = 2$, so $U^\sharp = 2$ and $\gamma^* = 3$. At $(10, 1/2)$: $s + 1 = 3/2 \le 5 = 2\Delta s^2$; $R = 5$ lies in the outer branch, $M(R) = 8$, $M(R)/(2s) = 8$, $\Delta s = 5$ and $\min(1/s, \Delta) = 2$, so $U^\sharp = 8$ and $\gamma^* = 9$. At $(3, 9/10)$: $s + 1 = 19/10 \le 243/50$; $R = 27/10$, $M(R)/(2s) = 1369/720$, $\Delta s = 27/10$ and $\min(1/s, \Delta) = 10/9$, so $U^\sharp = 27/10$ and $\gamma^* = 37/10$, which is $1 + \Delta|x_0|$. The bracket (5) at these three starts is $[5/2, 4]$, $[6, 11]$ and $[37/10, 4]$. $\square$

Proposition 9.2 (controls)**.** *Let $\Delta > 0$ and write H for the hypothesis of Theorem 1.6: $x_0 \ne 0$, $|x_0| \le 1$ and $|x_0| + 1 \le 2\Delta|x_0|^2$.*

1. (The bracket's upper end is not the value.) *Under H with $|x_0| < 1$, $\gamma^*(x_0, \Delta) < 1 + \Delta$.*
2. (Nothing is claimed for small Δ.) *If $\Delta < 1$ then $2\Delta s^2 < s + 1$ for every $0 < s \le 1$, so H holds nowhere on the window.*
3. (Non-vacuity.) *For every $\Delta > 1$ the start $x_0 = (\Delta + 1)/(2\Delta)$ satisfies H and lies strictly inside the window, $|x_0| < 1$.*
4. (Disjointness, and agreement where the two theorems meet.) *H forces $(2\Delta - 1)|x_0| \ge 1$, so Theorem 1.6's region never meets the interior of the plateau of Theorem 1.1; the two hypotheses hold together only at $\Delta = |x_0| = 1$, where $U^\sharp(1, 1) = 1$ and both give $\gamma^* = 2$.*
5. (Nothing below is guaranteeable, nothing above is forced.) *Under H no causal law guarantees any $c < 1 + U^\sharp(|x_0|, \Delta)$; and for every $x_0 \ne 0$ with $|x_0| \le 1$, no $c > 1 + U^\sharp(|x_0|, \Delta)$ is forced.*

Proof sketch. (1) Under H with $s = |x_0| < 1$ one has $\Delta s > 1$, while $s < 1$ gives $\Delta s < \Delta$; the middle branch obeys $M(R)/(2s) < 2R/(2s) = \Delta$ because $M(R) < 2R$ for $R > 0$ — which is the last clause of the source's own footnote; and $\min(1/s, \Delta) < \Delta$ because $1/s < \Delta$. So $U^\sharp < \Delta$ and Theorem 1.6 gives the claim. (2) is $2\Delta s^2 < 2s^2 \le s + 1$ for $0 < s \le 1$. (3) is arithmetic: at $s = (\Delta + 1)/(2\Delta)$, $2\Delta s^2 - (s + 1) = (\Delta^2 - \Delta)/(2\Delta) \ge 0$, and $s < 1$ because $\Delta > 1$. (4) is the identity $(2\Delta - 1)s - 1 = \big(2\Delta s^2 - (s + 1) + (1 - s^2)\big)/s$, whose numerator is nonnegative under H; the plateau requires $(2\Delta - 1)s \le 1$, so both hold only in the equality case, which forces $s = 1$ and then $\Delta = 1$. (5) is Theorem 9.1 and Theorem 7.2 against Definition 2.1. $\square$

10. What remains

10.1. **The band.** The whole remaining question about the deferred window is the band

$$\min\Big(1, \frac{1}{2\Delta - 1}\Big) \;<\; |x_0| \;<\; \frac{1 + \sqrt{1 + 8\Delta}}{4\Delta},$$

empty for $\Delta \le 1$, and $(0.333\ldots, 0.640\ldots)$, $(0.2, 0.5)$, $(0.0526\ldots, 0.25)$ at $\Delta = 2, 3, 10$; its right endpoint is $(1 + \sqrt{17})/8 = 0.6403\ldots$, $1/2$ and $1/4$ at those three. Two different gaps make it up. Below $s_2(\Delta)$ what is missing is a lower bound, and the obstruction is structural rather than a matter of computing longer: the quantity to certify is an infimum over a *continuum* of controller actions, which no finite grid delivers, and the diagnostic reported in Remark 6.1 says that the optimal action there is not the midpoint, so there is no closed-form selector to substitute for the infimum. Between $s_2(\Delta)$ and the threshold of Theorem 1.6 the upper bound is already the computed value and only the lower bound is short by the width of one case split: it is the return-to-rest branch of Theorem 9.1 that fixes the threshold, and improving it means relaxing over two steps rather than one.

10.2. **Higher dimensions.** The vector case is the source's own stated frontier and nothing in this paper bears on it.

11. Verification

Every lemma, proposition, theorem and corollary stated above has been formally verified in Lean 4 against *Mathlib* [8, 9]. The development is eight modules. Four — the game and the free move, the lower bound, the controller and its guarantees, and the assembly — carry 926 lines and 81 declarations, and cover everything in Sections 2–5 except Proposition 5.3. Four more — the prior as a parameter with the landing lemma, the exchange maximum and the sharpened cap, the strictness comparison, and Theorem B — carry 1025 lines and 81 declarations, and cover Proposition 5.3 together with Sections 7–9. The development contains no `sorry` and declares no axiom by hand. A print of the axiom set returns `propext`, `Classical.choice` and `Quot.sound` and nothing else, for *every* constant the eight modules put into the compiled environment: 112 from the first four and 116 from the four later ones, in each case their 81 declarations together with the equation, match and simp lemmas the compiler generates alongside them. (Version 1 of this paper reported the 88 constants of the first four that are not the compiler's own internal auxiliaries; the sweep here is over all of them.) In particular there is no `sorryAx`, and no compiled-evaluation axiom, since there is no finite object anywhere in this problem to decide. Every statement is real arithmetic; the development uses no decision procedure over a finite search space, no external solver and no certificate. The first four modules compile in about 56 seconds and the four later ones in about ten minutes between them, each at or below the cost of importing *Mathlib* itself. Two encoding points are worth recording because they are where a formalisation of a game can go wrong. Causality is structural, not asserted: a law is a function of the state history and the history's length is the time index, so there is no causality hypothesis that could be stated too weakly. And the adversary's quantifiers sit in the order the game demands — the parameter and the disturbance are chosen after the law, and the forcing statement exhibits the finite time at which the bound is met — so no supremum, and no exchange of a supremum with an infimum, occurs anywhere in the proofs.

Three things quoted above are deliberately outside the formal development and are labelled where they appear: the numerics of Remarks 6.1, 6.2 and 6.3, and the comparison of $s_2(\Delta)$ with the threshold of Theorem 1.6 in Section 6; the source's value $\gamma^*(x_0, \Delta) = \max(|x_0|, 1 + \Delta|x_0|)$ for $|x_0| \geq 1$, which is quoted from [1] and used only in Remark 1.5 and as boundary data in Remark 6.1; and the closed form of Theorem 1.6's hypothesis — the verified statement carries the inequality $|x_0| + 1 \leq 2\Delta|x_0|^2$ itself, and its restatement as $|x_0| \geq (1 + \sqrt{1 + 8\Delta})/(4\Delta)$, the positive root of that quadratic, is elementary algebra done on the page. Version 1 of this paper had to record one more: the strictness comparison in Proposition 1.3, which is now Proposition 5.3 and is verified with the rest.

The source of the development is currently held in a private repository: repository reference to be added at submission.

On the reference list. Each entry records how far the reading went. Every theorem, proposition, lemma, corollary, definition, remark and section number attributed to [1] above is the printed numbering of its compiled **v1** e-print. That numbering is version-sensitive: [1] numbers each kind of environment sequentially and without a section prefix, so a revision that inserts one statement shifts every later number of that kind.

References

[1] D. Ho, *Consistent model chasing is minimax optimal: the exact value of scalar adversarial adaptive control under large parametric uncertainty*, arXiv:2608.13651v1 [eess.SY], submitted 13 August 2026 (cross-listed cs.LG and math.OC). The source of the problem and of everything attributed above: the game (1)–(3), the value $1 + \Delta$ from rest, the midpoint law (4), the window lemma, the exposure invariant, the one-step reach bound, the value $\max(|x_0|, 1+\Delta|x_0|)$ for $|x_0| \geq 1$, and the deferral quoted in Section 1. Read here in the e-print's own LaTeX source and recompiled, so that the numbers quoted above are its printed v1 numbers; the deferral is in its Section 9 (*Discussion and outlook*), paragraph *Nonzero initial conditions.* arXiv record checked live on 3 September 2026: still one version, sole author, primary category eess.SY, journal-reference and DOI fields empty.

[2] A. Rantzer, *Minimax adaptive control for a finite set of linear systems*, arXiv:2011.10814; Proceedings of the 3rd Conference on Learning for Dynamics and Control, PMLR **144** (2021), 893–904. Abstract read on 3 September 2026 (its abstract states a bounded ℓ_2-gain); venue confirmed against dblp. Cited for the dynamic-game formulation of adaptive control with a finite-dimensional information state, and as an instance of the ℓ_2 criterion discussed in Section 1.7.

[3] A. Rantzer, *Minimax adaptive control for state matrix with unknown sign*, arXiv:1912.03550. Abstract read on 3 September 2026; cited as above. (The arXiv record carries no journal reference, only a note that the posting corrects a version published at the IFAC World Congress 2020; we cite the e-print.)

[4] O. Kjellqvist and A. Rantzer, *Output feedback minimax adaptive control*, arXiv:2409.04115. Abstract read on 3 September 2026; cited as above.

[5] A. Rantzer, *Minimax optimal dual control — the single input case*, arXiv:2604.18550. Author list, title and abstract read from the arXiv API on 3 September 2026; cited as above.

[6] A. Rantzer, *Minimax optimal adaptive control for systems on cones*, arXiv:2511.22758. Author list, title and abstract read from the arXiv API on 3 September 2026; cited as above.

[7] F. Bencherki and A. Rantzer, *Minimax adaptive control for finite sets of positive linear systems*, arXiv:2607.26816. Author list, title and abstract read from the arXiv API on 3 September 2026, which is also where the identifier–title pairing was checked; cited as the ℓ_1 instance of the same line, its abstract stating a bounded ℓ_1-gain from disturbances to errors.

[8] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science, Springer, 2021, 625–635, doi:10.1007/978-3-030-79876-5_37.

[9] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381, doi:10.1145/3372885.3373824.