

Regularity and Stability Properties of Selective SSMs with Discontinuous Gating

Nikola Zubić

Robotics and Perception Group
University of Zurich

zubic@ifi.uzh.ch

Davide Scaramuzza

Robotics and Perception Group
University of Zurich

sdavide@ifi.uzh.ch

Reviewed on OpenReview: <https://openreview.net/forum?id=7Vav53cDeN>

Abstract

Selective State-Space Models (SSMs) such as Mamba have become central to long-sequence modeling. Still, their stability is poorly understood: their state-space coefficients are modulated online by a token-dependent gating signal, making the recurrence neither linear time-invariant nor classically nonlinear. We study continuous-time selective SSMs through passivity, dissipativity, and Input-to-State Stability (ISS), explicitly separating the *selection signal* $x(\cdot)$ from the *driving input* $u(\cdot)$. We obtain four results: exponential forgetting under strict dissipativity; a canonical AUC_{loc} quadratic storage for the frozen-selection subsystem that accommodates discontinuous gating; a parametric LMI together with universal kernel constraints and “irreversible forgetting” under universal quadratic storage; and sufficient conditions for global ISS uniformly over admissible selection schedules. We then bridge to practice by deriving a sampled block LMI for the Mamba selective-scan core, which is used as a differentiable training-time regularizer. Across seven standard time-series datasets and four prediction horizons, the regularizer reduces sampled Mamba-core LMI violations by roughly 92% in 28/28 pairs at a clean-MSE cost of less than 0.018%. It improves internal Mamba passivity and state-norm diagnostics under injected perturbations. Our results turn classical control-theoretic tools into verifiable structural and training criteria for selective SSMs, while honestly scoping which guarantees transfer to a deep selective-scan architecture.

1 Introduction

Modern long-sequence models such as Mamba (Gu & Dao, 2024), S7 (Soydan et al., 2024), HGRN (Qin et al., 2024), and GLA (Yang et al., 2024) are usually described as state-space recurrences whose coefficients are not fixed but are computed online from each token. Concretely, in a continuous-time abstraction of these architectures, we will work with

$$\begin{cases} \dot{h}(t) = A(\Delta(t), x(t)) h(t) + B(\Delta(t), x(t)) u(t), \\ y(t) = C(\Delta(t), x(t)) h(t), \end{cases} \quad (1)$$

where $h(t) \in \mathbb{C}^N$ is the recurrent state, $\Delta(t)$ is a gating signal, $x(t)$ is a *selection* (or scheduling) signal that determines the state-space matrices, and $u(t)$ is the driving (port) input that appears in the energy supply rate. In token-selective architectures, it is natural to set $x(\cdot) \equiv u(\cdot)$ (token, gate, and drive collapse into a single signal), but, as we will see, it is analytically much cleaner to keep them separate.

The selectivity of A, B, C on $x(t)$, and the typically piecewise-constant nature of the gating $\Delta(t)$, place equation 1 in an awkward position. It is not a linear time-invariant system, nor a classical linear time-varying (LTV) system, because the time-varying coefficients are themselves data-dependent, and it is not a

classical smooth nonlinear control system either, because gating introduces *discontinuities*. Existing analyses of these architectures focus mostly on expressivity (Cirone et al., 2024; Zubic et al., 2025; Zubić et al., 2026) or on Lyapunov-exponent-style stability arguments that treat the recurrence as a fixed dynamical system per input (Halloran et al., 2024). What is missing is a unified *energy-based* account: when do these models forget initial conditions, how should we measure “how much of the state matters”, and what does any of this tell us about how to design or regularize a real Mamba implementation?

Passivity and dissipativity, in the sense of Willems (Willems, 1972; Van der Schaft, 2000), and Input-to-State Stability (ISS) (Sontag et al., 1989; Sontag & Wang, 1995; Jiang et al., 1999), were developed exactly for systems whose state interacts with an external port via a supply rate $\mathfrak{R}\langle u, y \rangle$. They are robust to time variation, they have a natural notion of energy storage, and recent work (Morandin & Hinsén, 2024) has shown that for passive LTV systems the minimal available storage is automatically quadratic, with controlled regularity, even when the coefficients are merely L_{loc}^p . This is the right setting for selective SSMs: gating produces L_{loc}^p (in fact piecewise-constant) coefficients, and we want to talk about energy stored in the recurrent state $h(t)$.

A key conceptual move in this paper is to keep two signals separate. The selection signal $x(\cdot)$ *schedules* the dynamics through A, B, C , while the port input $u(\cdot)$ *drives* the dynamics and appears in the supply rate $\mathfrak{R}\langle u, y \rangle$. In a token-selective architecture, both roles are played by the same vector, $x \equiv u$, but treating them as one signal hides the structure of the problem: passivity and ISS quantifiers are statements about *all admissible* u , while structural results about gating are statements about *all admissible* x *at almost every time*. Once these quantifiers are separated, the LMI we will derive admits a natural reading as a parametric condition that must hold for every selection value x that gating can produce, while u is absorbed by quadratic completion.

Standard ingredients (used)	What this paper adds (new)
Dissipative systems and passivity (Willems, 1972; Van der Schaft, 2000).	A clean separation of selection $x(\cdot)$ vs. port input $u(\cdot)$ for selective SSMs, and a derivation of the implied passivity constraints.
ISS theory and ISS-Lyapunov functions (Sontag et al., 1989; Jiang et al., 1999).	Global ISS w.r.t. $u(\cdot)$ uniformly over admissible selection schedules, under a uniform LMI on (A, B, C) in x .
Quadratic storage for passive LTV systems with L_{loc}^p coefficients (Morandin & Hinsén, 2024).	Specialisation to the <i>frozen-selection</i> subsystem $x(t) \equiv 0$ of a selective SSM, yielding $Q_0 \in \text{AUC}_{\text{loc}}^p$ and rank-monotonicity in the gated, switching regime.
KYP-type LMIs for LTV systems with given coefficients (Kalman, 1963; Anderson & Vongpanitlerd, 2013).	A <i>parametric</i> LMI in x that must hold uniformly over all admissible selection values, and the derived universal kernel/unobservability and “irreversible forgetting” constraints.
Continuous-time analyses of deep SSMs (Halloran et al., 2024; Gu et al., 2022).	A discrete one-step dissipativity inequality and sampled block LMI for the Mamba selective-scan recurrence, used as a differentiable regularizer on a real forecasting architecture (Section 6 and Section 8).

Table 1: Standard control-theoretic ingredients we leverage, vs. what this paper contributes for selective SSMs.

Concretely, we make the following contributions.

1. **A clean two-signal model.** We formulate continuous-time selective SSMs as port-Hamiltonian-style systems with selection signal $x(\cdot)$ and driving input $u(\cdot)$, under mild L_{loc}^p regularity that is satisfied by piecewise-constant gating (Section 3).
2. **Forgetting from strict dissipativity.** We show that state-strict dissipativity ($\beta > 0$) plus quadratic bounds on a storage functional implies exponential decay of homogeneous trajectories,

uniformly over admissible selection schedules (Theorem 4.1). This is our formal statement of “forgetting” for selective SSMs.

3. **Frozen-selection structure.** We show that freezing the selection ($x(t) \equiv 0$) yields a genuinely passive LTV input–output subsystem, whose minimal available storage is necessarily quadratic with AUC_{loc} regularity, even under discontinuous gating (Theorem 4.2). This identifies a canonical intrinsic energy structure of selective SSMs and explains why principled initializations such as HiPPO (Gu et al., 2020) matter (Section 4).
4. **Structural consequences of universal quadratic storage.** Under the strong but analytically informative hypothesis that a *single* quadratic storage certifies passivity for *all* admissible selection schedules, we derive a parametric LMI in x (Theorem 5.2) and universal kernel/unobservability conditions on gating, and we formalize “irreversible forgetting”: once a state direction is energy-less, structure prevents any selection from making it observable again (Theorem 5.4).
5. **Robust stability under driving inputs.** We give sufficient conditions for global ISS with respect to $u(\cdot)$, uniformly over admissible selection schedules (Theorem 7.1). This is the practically useful counterpart of the structural results.
6. **Continuous-to-discrete bridge for Mamba.** We translate the continuous-time LMI into a one-step discrete dissipativity inequality and a *sampled block LMI* on the Mamba selective-scan recurrence (Section 6). The largest positive eigenvalue of this matrix is differentiable with respect to the model parameters, yielding an explicit Mamba-core regularizer.
7. **Empirical validation on a real Mamba/SST architecture.** On seven standard forecasting datasets (ETTh1/2, ETTm1/2, weather, electricity, traffic) and four prediction horizons, we evaluate the discrete LMI regularizer inside SST (Xu et al., 2025), which uses Mamba as its long-range expert. Across all 28 dataset–horizon pairs we observe roughly 92% reduction in mean Mamba-core LMI violation while clean MSE remains within 0.018% of the unregularized baseline. We also show a conservativeness ablation in which a learnable diagonal storage $P = \text{diag}(q)$ Pareto-improves over $P = I$ in 28/28 pairs at no measurable accuracy cost, and a robustness study under injected perturbations (Gaussian noise, spikes, block masks, amplitude shifts, last-value-hold) where the regularizer improves internal Mamba passivity and state-norm diagnostics but does not by itself produce a robust-MSE improvement (Section 8).
8. **Honest scoping of a router-energy ablation.** We additionally test whether the SST router can learn an *energy-aware* long/short selection schedule by adding a stability-aware penalty. The penalty preserves clean accuracy and consistently nudges the Mamba branch downward, but the train-time correlation between Mamba weight and local LMI violation is essentially zero on aggregate and inconsistent in sign across datasets. We report this as a neutral result and as a concrete pointer to perturbation-aware router training (Section 8, Appendix C.5).

Reading guide and what is structural vs. practical. Sections 3–7 are the continuous-time analysis. Section 4 establishes baseline forgetting and the canonical quadratic structure of the frozen-selection subsystem; this is the part that justifies, in energy-based terms, why initializations such as HiPPO are well-behaved. Section 5 examines what happens when a single quadratic storage is required to certify passivity uniformly over selection schedules; the resulting parametric LMI and irreversible-forgetting picture is a strong but useful sufficient lens, not a claim that every Mamba in the wild satisfies it. Section 7 is the practically useful global ISS result. Section 6 connects this to the actual Mamba selective-scan recurrence, and Section 8 reports the empirical evaluation on SST.

We emphasize from the outset that our empirical claim is intentionally narrow. The continuous-time theorems give structural and ISS guarantees under explicit hypotheses. The discrete LMI we derive in Section 6 is a sampled, locally-linearized criterion on the Mamba selective-scan core, not a passivity certificate for the full deep network with normalization, residual connections, and routing. Accordingly, in Section 8 we report what the regularizer actually does: it substantially reduces sampled Mamba-core LMI violations and improves internal stability diagnostics with negligible clean-accuracy cost, and we are explicit about what it does *not*

yet do (e.g., produce a robust-MSE improvement, or induce a strongly energy-aware router schedule). All formal proofs are deferred to Appendix A, the original toy simulations that illustrate the continuous-time theory are in Appendix B, and extended experimental details are in Appendix C.

2 Related Work

The analysis of stability and energy-based properties in dynamical systems has a rich history, and most of the technical ingredients we use are classical. Our goal in this section is to make precise which of those ingredients we leverage and where our contribution lies relative to them. The high-level summary is in Table 1 of Section 1.

2.1 Dissipativity, Passivity, and LTV Systems

The foundational theory of dissipative systems, pioneered by Willems (1972), provides a general framework for analyzing systems based on energy-like storage functions and supply rates. Passivity, a special case where the supply rate is the input–output inner product, is central to understanding robust stability and interconnection (Van der Schaft, 2000). The Kalman–Yakubovich–Popov (KYP) lemma established a vital link between frequency-domain passivity conditions and state-space properties for LTI systems (Kalman, 1963; Popov, 1961), and was later extended to LTV systems, often involving time-varying Riccati equations or differential/integral inequalities (Anderson & Vongpanitlerd, 2013).

Particularly relevant to our setting is recent work by Morandin & Hinsin (2024), which rigorously investigates quadratic storage functions for passive LTV systems under *minimal regularity assumptions* on the coefficients (A, B, C, D) . Two of their results matter for us. First, every passive LTV input–output system in their setting admits a minimal quadratic available storage induced by a matrix function $Q \in \text{AUC}_{\text{loc}}$. Second, this minimal Q has weakly non-increasing rank in time. Our use of these results is specific: we apply them to the *frozen-selection* input–output subsystem of a selective SSM (Section 4), where they immediately yield a canonical intrinsic energy structure of AUC_{loc} regularity even when the gating signal $\Delta(t)$ is piecewise-constant. We do not extend this LTV theory, what we add is the recognition that selective SSMs contain such a passive LTV subsystem after we separate the selection signal from the port input, and the structural consequences this has on gating (Section 5).

2.2 Stability of Time-Varying and Switched Systems

The input-selectivity of modern SSMs creates dynamics that behave like switched systems, where the “switching signal” is the input data itself. This necessitates tools for analyzing systems with discontinuous parameter changes (Liberzon, 2003), such as methods involving common or multiple Lyapunov functions (Branicky, 1998). The formalisms of Filippov (1988) and Coddington et al. (1956) provide the theoretical bedrock for guaranteeing that solutions to our system exist under the mild L^p_{loc} regularity we assume, which is, in particular, satisfied by piecewise-constant gating. Our work builds directly on this foundation. We move beyond proving solution existence to ask how a single, coherent energy structure, an AUC_{loc} quadratic storage function, can persist across these data-driven switches, and what structural constraints this persistence imposes on the gating mechanism of a selective SSM.

2.3 Input-to-State Stability (ISS) for Nonlinear Systems

The ISS framework (Sontag et al., 1989; Sontag & Wang, 1995; Jiang et al., 1999) provides a robust notion of stability for nonlinear systems subject to external inputs, characterized by ISS-Lyapunov functions. It quantifies how the system state is affected by both initial conditions and input magnitudes. While classical ISS theory often assumes smooth dynamics, its principles transfer naturally to selective SSMs after we make the two-signal decomposition explicit: the selection signal $x(\cdot)$ renders the dynamics schedule-dependent, and the port input $u(\cdot)$ enters as the perturbation in the supply rate. Our contribution in this direction (Theorem 7.1) adapts ISS-Lyapunov reasoning by seeking a common quadratic Lyapunov function whose decay condition holds *uniformly* across all admissible selection values, giving a specific and verifiable condition

under which the selective SSM is globally ISS w.r.t. $u(\cdot)$. Conceptually, this is a direct application of standard ISS technology; the work is in stating the right uniform-in- x hypothesis.

2.4 Discretization, Sampled-Data, and Mamba-style Selective Scans

A separate strand of control-theoretic work concerns the conditions under which continuous-time stability or passivity properties survive sampling and discretization. For LTV and switched systems, sampled-data passivity has been studied extensively (Liberzon, 2003; Anderson & Vongpanitlerd, 2013). The exact discretization used in Mamba’s selective scan (Gu & Dao, 2024) can be viewed as a particular zero-order-hold-style sampling rule applied per token, with input-dependent step size $\Delta(t)$. In Section 6 we use this view to write the Mamba selective-scan recurrence as a one-step affine update, and to derive a sampled block LMI on its parameters that plays the role of the continuous-time LMI from Theorem 5.2. We do not claim novelty in the general principle of sampled-data passivity, what is new is the explicit instantiation for the Mamba selective-scan core and its use as a differentiable training-time regularizer.

2.5 Stability and Dynamics of Modern SSMs in Deep Learning

Recent deep learning SSMs such as S4 (Gu et al., 2022), S5 (Smith et al., 2023), and particularly Mamba (Gu & Dao, 2024) (and its variants), have shown remarkable performance, and theoretical analyses are emerging. Halloran et al. (2024) analyzes Mamba’s stability via Lyapunov exponents, showing non-positive maximal exponents and inferring robustness to small perturbations. Our work is complementary: instead of an exponent-based diagnosis of an already-trained Mamba, we provide explicit dissipativity-based conditions and a discrete LMI that can be evaluated, and used as a regularizer, during training of a real Mamba-based architecture (Section 6, Section 8). Other works connect selective SSMs to controlled differential equations and rough-path theory to explain expressivity (Kidger et al., 2020; Lyons, 2014; Cirone et al., 2024; Zubic et al., 2025; Zubić et al., 2026). While the present paper takes a different (energy-based, control-theoretic) angle, these works underscore that selective SSMs sit at a rich mathematical intersection, and we view our contribution as another concrete bridge between that mathematics and the design of selective-scan architectures.

3 Preliminaries: Selective SSMs and Passivity Framework

This section introduces the continuous-time selective SSMs that are the focus of our analysis, along with the fundamental control-theoretic concepts that underpin our approach. To make the rest of the paper easier to read for an ML-oriented audience, we first introduce the small piece of notation that will recur throughout (Section 3.1), then explain selective SSMs intuitively (Section 3.2) before formalizing them as the system equation 2, and finally give the energy-based definitions (passivity, dissipativity, quadratic storage) that the analysis is built on.

3.1 Notation at a Glance

We collect the recurring notation in Table 2. The most important point is the distinction between two signals: $x(\cdot)$ is the *selection (scheduling) signal* that determines the time-varying coefficients A, B, C at each time, while $u(\cdot)$ is the *driving (port) input* that appears in the supply rate $\Re\langle u, y \rangle$. In token-selective architectures, these are typically the same vector ($u \equiv x$), but the analysis is much cleaner when we keep them separate, see the remark after equation 2. For matrices and vectors, M^H denotes the Hermitian (conjugate) transpose; $\|v\|$ is the standard Euclidean norm; $\langle u, v \rangle = v^H u$ is the inner product; and $\mathbf{S}_+^n(\mathbb{C})$ denotes the cone of $n \times n$ Hermitian positive-semidefinite matrices, with $Q \succeq 0$ used as shorthand for $Q \in \mathbf{S}_+^n(\mathbb{C})$.

3.2 Selective SSMs, Intuitively

Modern selective architectures such as Mamba (Gu & Dao, 2024), HGRN (Qin et al., 2024), and GLA (Yang et al., 2024) can all be described, at a continuous-time level of abstraction, in two steps:

Symbol	Meaning
$\mathbb{T} \subseteq [0, +\infty)$	Continuous-time interval on which the system evolves.
$h(t) \in \mathbb{C}^N$	Recurrent state of the selective SSM.
$x(t) \in \mathbb{C}^{d_{\text{in}}}$	<i>Selection (scheduling) signal</i> that parametrizes A, B, C via gating.
$u(t) \in \mathbb{C}^{d_{\text{in}}}$	<i>Driving (port) input</i> ; appears in the supply rate $\Re\langle u, y \rangle$.
$y(t) \in \mathbb{C}^{d_{\text{in}}}$	Output of the selective SSM.
$\Delta(t)$	Auxiliary gating signal modulating A, B, C together with $x(t)$.
$A_{\text{eff}}(t), B_{\text{eff}}(t), C_{\text{eff}}(t)$	Effective time-varying matrices obtained by substituting the current $\Delta(t)$ and $x(t)$ into A, B, C (Section 3.4).
$V(t, h)$	Storage function (energy-like, non-negative).
$V_Q(t, h) = \frac{1}{2}h^H Q(t)h$	Quadratic storage induced by a matrix function $Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$.
$\beta \geq 0$	State-strict dissipativity rate; $\beta > 0$ implies intrinsic energy loss.
AUC_{loc}	Locally absolutely upper semicontinuous matrix functions (Morandin & Hinsen, 2024); supports controlled jumps.
$\text{Ker}(Q(t))$	Kernel of the storage matrix at time t (“energy-less” state directions).

Table 2: Notation used throughout the paper. The single most important convention is the separation of the selection signal $x(\cdot)$ (which schedules A, B, C) from the driving (port) input $u(\cdot)$ (which appears in the supply rate).

- (i) A small neural network reads the current token (or, more generally, the current selection signal $x(t)$) and an auxiliary gating signal $\Delta(t)$, and produces, at every time t , a triple of matrices $(A, B, C)(\Delta(t), x(t))$. These matrices are not learned per token from scratch but are computed online from the inputs.
- (ii) These data-dependent matrices are then plugged into a continuous-time linear recurrence on a hidden state $h(t)$, so that the recurrence is locally linear in h but globally non-stationary because (A, B, C) change with $x(t)$ and $\Delta(t)$.

This is what makes selective SSMs distinctive: the recurrence is *linear in the state* but *schedule-dependent*, with the schedule chosen by the data. In particular, the gating signal $\Delta(t)$ is typically piecewise-constant (one value per token), so the effective coefficient matrices undergo discontinuous jumps. This is the regime where AUC_{loc} -style assumptions on the storage matrix become natural, rather than smoothness assumptions.

In practice, the selection signal $x(t)$ and the input that drives the system are often the same vector: a token, treated both as a query for the gating network and as the recurrent input. We will keep the two roles separate at the level of notation, because passivity statements are clearest when “what schedules A, B, C ” and “what appears in the supply rate” are not conflated. Whenever we want the token-selective regime, we simply specialize to the coupling $u(\cdot) \equiv x(\cdot)$ at the end.

3.3 System Definition and Assumptions

We consider dynamical systems evolving in continuous time $t \in \mathbb{T}$, where $\mathbb{T} \subseteq [0, +\infty)$ is a time interval. The internal state of the system at time t is denoted by $h(t) \in \mathbb{C}^N$ (or $\mathbb{R}^N$ in real-valued cases). A key point for our passivity analysis is to distinguish between: (i) a *driving (port) input* $u(t) \in \mathbb{C}^{d_{\text{in}}}$, which is the signal that appears in the supply rate $\Re\langle u, y \rangle$, and (ii) a *selection (scheduling) signal* $x(t) \in \mathbb{C}^{d_{\text{in}}}$ (e.g. the token embedding) that parametrizes the time-varying matrices via input-dependent gating.¹

The system produces an output $y(t) \in \mathbb{C}^{d_{\text{in}}}$. For matrices and vectors, M^H denotes the Hermitian (conjugate) transpose. The standard Euclidean norm of a vector v is $\|v\|$, and the inner product of two vectors u, v is $\langle u, v \rangle = v^H u$. We denote the space of $n \times n$ Hermitian positive-semidefinite matrices by $\mathbf{S}_+^n(\mathbb{C})$; for $Q \in \mathbf{S}_+^n(\mathbb{C})$, we write $Q \succeq 0$.

¹Throughout Sections 3–7 we assume $d_{\text{out}} = d_{\text{in}}$ so that $\langle u, y \rangle$ and the passivity LMI blocks are dimensionally well-defined; this is the common setting in sequence models where the recurrent channel has matched input/output width.

Our work focuses on Continuous-Time Selective SSMs. Inspired by architectures such as Mamba (Gu & Dao, 2024), their defining characteristic is that the system’s core parameters are dynamically modulated by the selection signal $x(t)$ and an auxiliary gating signal $\Delta(t)$. Formally, the dynamics are given by:

$$\begin{cases} \dot{h}(t) = A(\Delta(t), x(t)) h(t) + B(\Delta(t), x(t)) u(t), \\ y(t) = C(\Delta(t), x(t)) h(t). \end{cases} \quad (2)$$

The equation 2 should be read as “a continuous-time abstraction of a selective scan”: at each time t , the gating mechanism produces $(\Delta(t), x(t))$, these are fed into the per-architecture maps that produce the matrices A, B, C , and the resulting linear update acts on the recurrent state $h(t)$ with $u(t)$ as the driving input. We will refer to the substituted matrices $A_{\text{eff}}(t) := A(\Delta(t), x(t))$, $B_{\text{eff}}(t) := B(\Delta(t), x(t))$, and $C_{\text{eff}}(t) := C(\Delta(t), x(t))$ as the *effective* time-varying matrices: they are what the recurrence actually sees once the gating mechanism has produced its outputs at time t . The next subsection states the regularity assumption we make on these effective matrices to guarantee well-posedness.

Remark 3.1 (Coupled-input selective SSMs). In many token-selective architectures, the same signal both *selects* the parameters and *drives* the dynamics; this corresponds to the special case $u(\cdot) \equiv x(\cdot)$. In this paper, we keep $u(\cdot)$ and $x(\cdot)$ conceptually distinct to make the quantifiers in the passivity analysis explicit. Any condition proved for all admissible pairs $(x(\cdot), u(\cdot))$ remains valid (as a sufficient condition) under the coupling $u = x$.

3.4 Well-Posedness and Function Space Assumptions

To ensure the system equation 2 is well-defined even with discontinuous parameter changes, we adopt mild regularity conditions. For any admissible input trajectory $x(\cdot)$, let $A_{\text{eff}}(t) := A(\Delta(t), x(t))$ (and similarly for $B_{\text{eff}}, C_{\text{eff}}$). We assume the effective matrices belong to standard Lebesgue spaces: $A_{\text{eff}}(\cdot) \in L^1_{\text{loc}}$, $B_{\text{eff}}(\cdot) \in L^2_{\text{loc}}$, and $C_{\text{eff}}(\cdot) \in L^2_{\text{loc}}$, with the selection signal $x(\cdot) \in L^2_{\text{loc}}$ and the driving input $u(\cdot) \in L^2_{\text{loc}}$. These conditions ensure that the ODE for $h(t)$ satisfies Carathéodory conditions, guaranteeing the local existence and uniqueness of an absolutely continuous solution $h(\cdot) \in W^{1,1}_{\text{loc}}$ (Filippov, 1988; Coddington et al., 1956).

Remark 3.2 (Feasibility of Regularity Assumptions). These assumptions are mathematically mild and readily satisfied by modern selective SSMs. In token-based architectures such as Mamba, the selection mechanism yields system matrices that are **piecewise-constant**. Such functions are well-behaved members of the L^p_{loc} spaces, confirming that our framework is grounded in practical implementations while being general enough for future architectures.

3.5 Passivity and Dissipativity

We analyze the system equation 2 using the powerful energy-based frameworks of passivity and dissipativity (Willems, 1972). Intuitively, a passive system cannot generate its own energy. It can only store or dissipate energy supplied from its input. A strictly dissipative system is even stronger, as it inherently loses energy over time. These concepts are formalized in the following definitions.

Definition 3.3 (Storage Function). A function $V : \mathbb{T} \times \mathbb{C}^N \rightarrow \mathbb{R}$ is a *storage function* if it is non-negative, i.e., $V(t, h) \geq 0$ for all (t, h) , and typically $V(t, 0) = 0$.

Definition 3.4 (Passivity). The system equation 2 is *passive* if there exists a storage function V such that for all admissible trajectories on any interval $[t_0, T]$:

$$V(T, h(T)) - V(t_0, h(t_0)) \leq \int_{t_0}^T \text{Re} \langle u(\tau), y(\tau) \rangle d\tau. \quad (3)$$

The term $\text{Re} \langle u(\tau), y(\tau) \rangle$ is the instantaneous power, or *supply rate*, provided to the system.

Definition 3.5 (Strict Dissipativity). The system is *strictly dissipative* with rate $\beta > 0$ if it satisfies the stronger inequality:

$$V(T, h(T)) - V(t_0, h(t_0)) \leq \int_{t_0}^T \text{Re} \langle u(\tau), y(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau. \quad (4)$$

A positive β implies an intrinsic rate of energy dissipation, which is crucial for proving strong stability properties like exponential decay. Given the absolute continuity of $h(\cdot)$ and local Lipschitzness of V in h , these integral inequalities have corresponding differential forms: $\frac{d}{dt}V \leq \text{Re} \langle u, y \rangle$ for passivity, and $\frac{d}{dt}V \leq \text{Re} \langle u, y \rangle - \beta \|h\|^2$ for strict dissipativity.

3.6 Quadratic Storage Functions and AUC Regularity

A significant portion of our analysis focuses on *quadratic storage functions* of the form $V_Q(t, h) = \frac{1}{2} h^H Q(t) h$, where $Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$ is a time-varying matrix. The regularity of $Q(t)$ is critical in our setting, as the gating mechanism can induce discontinuities in the system dynamics.

To handle this robustly, we consider $Q(t)$ to belong to the class of Locally Absolutely Upper Semicontinuous matrix functions, denoted $Q \in \text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$, following the framework of (Morandin & Hinsin, 2024). For our purposes, AUC_{loc} allows controlled discontinuities and ensures $\dot{Q}(t)$ exists a.e. Under the additional hypothesis that Q induces a universal quadratic storage for a passive LTV subsystem implies further structural constraints, including rank monotonicity (Theorem 5.1). This property is the foundation for the “irreversible forgetting” we analyze later. Formally, a function $Q \in \text{AUC}_{\text{loc}}$ is characterized by being of locally bounded variation (BV_{loc}), having a derivative $\dot{Q}(t)$ that exists almost everywhere, and satisfying certain integral and jump conditions. The key property that enforces rank monotonicity is that the singular part of its decomposition must be weakly monotonically decreasing in the Loewner order (e.g., at any jump, $\lim_{\tau \rightarrow t^-} Q(\tau) \succeq Q(t)$). This mathematical structure is precisely what makes the framework suitable for analyzing switched systems while retaining essential properties about their energy evolution.

4 Stability and Structural Properties from Passivity

We establish foundational stability guarantees and explore the core structural properties of passive selective SSMs. We begin by analyzing the *homogeneous frozen-selection* case, where the selection signal is fixed to a reference value $x(t) \equiv 0$ and the port input is turned off $u(t) \equiv 0$. The resulting dynamics $\dot{h}(t) = A(\Delta(t), 0)h(t)$ are determined entirely by the model’s initialization, allowing us to isolate and understand its intrinsic stability and memory properties. As we demonstrate in our simulation study in Appendix B.1, this baseline behavior is a powerful predictor of a model’s practical performance, and it underpins the empirical observation that the LMI-regularizer introduced in Section 6 preserves clean MSE in our SST/Mamba experiments (Section 8). We first show how strict dissipativity leads to exponential forgetting, and then prove that the unforced dynamics possess a fundamental, well-behaved quadratic energy structure.

4.1 Exponential Decay from Strict Dissipativity

The most fundamental stability guarantee arises when the system intrinsically dissipates energy faster than it stores it, even with zero input. This leads to exponential convergence of the state to the origin.

Theorem 4.1 (Exponential Decay from Strict Dissipativity). Consider the continuous-time selective state-space model defined in Eq. equation 2. Suppose there exists a storage functional $V : [0, \infty) \times \mathbb{C}^N \rightarrow \mathbb{R}_{\geq 0}$ satisfying:

- (i) **Strict Dissipativity Inequality:** For some constant $\beta > 0$, every admissible pair $(x(\cdot), u(\cdot))$ and the corresponding state–input–output trajectory $\{h(\tau), u(\tau), y(\tau)\}$ on any interval $[t_0, T]$ satisfies:

$$V(T, h(T)) - V(t_0, h(t_0)) \leq \int_{t_0}^T \text{Re} \langle u(\tau), y(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau. \quad (5)$$

- (ii) **Regularity and Quadratic Bounds:** $V(t, h)$ is locally Lipschitz in h , absolutely continuous in t along system trajectories, and there exist constants $k_2 \geq k_1 > 0$ such that for all $t \geq 0$ and $h \in \mathbb{C}^N$:

$$k_1 \|h\|^2 \leq V(t, h) \leq k_2 \|h\|^2. \quad (6)$$

Then, for any admissible selection schedule $x(\cdot)$, the corresponding homogeneous (zero port-input) trajectory (i.e., when $u(t) \equiv 0$ for $t \geq 0$) exhibits exponential decay: there exist constants $C \geq 1$ and $\gamma > 0$ such that for any initial state $h(0)$, the solution satisfies:

$$\|h(t)\| \leq C e^{-\gamma t} \|h(0)\| \quad \text{for all } t \geq 0. \quad (7)$$

Proof Sketch. The proof considers the system without port input, i.e., $u = 0$ (so the supply term $\text{Re} \langle u, y \rangle$ vanishes). In this case, the strict dissipativity inequality simplifies to $\frac{dV}{dt} \leq -\beta \|h\|^2$. Since the energy function V is assumed to be quadratically bounded, satisfying $k_1 \|h\|^2 \leq V \leq k_2 \|h\|^2$, this implies $\frac{dV}{dt} \leq -\gamma V$ for some constant $\gamma > 0$. This is a classic differential inequality whose solution is an exponential decay. The bounds on V then directly translate this exponential decay of energy into an exponential decay of the state norm $\|h(t)\|$. Check Appendix A.1 for the full proof. $\square$

ML implication. *For a selective SSM, this is the formal statement of “forgetting”: as long as the architecture admits any storage functional V with strict dissipativity ($\beta > 0$) and quadratic bounds, the recurrent state decays exponentially in the absence of external drive, irrespective of which selection schedule $x(\cdot)$ the gating produces. Whether such a V exists is a property of the architecture and its initialization; the storage functional we will exhibit in Section 4.2 (and use as the reference energy for HiPPO-style initialization in Appendix B.1) is one concrete witness.*

4.2 Inherent Quadratic Structure and Regularity in Frozen-Selection LTV Dynamics

While Theorem 4.1 guarantees stability under strict dissipativity, it does not specify the form of the storage function V . We now connect selective SSMs to the passive LTV framework of Morandin & Hinsen (2024). The key point is that the results of Morandin & Hinsen (2024) apply to *input-output* LTV systems of the form

$$\dot{h}(t) = A(t)h(t) + B(t)u(t), \quad y(t) = C(t)h(t) + D(t)u(t), \quad (8)$$

with mild regularity on (A, B, C, D) and passivity defined via the supply rate $\Re \langle u, y \rangle$ (Morandin & Hinsen, 2024, Eq. (1), Def. 1.1). Accordingly, instead of removing the input channel, we *freeze the selection signal* (e.g., set $x(t) \equiv 0$) and obtain a genuine LTV input-output subsystem with port input $u(\cdot)$. For passive LTV systems, the available storage is finite and coincides with a minimal quadratic storage $V_a(t, h) = \frac{1}{2} h^H Q(t) h$ with $Q \in \text{AUC}_{\text{loc}}$ (Morandin & Hinsen, 2024, Def. 4.1, Cor. 4.3). Evaluating this quadratic storage along homogeneous (zero port-input) trajectories $u \equiv 0$ yields a canonical quadratic energy functional for the frozen-selection baseline dynamics (i.e., $x \equiv 0$ and $u \equiv 0$).

To prove this, we leverage the concept of the *available storage* function, which represents the maximum energy that can be extracted from the system from a given state. It is a standard result in dissipativity theory that if any storage function exists, then the available storage function is well-defined and is itself a valid storage function (see Appendix A.2). The following theorem connects this concept to the structure of selective SSMs.

Theorem 4.2 (Existence and AUC_{loc} regularity of a minimal quadratic storage for the frozen-selection LTV subsystem). Consider the selective SSM equation 2 with selection signal $x(\cdot)$ and driving input $u(\cdot)$. Assume the coefficient maps satisfy the standard well-posedness regularity: for admissible $x(\cdot)$ we have

$$A(\Delta(\cdot), x(\cdot)) \in L^1_{\text{loc}}(\mathbb{T}), \quad B(\Delta(\cdot), x(\cdot)) \in L^2_{\text{loc}}(\mathbb{T}), \quad C(\Delta(\cdot), x(\cdot)) \in L^2_{\text{loc}}(\mathbb{T}),$$

and $u(\cdot) \in L^2_{\text{loc}}(\mathbb{T})$. (Here $D \equiv 0$, hence $D \in L^\infty_{\text{loc}}$ holds trivially.)

Suppose there exists a storage functional $V : \mathbb{T} \times \mathbb{C}^N \rightarrow \mathbb{R}_{\geq 0}$ with $V(t, 0) = 0$ such that for some $\beta \geq 0$ the dissipativity inequality

$$V(T, h(T)) - V(t_0, h(t_0)) \leq \int_{t_0}^T \text{Re} \langle u(\tau), y(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau \quad (9)$$

holds for all admissible pairs $(x(\cdot), u(\cdot))$ and all corresponding trajectories of equation 2. (In particular, dropping the nonpositive term $-\beta \int \|h\|^2$ yields the standard passivity inequality with supply rate $\Re \langle u, y \rangle$.)

Fix the *frozen-selection* signal $x(t) \equiv 0$ and define the associated LTV input–output system

$$\begin{cases} \dot{h}(t) = A_0(t)h(t) + B_0(t)u(t), \\ y_0(t) = C_0(t)h(t), \end{cases} \quad \text{where } A_0(t) = A(\Delta(t), 0), B_0(t) = B(\Delta(t), 0), C_0(t) = C(\Delta(t), 0). \quad (10)$$

Then the following hold:

- (a) **Passivity of the frozen-selection LTV system.** The LTV system equation 10 is passive in the sense of Morandin & Hinsén (2024, Def. 1.1). Indeed, restricting equation 9 to $x(\cdot) \equiv 0$ yields a dissipation inequality for every state–input–output solution (h, u, y_0) of equation 10.
- (b) **Existence of minimal quadratic available storage.** Let $V_{a,0}$ denote the available storage of equation 10 in the sense of Morandin & Hinsén (2024, Def. 4.1). Since equation 10 is passive and satisfies the regularity assumptions of Morandin & Hinsén (2024, Eq. (1)), its available storage is finite and coincides with a quadratic storage,

$$V_{a,0}(t, h) = \frac{1}{2}h^H Q_0(t)h,$$

for some matrix function $Q_0 : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$, and $V_{a,0}$ is minimal among all storage functions (Morandin & Hinsén, 2024, Cor. 4.3).

- (c) **AUC_{loc} regularity of $Q_0(t)$.** The inducing matrix function can be chosen with AUC_{loc} regularity:

$$Q_0 \in \text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C})),$$

again by Morandin & Hinsén (2024, Cor. 4.3).

Finally, the homogeneous (unforced) dynamics of the frozen-selection system correspond to setting $u \equiv 0$ in equation 10, i.e., $\dot{h} = A_0(t)h$. If, in addition, V satisfies the hypotheses of Theorem 4.1 (strict dissipativity $\beta > 0$ and quadratic bounds), then these homogeneous trajectories decay exponentially as stated in Theorem 4.1.

Proof Sketch. The frozen-selection system equation 10 is an LTV input–output system of the form

$$\dot{h}(t) = A_0(t)h(t) + B_0(t)u(t), \quad y_0(t) = C_0(t)h(t) + D_0(t)u(t),$$

with $D_0 \equiv 0$. Under our assumptions we have $A_0 \in L_{\text{loc}}^1$, $B_0, C_0 \in L_{\text{loc}}^2$, and $D_0 \in L_{\text{loc}}^\infty$, hence equation 10 fits the setting of Morandin & Hinsén (2024, Eq. (1)).

Since equation 9 holds for all admissible $(x(\cdot), u(\cdot))$, it holds in particular when restricting to $x(\cdot) \equiv 0$. Dropping the nonpositive term $-\beta \int \|h\|^2$ yields the passivity dissipation inequality of Morandin & Hinsén (2024, Def. 1.1) for equation 10, so the frozen-selection LTV system is passive.

For passive LTV systems, Morandin & Hinsén (2024, Cor. 4.3) shows that the available storage (Morandin & Hinsén, 2024, Def. 4.1) is finite and coincides with a minimal quadratic storage $V_{a,0}(t, h) = \frac{1}{2}h^H Q_0(t)h$ with $Q_0 \in \text{AUC}_{\text{loc}}$. Finally, the homogeneous dynamics correspond to fixing $u \equiv 0$ (cf. Morandin & Hinsén (2024, Sec. 2.1)), and exponential decay under strict dissipativity follows from Theorem 4.1. $\square$

ML implication. *This is the canonical intrinsic energy of a selective SSM at its initialization. Once the gating mechanism produces piecewise-constant matrix updates ($\Delta(t)$ is piecewise-constant per token), the matrix $Q_0(t)$ that measures “how much state has been usefully stored” is in AUC_{loc}: it can jump, but only in controlled (Loewner-monotone-decreasing) ways. This is exactly the mathematical regularity needed for principled initializations such as HiPPO (Gu et al., 2020) to behave well as a baseline energy landscape on which the input-driven selection then operates. Empirically, we observe in Appendix B.1 that initializations admitting such a well-conditioned Q_0 (HiPPO-style) yield slow, stable energy decay under input removal, while initializations failing this admit no valid Q_0 and diverge. The discrete-time analogue of this storage matrix becomes the matrix P that we regularize in Section 6 and Section 8.*

5 Constraints Imposed by Universal Quadratic Passivity on Gating Mechanisms

Building on the foundational properties of the unforced system, this section studies the inherent design constraints of selective SSMs. We investigate the consequences of a strong but informative assumption: that a *single*, universal quadratic storage function $V_Q(t, h)$ guarantees passivity across all possible selection-driven (schedule-dependent) dynamics. The power of this assumption lies not in its generality, but in the strict, necessary consequences it imposes on the gating mechanism. This approach allows us to derive concrete constraints in the form of LMIs and to formalize a notion of "irreversible forgetting".

We emphasize from the outset that the universal-storage hypothesis is a *sufficient* structural condition: it is what we use to expose what passivity “would have to” force on a selective SSM if a single quadratic energy were sufficient to certify it for every gating decision. We do *not* claim that every Mamba in the wild satisfies it; rather, the hypothesis isolates a clean structural picture (a parametric LMI in x , kernel/unobservability constraints, and irreversible forgetting) which then motivates, in Section 6, the discrete LMI we use as a differentiable training-time regularizer on a real Mamba-based architecture, where the LMI is enforced as a soft constraint rather than as a fact about the architecture.

5.1 Regularity and Rank Monotonicity of Universal Quadratic Storage

If a single quadratic function V_Q is capable of certifying passivity regardless of the input-driven variations in system parameters, the matrix $Q(t)$ defining this function must possess inherent structural regularity and obey specific monotonicity properties.

Theorem 5.1 (Regularity and Rank of Universal Quadratic Storage). Consider the continuous-time selective SSM equation 2 under the standard regularity assumptions. Suppose there exists a time-varying quadratic storage function $V_Q(t, h) = \frac{1}{2}h^H Q(t)h$ (with $Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$) that satisfies the dissipativity inequality (which implies passivity by dropping the nonpositive term):

$$V_Q(T, h(T)) - V_Q(t_0, h(t_0)) \leq \int_{t_0}^T \operatorname{Re} \langle u(\tau), y(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau \quad (11)$$

for some $\beta \geq 0$, valid for *all* admissible pairs $(x(\cdot), u(\cdot))$ and all corresponding state–input–output trajectories $\{h(\tau), u(\tau), y(\tau)\}$ of equation 2. Then, the matrix function $Q(t)$ must satisfy the following properties:

- (a) **AUC Regularity:** $Q(t)$ must belong to the class of locally absolutely upper semicontinuous matrix functions, i.e., $Q \in \text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$.
- (b) **Rank Monotonicity:** The rank of $Q(t)$, denoted $r(t) = \text{rank}(Q(t))$, must be weakly monotonically non-increasing over time $t \in \mathbb{T}$.

Proof Sketch. The logic here is similar to the previous theorem but applied to the storage matrix $Q(t)$ itself rather than the minimal one. If a single quadratic function V_Q works for all admissible pairs $(x(\cdot), u(\cdot))$, then in particular it works for the frozen-selection LTV subsystem obtained by fixing $x(\cdot) \equiv 0$. Hence V_Q is a valid quadratic storage function for that passive LTV system, and the structural results of Morandin & Hinsen (2024) imply $Q \in \text{AUC}_{\text{loc}}$ and that $t \mapsto \text{rank}(Q(t))$ is weakly non-increasing. The same theoretical results from Morandin & Hinsen (2024) that impose structure on the minimal storage function also apply to any valid quadratic storage function. This forces $Q(t)$ to have AUC_{loc} regularity and, crucially, a non-increasing rank over time. Check Appendix A.4 for the full proof. $\square$

ML implication. For a selective SSM, this says: if there is any single quadratic “energy meter” V_Q that the architecture cannot beat regardless of the data-dependent gating decisions, then the meter itself must obey two structural laws. It is allowed to jump (gating produces piecewise-constant updates), but only in AUC_{loc} ways, and the rank of the matrix $Q(t)$ that defines it cannot grow over time. This is the precise sense in which a selective SSM, viewed through a single passivity certificate, can only ever lose dimensions of usable energy storage. The next two theorems unpack what this rank-non-increasingness means operationally for gating.

5.2 Parametric LMI and Kernel Constraints on Gating

The existence of a universal quadratic storage function V_Q translates into concrete algebraic constraints that the selection-dependent matrices $A(t, x), B(t, x), C(t, x)$ must satisfy *parametrically* for all admissible selection values x . Equivalently, for each fixed selection value x and almost every time t , the resulting quadratic dissipation inequality must hold for all states h and all port inputs u . These constraints take the form of a parametric Linear Matrix Inequality (LMI) and lead to important conditions on how the output matrix $C(t, x)$ interacts with the kernel of the storage matrix $Q(t)$.

Theorem 5.2 (Passivity Constraints on Selection via a Parametric LMI). Consider the selective SSM equation 2 with selection signal $x(\cdot)$ and driving input $u(\cdot)$ under the standard regularity assumptions. Assume the system’s passivity is guaranteed by a *single*, time-varying quadratic storage function

$$V_Q(t, h) = \frac{1}{2} h^H Q(t) h, \quad Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C}),$$

such that for some $\beta \geq 0$ the dissipativity inequality

$$V_Q(T, h(T)) - V_Q(t_0, h(t_0)) \leq \int_{t_0}^T \operatorname{Re} \langle u(\tau), y(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau \quad (12)$$

holds for *all* initial states and for *all* admissible pairs $(x(\cdot), u(\cdot))$.

Then, for almost every $t \in \mathbb{T}$ and for every admissible selection value $x \in \mathbb{C}^{d_{\text{in}}}$, the following parametric LMI holds:

$$\mathcal{L}(t, x) := \begin{bmatrix} \dot{Q}(t) + Q(t)A(t, x) + A(t, x)^H Q(t) + 2\beta I & Q(t)B(t, x) - C(t, x)^H \\ B(t, x)^H Q(t) - C(t, x) & 0 \end{bmatrix} \preceq 0, \quad (13)$$

where $A(t, x) = A(\Delta(t), x)$, $B(t, x) = B(\Delta(t), x)$, $C(t, x) = C(\Delta(t), x)$, and $\dot{Q}(t)$ exists a.e.

In particular, the following universal kernel condition must hold:

$$\forall v \in \operatorname{Ker}(Q(t)), \quad C(\Delta(t), x)v = 0 \quad \text{for all admissible } x \text{ at time } t, \text{ a.e. } t \in \mathbb{T}. \quad (14)$$

Remark 5.3 (Coupled case $u = x$). If the architecture uses a single signal (i.e., $u(\cdot) \equiv x(\cdot)$), then equation 13 remains a strong *sufficient* condition for passivity and motivates the regularizer used in Appendix B.3. Necessity holds under the stronger schedule-independent hypothesis stated in the theorem (passivity for all admissible pairs $(x(\cdot), u(\cdot))$).

Proof Sketch. This proof translates the differential form of the passivity inequality, $\frac{dV_Q}{dt} \leq \dots$, into a matrix inequality. By substituting the system dynamics and the quadratic form of V_Q , we arrive at an expression that must be nonpositive for all states h and port inputs u , for each fixed selection value x (and for a.e. time t). Such an expression can be elegantly represented as a Linear Matrix Inequality (LMI). For the LMI to hold for all h and u (uniformly over admissible selection values x), specific conditions must be met. By choosing a state v from the kernel of Q (where $Qv = 0$), the LMI simplifies dramatically, revealing that the term involving the output matrix C must be zero to prevent the inequality from being violated. This leads to the necessary condition that $C(t, x)v = 0$. Check Appendix A.5 for the full proof. $\square$

ML implication. Equation 13 is the formal statement of the practical regularizer we use later. It is a parametric LMI: it must hold not just at one selection value, but uniformly over every selection value x that the gating network can produce at time t . This is exactly what we sample and penalize during training of a real Mamba in Section 6: each token produces a selection-dependent matrix $\mathcal{L}(t, x)$, and we drive its largest positive eigenvalue toward zero. The kernel condition equation 14 adds a striking structural reading: any state direction the storage matrix considers “energy-less” (in $\operatorname{Ker} Q(t)$) must be unobservable through $C(t, x)$, regardless of which selection value x the gate emits. In Mamba terms, the gating cannot decide to “read out” a direction that the energy structure has already discarded.

5.3 Irreversible Forgetting and Energy Consistency in Kernel Subspace

The constraints derived above, particularly the non-increasing rank of $Q(t)$ and the universal kernel condition for $C(t, x)$, suggest a form of structural irreversibility associated with the modes that fall into the kernel of the storage matrix $Q(t)$. We now explore this concept of "irreversible forgetting" from the perspective of the universal storage function V_Q and examine the energy balance required within this forgotten subspace.

Theorem 5.4 (Inertness of Forgotten Modes and Gating Constraints from Passivity). Consider the continuous-time selective state-space model equation 2 under standard regularity assumptions. Assume its passivity is guaranteed by a *single*, time-varying quadratic storage function $V_Q(t, h) = \frac{1}{2}h^H Q(t)h$ with $Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$, satisfying the dissipativity inequality equation 12 for some $\beta \geq 0$ and for *all* admissible pairs $(x(\cdot), u(\cdot))$. Then, the following properties regarding "forgotten modes" (states in $\text{Ker}(Q(t))$) hold:

1. **Irreversibility in storage dimension (rank monotonicity):** By Theorem 5.1, the matrix function Q can be chosen in $\text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$ and the rank $r(t) = \text{rank}(Q(t))$ is weakly monotonically non-increasing over time. Equivalently, since $Q(t) \in \mathbb{C}^{N \times N}$, $\dim \text{Ker}(Q(t)) = N - r(t)$ is weakly non-decreasing.
2. **Energy Consistency within the Kernel:** For almost every $t \in \mathbb{T}$ where $\dot{Q}(t)$ exists, every direction $v \in \text{Ker}(Q(t))$ satisfies the kernel energy inequality

$$v^H \dot{Q}(t)v + 2\beta \|v\|^2 \leq 0. \quad (15)$$

This ensures the evolution $\dot{Q}(t)$ within the kernel is consistent with the required dissipation β . If $\beta = 0$, $h^H \dot{Q}h \leq 0$; if $\beta > 0$, $h^H \dot{Q}h \leq -2\beta \|h\|^2 < 0$ (for $h \neq 0$).

3. **Constraint on selection schedules:** Under the universal passivity hypothesis, the parametric LMI equation 13 must hold for almost every t and for every admissible selection value x . Hence any selection schedule $x(\cdot)$ that attains values for which equation 13 fails on a set of times of positive measure would contradict the hypothesis. The dynamics must respect the energy structure defined by $Q(t)$ and $\dot{Q}(t)$ within the kernel.

Remark 5.5 (On strict dissipativity and rank-deficient storage). If $\beta > 0$ and $Q(t) \succeq 0$, then the kernel inequality $v^H \dot{Q}(t)v + 2\beta \|v\|^2 \leq 0$ implies that any nonzero $v \in \text{Ker}(Q(t))$ would satisfy $v^H \dot{Q}(t)v < 0$ at times where $\dot{Q}(t)$ exists, which would contradict the requirement $Q(t + \varepsilon) \succeq 0$ for small $\varepsilon > 0$. Thus, under state-strict dissipativity ($\beta > 0$), $Q(t)$ must be positive definite for almost every t , so the kernel-based "irreversible forgetting" discussion is primarily meaningful in the passive case $\beta = 0$.

Proof Sketch. This theorem's proof is a direct consequence of the preceding results.

1. **Kernel Non-Shrinking:** This follows immediately from the rank monotonicity established in Theorem 5.1; if the rank cannot increase, the dimension of the kernel ($N - \text{rank}$) cannot decrease.
2. **Energy Consistency:** This is derived by taking the (1,1) block of the LMI from Theorem 5.2 and evaluating it for a state h in the kernel of Q . This isolates the term $h^H \frac{dQ}{dt} h$ and shows it must be negative enough to account for the required energy dissipation.
3. **Constraint on Dynamics:** This is a logical conclusion: any selection signal x that would induce dynamics violating the LMI contradicts the fundamental assumption of universal passivity, and is therefore inadmissible.

Check Appendix A.6 for the full proof. □

ML implication. Read in the passive case ($\beta = 0$, per the Remark above), this is the "irreversible forgetting" picture: once a state direction enters $\text{Ker } Q(t)$, the kernel cannot shrink, the LMI prevents any selection from making that direction observable through $C(t, x)$, and the energy structure within the kernel must continue to dissipate consistently. For a Mamba-style selective SSM, this offers a structural lens on what gating "cannot undo" under a single passive certificate, and it foreshadows, on the practical side, why a discrete LMI regularizer monotonically reduces violations across training without the model recovering them

later (Section 8, Appendix C). We do not claim every selective SSM in the wild lives in this strong universal-storage regime; rather, the lens identifies the structural bias that the discrete LMI of Section 6 encourages a real Mamba to acquire during training.

6 From Continuous Theory to Discrete Mamba

The continuous-time results of Sections 4–7 give us a clear structural picture: under appropriate dissipativity assumptions, selective SSMs must satisfy a parametric LMI in the selection value x , with kernel/unobservability and rank-monotonicity consequences. In practice, however, modern selective architectures such as Mamba (Gu & Dao, 2024) do not run a continuous-time ODE; they run a discrete-time selective scan with an input-dependent step size Δ_k per token. To make the theory actionable for these architectures, this section translates the LMI of Theorem 5.2 into a sampled-data inequality on the Mamba selective-scan core, and into a differentiable training-time regularizer that we will evaluate empirically in Section 8.

We do not claim novelty in the general principle of sampled-data passivity (Liberzon, 2003; Anderson & Vongpanitlerd, 2013). What is new here is the explicit instantiation of the Mamba selective-scan recurrence and its integration into the training loop of a real architecture. Throughout, we work in the coupled-input regime $u_k \equiv x_k$ (token, gate, and drive collapse into one signal), since this is the regime real Mamba uses; under that coupling, Theorem 5.2 gives a strong sufficient condition for one-step dissipativity rather than a necessity statement, and we use it accordingly.

6.1 Discrete Selective-Scan Dynamics

In the standard Mamba selective scan, at each token index $k \in \{0, 1, \dots, T-1\}$ the model produces from the current input $u_k \in \mathbb{C}^{d_{\text{in}}}$ a token-dependent step size $\Delta_k \in \mathbb{R}_{>0}$ and continuous-time matrices $A(\Delta_k, u_k)$, $B(\Delta_k, u_k)$, $C(\Delta_k, u_k)$, then discretizes the per-token continuous update via a zero-order-hold-style rule. Concretely, this yields a discrete recurrence

$$\begin{cases} h_{k+1} = F_k h_k + G_k u_k, \\ y_k = C_k h_{k+1} + D_k u_k, \end{cases} \quad k = 0, 1, \dots, T-1, \quad (16)$$

where the per-token matrices (F_k, G_k, C_k, D_k) are computed from (Δ_k, u_k) by the architecture’s discretization rule. We stress two structural facts about equation 16:

- (i) For each fixed k , equation 16 is an affine update in (h_k, u_k) with parameters (F_k, G_k, C_k, D_k) that can be read off the Mamba selective-scan code path. In our experiments, these are exactly the per-step quantities that the Mamba block computes during the forward pass.
- (ii) The token-to-token dependence of (F_k, G_k, C_k, D_k) on u_k is the discrete counterpart of the continuous-time selectivity of (A, B, C) on $x(t)$. The selection structure of the recurrence is preserved under sampling.

We take equation 16 as our discrete model of the Mamba selective-scan core and develop a passivity criterion on its parameters.

6.2 One-Step Dissipativity Inequality

Following the continuous-time framework, we ask whether equation 16 admits a quadratic storage that decays under the supply rate, sampled at the token grid. Let $P \in \mathbf{S}_+^N(\mathbb{C})$ be a positive-semidefinite matrix and define the sampled storage

$$V(h) = \frac{1}{2} h^\top P h.$$

A natural sampled analog of strict dissipativity is the requirement that, for every k and every (h_k, u_k) ,

$$V(h_{k+1}) - V(h_k) \leq u_k^\top y_k - \beta \|h_k\|^2, \quad (17)$$

for some $\beta \geq 0$. Inequality equation 17 is the discrete analog of the differential dissipativity bound used in the proof of Theorem 5.2.

Proposition 6.1 (Sampled block LMI for the Mamba selective-scan core). For the discrete recurrence equation 16 with quadratic storage $V(h) = \frac{1}{2}h^\top Ph$ and dissipation parameter $\beta \geq 0$, the one-step dissipativity inequality equation 17 holds for every (h_k, u_k) if and only if the per-step block matrix

$$\mathcal{L}_k(P, \beta) := \begin{bmatrix} F_k^\top P F_k - P + 2\beta I & F_k^\top P G_k - F_k^\top C_k^\top \\ G_k^\top P F_k - C_k F_k & G_k^\top P G_k - C_k G_k - G_k^\top C_k^\top - 2D_k \end{bmatrix} \preceq 0 \quad (18)$$

is negative semidefinite for that token. The matrix $\mathcal{L}_k(P, \beta)$ is differentiable in the model parameters that produce (F_k, G_k, C_k, D_k) and in P .

Proof Sketch. Substitute $h_{k+1} = F_k h_k + G_k u_k$ into $V(h_{k+1})$ and $y_k = C_k h_{k+1} + D_k u_k = C_k F_k h_k + (C_k G_k + D_k) u_k$ into the supply rate $u_k^\top y_k$. The inequality equation 17 becomes a quadratic form in the augmented vector $z_k = [h_k^\top, u_k^\top]^\top$:

$$z_k^\top \mathcal{L}_k(P, \beta) z_k \leq 0 \quad \text{for all } z_k,$$

which is equivalent to $\mathcal{L}_k(P, \beta) \preceq 0$. Differentiability in (F_k, G_k, C_k, D_k, P) is immediate from the polynomial structure of the entries. $\square$

Reading. Proposition 6.1 is exactly the discrete sampled instance of the parametric LMI in Theorem 5.2: the role of $Q(t)$ is played by the constant matrix P , the role of $\dot{Q}(t) + QA + A^\top Q$ is played by $F_k^\top P F_k - P$, and the off-diagonal $QB - C^\top$ is replaced by $F_k^\top P G_k - F_k^\top C_k^\top$. Forcing $\mathcal{L}_k(P, \beta) \preceq 0$ at every token is the discrete analog of forcing the parametric LMI to hold pointwise in x .

Scope of the criterion. A few cautious words about what equation 18 does and does not say. First, (F_k, G_k, C_k, D_k) are local linearizations of the Mamba selective-scan core at token k ; they are not the parameters of the full deep network with normalization, residual connections, and routing. Imposing $\mathcal{L}_k(P, \beta) \preceq 0$ is therefore a structural statement about the Mamba core, not a passivity certificate for the surrounding architecture. Second, in line with our coupled-input framing, equation 18 is a strong sufficient condition for the one-step dissipativity inequality equation 17: once the LMI holds at every k , the storage $V(h)$ is non-increasing along trajectories driven by $u_k \equiv x_k$ up to the supply term, regardless of the gating rule that produced (Δ_k, u_k) . Third, the matrix P is a design choice; the simplest case $P = I$ already gives a useful (Euclidean) energy meter, and we will see in Section 8 that learnable diagonal $P = \text{diag}(q)$ produces a small but consistent Pareto improvement.

6.3 Practical Mamba-Core regularizer

Proposition 6.1 suggests a clean training-time mechanism. Given a real Mamba forward pass on a batch of tokens, we can read off (F_k, G_k, C_k, D_k) at every step (they are exactly the per-token matrices used in the selective scan), choose a storage matrix P and a dissipation parameter $\beta \geq 0$, form the per-step block matrix $\mathcal{L}_k(P, \beta)$, and penalize its largest positive eigenvalue. Concretely, define

$$\mathcal{R}_{\text{LMI}}(\theta; P, \beta) := \mathbb{E}_k [\max(0, \lambda_{\max}(\mathcal{L}_k(P, \beta)))] , \quad (19)$$

where θ collects the trainable parameters that determine (F_k, G_k, C_k, D_k) in the Mamba core, and the expectation is taken over tokens (and optionally Mamba channels). The total training loss is then

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{L}_{\text{task}}(\theta) + \lambda_{\text{LMI}} \mathcal{R}_{\text{LMI}}(\theta; P, \beta),$$

with $\lambda_{\text{LMI}} \geq 0$ a single weight. We will refer to $\mathcal{R}_{\text{LMI}}$ as the *discrete LMI regularizer*.

A few practical remarks. Computing $\lambda_{\max}(\mathcal{L}_k(P, \beta))$ on small per-token blocks is cheap: in the architectures we consider in Section 8, $\mathcal{L}_k(P, \beta)$ is at most a few tens of dimensions, and its largest eigenvalue (with autograd-friendly subgradients) is computed via standard symmetric eigendecomposition. The choice of P controls how conservative the criterion is. We use $P = I$ as our main practical method and report a learnable

diagonal $P = \text{diag}(q)$ as an ablation; we found that the optional projection step that re-PSDifies P after each update is essentially inactive in our runs, so we do not present it as a central component. Finally, the regularizer is purely token-local at the level of the Mamba core; it does not require backpropagating through the entire sequence to evaluate, and its computational cost is negligible compared with the Mamba forward pass.

What this section does and does not assert. The reader should take three things from this bridge. (i) Proposition 6.1 is the discrete, sampled-data form of the continuous parametric LMI of Theorem 5.2, instantiated explicitly for the Mamba selective-scan recurrence. (ii) The criterion is a Mamba-core property: it constrains the per-step matrices that the Mamba block produces, not the surrounding deep network. (iii) The criterion is differentiable, so it can be used as a soft constraint during training. Whether enforcing it in this way actually changes how a real Mamba behaves — on clean test data, on input perturbations, and on routing-energy signals — is an empirical question, which we now answer in Section 8 on seven standard time-series datasets, with extended detail in Appendix C.

7 Robust Stability under Input-Driven Dynamics

Sections 4–5 laid out the structural side of selective SSMs — forgetting from strict dissipativity, the canonical quadratic structure of the frozen-selection subsystem, and the parametric-LMI consequences of universal quadratic storage. Section 6 then translated those structural results into a sampled, differentiable criterion for the Mamba selective-scan core. We now return to the continuous-time setting and analyze the selective SSM in the fully general and realistic case where it is driven by persistent, non-zero *port inputs* $u(\cdot)$ under an exogenous selection schedule $x(\cdot)$. This section provides sufficient conditions for Input-to-State Stability (ISS), the gold standard for robust stability in nonlinear systems. To achieve this strong guarantee, we must impose conditions that ensure the system’s internal dynamics are powerful enough to overcome any disturbance from the input. These conditions, while strong, provide a clear, verifiable pathway to designing certifiably robust models and are the continuous-time counterparts of the practical regularizer used in Section 8.

7.1 Global ISS from Uniform Dissipativity

To guarantee robust stability in the presence of arbitrary bounded inputs, basic passivity is not enough. We need to ensure the system’s internal dynamics are uniformly contractive, that is, they dissipate energy at a guaranteed rate, regardless of the operating mode selected by the *selection signal* $x(\cdot)$. This ensures the system can actively counteract the energy being injected by the *port input* $u(\cdot)$.

With this subsection, we present such a condition based on the existence of a common quadratic Lyapunov function demonstrating uniform dissipativity.

Theorem 7.1 (Global Stability from Uniform Local Dissipativity). Consider the continuous-time selective state-space model equation 2 under the standard regularity assumptions. Assume further that:

- (i) There exists a time-varying quadratic Lyapunov function candidate $V_Q(t, h) = \frac{1}{2}h^H Q(t)h$, where $Q : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$ satisfies:
 - $Q \in W_{\text{loc}}^{1,1}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$ (absolutely continuous locally).
 - $Q(t)$ is uniformly positive definite and bounded: there exist constants $k_2 \geq k_1 > 0$ such that $k_1 I \preceq Q(t) \preceq k_2 I$ for all $t \in \mathbb{T}$.
- (ii) The dynamics satisfy a uniform dissipativity condition with respect to V_Q : There exists a constant $\delta > 0$ such that for almost every $t \in \mathbb{T}$ and for **all** admissible *selection* values $x \in \mathbb{C}^{d_{\text{in}}}$ that can occur at time t , the following matrix inequality holds:

$$\dot{Q}(t) + Q(t)A(\Delta(t), x) + A(\Delta(t), x)^H Q(t) \preceq -2\delta Q(t) \quad (20)$$

(This implies that for any fixed selection value x , the homogeneous LTV system $\dot{h} = A(\Delta(t), x)h$ satisfies $\dot{V}_Q \leq -2\delta V_Q$ and hence is uniformly exponentially stable; equivalently, $\|h(t)\| \leq \sqrt{k_2/k_1} e^{-\delta(t-t_0)} \|h(t_0)\|$.)

- (iii) The input matrix $B(\Delta(t), x)$ is uniformly bounded: There exists a constant $M_B > 0$ such that $\|B(\Delta(t), x)\|_2 \leq M_B$ for all t and admissible x .

Then, the selective state-space system equation 2 is globally ISS with respect to the *port input* $u(t)$, uniformly over all admissible selection schedules $x(\cdot)$. Specifically, there exist constants $\tilde{C} \geq 1$, $\tilde{\gamma} > 0$, and a class $\mathcal{K}$ gain function σ such that for any initial state $h(t_0)$, any admissible selection schedule $x(\cdot)$, and any admissible port input $u(\cdot)$, the solution satisfies:

$$\|h(t)\| \leq \tilde{C} e^{-\tilde{\gamma}(t-t_0)} \|h(t_0)\| + \sigma(\|u\|_{[t_0, t], \infty}), \quad \|u\|_{[t_0, t], \infty} := \operatorname{ess\,sup}_{t_0 \leq \tau \leq t} \|u(\tau)\|. \quad \text{for all } t \geq t_0. \quad (21)$$

Proof Sketch. We analyze the time derivative of the Lyapunov function V_Q along the system trajectories. The derivative splits into two parts: one from the internal dynamics (A matrix) and one from the external *port input* (B matrix). The uniform dissipativity assumption provides a strong negative bound on the internal part ($\leq -2\delta V_Q$). The input part is bounded using the Cauchy-Schwarz inequality. Combining these results in a differential inequality of the form $\frac{dV}{dt} \leq -aV + b\sqrt{V}$. By a change of variables, $\Psi = \sqrt{V}$, this transforms into a standard linear differential inequality, $\frac{d\Psi}{dt} \leq -c\Psi + d$, whose solution is bounded. Translating this bound back to the state norm $\|h(t)\|$ yields the classic Input-to-State Stability (ISS) estimate. Check Appendix A.7 for the full proof. $\square$

ML implication. *Theorem 7.1 is the practically useful version of the structural results in Section 5. Where Theorem 5.2 characterized what passivity would have to enforce on a selective SSM under a single quadratic certificate, the ISS condition equation 20 is what robustness actually requires: a uniform exponential decay rate of the homogeneous dynamics, in matrix-inequality form, that holds across every selection value gating can produce. The discrete LMI regularizer of Section 6 can be read as a soft enforcement of exactly this kind of uniform-in- x matrix inequality on the Mamba selective-scan core: the regularizer drives $\lambda_{\max}(\mathcal{L}_k(P, \beta))$ toward zero across all tokens, which is the sampled analogue of asking equation 20 to hold for every selection value at almost every time. The empirical question this raises — does enforcing the discrete analog actually improve the model’s behavior under driving inputs and perturbations? — is what Section 8 is designed to answer.*

Remark 7.2 (Relationship to the Passivity LMI). The ISS condition in Eq. equation 20 is conceptually stronger than the passivity LMI from Theorem 5.2. The passivity LMI describes an **energy balance**, ensuring the system does not generate energy. In contrast, the ISS condition demands a forced **energy decay**, ensuring the system’s internal dynamics actively dissipate energy at a uniform exponential rate.

Formally, if the ISS condition holds (with $\delta > 0$ and a positive definite Q), it is sufficient to satisfy the internal stability portion (the (1,1) block) of the passivity LMI. However, a system can be passive (energy-balanced) without being uniformly contractive in this stricter sense. Thus, the ISS condition is a specialized and powerful tool for proving robust stability, while the passivity LMI is a more general tool for analyzing a system’s energy flow. In Section 8, we therefore use the more permissive passivity LMI (in its sampled form from Section 6) as the training-time regularizer, and treat the stronger ISS-style condition as the gold standard that the regularizer is empirically pushing the Mamba core toward.

8 Empirical Validation on SST/Mamba

The continuous-time analysis of Sections 4–7 and the discrete bridge of Section 6 make a concrete prediction: penalizing the largest positive eigenvalue of the sampled block matrix $\mathcal{L}_k(P, \beta)$ from equation 18 during training should constrain the Mamba selective-scan core toward sampled passivity, ideally without harming downstream forecasting accuracy. This section empirically evaluates that prediction across a real Mamba-based architecture on seven standard time-series forecasting datasets and four prediction horizons. Our goal

is not to set new MSE records on these benchmarks; it is to ask, with honest scoping, whether the discrete LMI regularizer in equation 19 actually does what the theory says it should, and what limits of the framework the experiments expose. Extended per-dataset tables, training hyperparameters, and additional ablations are deferred to Appendix C.

8.1 Experimental Setup

Architecture. We use the SST forecasting architecture (Xu et al., 2025), which is a routed mixture of a Mamba long-range expert and a short-range attention branch. We apply our regularizer only to the Mamba branch: at each token k , the per-step matrices (F_k, G_k, C_k, D_k) are extracted from the Mamba selective-scan core during the forward pass, the block matrix $\mathcal{L}_k(P, \beta)$ from equation 18 is formed, and the regularizer equation 19 is added to the standard task loss. Importantly, this means the regularizer is a Mamba-core diagnostic rather than a passivity certificate for the full SST network (with normalization, residual connections, and routing); we discuss this scoping carefully throughout.

Datasets and horizons. We evaluate on seven standard forecasting datasets: ETTh1, ETTh2, ETTm1, ETTm2, weather, electricity, and traffic, with prediction horizons $\{96, 192, 336, 720\}$, giving $7 \times 4 = 28$ dataset-horizon pairs. The look-back window is 672 tokens; the label length is 336.

Variants. We compare:

- **SST baseline.** The unregularized SST architecture.
- **SST + discrete LMI ($P = I$).** SST trained with $\mathcal{R}_{\text{LMI}}$ at $P = I$, $\lambda_{\text{LMI}}=10^{-2}$, $\beta=10^{-3}$. This is our main practical method.
- **SST + discrete LMI (diag- Q).** As above with $P = \text{diag}(q)$ a learnable diagonal storage (initialized at $q = 1$). Reported as a conservativeness ablation.

Additional ablations (a milder λ_{LMI} regime, a projection step, and a router-energy variant) appear in Sections 8.4–8.5 and Appendix C.

Diagnostics. At evaluation time we record clean test MSE and MAE, plus four Mamba-core diagnostics: the mean sampled LMI violation $\mathbb{E}_k[\max(0, \lambda_{\max}(\mathcal{L}_k))]$, the maximum sampled $\lambda_{\max}(\mathcal{L}_k)$, the mean Mamba-core SSM state norm, and (when relevant) the router’s mean Mamba weight m_{weight} .

8.2 Discrete LMI Regularization Reduces Mamba-Core Passivity Violations at No MSE Cost

We first evaluate, on clean test data, whether enforcing $\mathcal{L}_k(P, \beta) \preceq 0$ during training actually shifts the Mamba core toward sampled passivity. Table 3 summarizes the aggregate result across all 28 dataset-horizon pairs.

Variant	Clean MSE	Clean MAE	Mean LMI viol.	Max $\lambda_{\max}(\mathcal{L}_k)$
SST baseline	0.385804	0.390236	0.035824	2.6012
SST + discrete LMI ($P = I$)	0.385815	0.390244	0.002882	1.8972
SST + discrete LMI (diag- Q)	0.385815	0.390243	0.002877	1.8975

Table 3: Aggregate clean-test results across 28 dataset-horizon pairs, mean over runs. The discrete LMI regularizer reduces the sampled Mamba-core LMI violation by approximately 92% relative to baseline, while clean MSE differs by less than 0.003%.

The aggregate picture is unambiguous. Mean Mamba-core LMI violation drops from 0.0358 to 0.0029 under $P = I$ regularization, an approximately 92% reduction; the largest positive eigenvalue $\lambda_{\max}(\mathcal{L}_k)$ across all tokens and runs drops from 2.6012 to 1.8972. Clean MSE, by contrast, changes by approximately 0.0012% on

average, with a maximum per-run delta of 0.018%. Crucially, the violation reduction holds in 28/28 paired dataset-horizon comparisons, with the largest absolute reductions on the datasets where the unregularized baseline violates the criterion most (traffic, electricity), and the smallest on datasets where the baseline already nearly satisfies it (the ETTh family). Figure 1 shows this per-dataset structure.




Figure 1: Mean sampled Mamba-core LMI violation per dataset, averaged across the four prediction horizons. The discrete LMI regularizer reduces the violation on every dataset, with the largest reductions exactly where the unregularized baseline violates the sampled criterion most strongly.

We read this as the cleanest empirical statement of what Section 6 predicted: the regularizer does what it is designed to do, and the cost in forecasting accuracy is in the noise. Two notes on scoping. First, although $\lambda_{\max}(\mathcal{L}_k)$ is reduced uniformly, it is not driven below zero; the regularizer produces a substantial and consistent *reduction in violation*, not a certified passivity statement. Second, this is a Mamba-core property; the surrounding SST network (normalization, residuals, routing) is not constrained by the regularizer.

8.3 Conservativeness Ablation: $P = I$ vs Learnable Diagonal P

A natural reviewer concern is that $P = I$ might be unduly conservative as a storage choice. Section 6 explicitly noted that any $P \succeq 0$ is admissible, so we test whether replacing $P = I$ with a learnable diagonal $P = \text{diag}(q)$ buys anything. Aggregate results are in Table 4.

Variant	Clean MSE	Mean LMI viol.	q_{mean}	$q_{\text{log-std}}$
SST + discrete LMI ($P = I$)	0.385815	0.002882	—	—
SST + discrete LMI (diag- Q)	0.385815	0.002877	1.0008	1.1×10^{-3}

Table 4: Conservativeness ablation: replacing the storage $P = I$ with a learnable diagonal $P = \text{diag}(q)$ produces a strict (if small) Pareto improvement on 28/28 pairs at no measurable accuracy cost.

Across all 28 paired runs, the diagonal- Q variant has lower mean LMI violation than $P = I$, with no MSE degradation (paired-run violation reductions range from -0.10% to -0.33% ; MSE deltas in the noise floor at $\sim 10^{-5}$). Per-run reductions are strict: every paired run sits in the favourable corner of the (violation reduction, MSE delta) plane. The learned q moves visibly away from identity ($q_{\text{mean}} \approx 1.0008$, with $q_{\text{log-std}} \approx 1.1 \times 10^{-3}$ and $q_{\text{max}} \approx 1.0084$ on the dataset where it pushes hardest, traffic), confirming that the diagonal storage is being used. The improvement is modest in absolute terms; we read this as evidence that the framework can support less-conservative storage metrics without harming accuracy, and we use $P = I$ as the default practical choice while reporting diag- Q as the cleaner conservativeness check.

8.4 Robustness Diagnostics under Injected Perturbations

So far, our empirical claims concern clean test inputs. We next ask whether the regularizer changes the Mamba core’s behavior when inputs are corrupted. We use five synthetic perturbation modes applied only to the test inputs (target sequences are kept clean): Gaussian noise, spike injections, block masks, amplitude shifts, and last-value-hold; each variant is evaluated three times across seeds. Aggregate results, averaged across the five non-clean modes and all 28 dataset–horizon pairs, are in Table 5.

Variant	Robust MSE	MSE deg. %	Mean LMI viol.	Mean max state norm
SST baseline	0.70808	+126.76%	0.02345	0.26651
SST + discrete LMI ($P = I$)	0.70809	+126.76%	0.00512	0.25588
SST + discrete LMI (diag- Q)	0.70809	+126.76%	0.00512	0.25588

Table 5: Robustness diagnostics under five injected perturbation modes (Gaussian noise, spikes, block masks, amplitude shifts, last-value-hold), averaged across 28 dataset–horizon pairs and 3 seeds. The regularizer reduces internal Mamba-core LMI violations and state norms uniformly, but does not change robust forecasting MSE.

The result is sharply two-sided, and we report it as such. On internal Mamba-core diagnostics under perturbation, the regularizer is a strict Pareto improvement: mean LMI violation drops from 0.02345 to 0.00512 ($\sim 60\%$ reduction; lower in 35/35 dataset–mode pairs), mean max SSM state norm from 0.26651 to 0.25588 ($\sim 3.4\%$, lower in 35/35 pairs), and the worst observed state norm in the experiment from approximately 0.6065 to 0.5387. These improvements hold uniformly across all five perturbation modes, with the largest LMI-violation reductions on the hardest perturbation (amplitude shift, $\sim 345\%$ MSE degradation against clean for the baseline). On forecasting MSE under perturbation, by contrast, the regularizer produces no measurable improvement: robust MSE deltas vs. baseline are at +0.0008%, well within numerical noise. The Pareto picture under amplitude shift, the most challenging perturbation, is shown in Figure 2.




Figure 2: Pareto plot under amplitude perturbation, aggregated across datasets. Each point is one variant. The four LMI-regularized variants cluster strictly inside the favorable corner (lower mean Mamba-core LMI violation *and* lower mean max SSM state norm) of the (state norm, LMI violation) plane, while baseline and projection-only fall in the unfavorable corner.

We are intentionally precise about how we frame this. The regularizer improves internal Mamba stability and passivity diagnostics under perturbed inputs without degrading forecasting accuracy. It does not, in this run, produce a robust-MSE improvement; this is an honest negative on the practical task metric. Two plausible reasons are worth flagging. First, the regularizer is a Mamba-core diagnostic; the surrounding SST network (normalization, residual paths, routing, and output heads) absorbs much of the perturbation effect on MSE, so improving the Mamba core internally does not by itself translate into improved end-to-end MSE under corruption. Second, the perturbations are injected at test time only, so there is no opportunity for the regularized model to align its representation with the perturbed regime. We discuss both in Section 8.6.

8.5 Router-Stability Ablation: A Neutral Result

The SST router decides, per sample, how much to weight the Mamba branch versus the short-range branch. A natural conjecture is that adding an energy-aware penalty on the router — one that discourages high Mamba weight on samples with high local LMI violation — could induce a learned “avoid the Mamba branch when the Mamba core is locally unstable” rule, of the kind a human reading Section 7 would write. We tested this with a small additional regularizer on the router during training. Aggregate results over the same 28 dataset–horizon pairs are in Table 6.

Variant	Clean MSE	Mean m_w	Train Corr(m_w , viol.)	Mean LMI viol.
SST + discrete LMI ($P = I$)	0.385815	0.30801	–	0.00754
SST + discrete LMI + router	0.385815	0.30799	+0.0043	0.00754

Table 6: Router-stability ablation: adding the energy-aware router penalty preserves clean MSE in 28/28 runs and consistently nudges the Mamba weight downward, but the train-time correlation between Mamba weight and LMI violation is essentially zero on aggregate, with inconsistent sign across datasets (−0.107 on ETTh1 to +0.054 on ETTm1). We report this as a clean-performance neutrality result rather than as evidence of a learned energy-aware routing schedule. We use $m_w := m_{\text{weight}}$ for brevity.

The result is honest neutral. Clean MSE is preserved exactly within tolerance (28/28 within 1%), and the Mamba weight is lower than the LMI-only variant in 28/28 runs — directionally aligned with the energy-aware reading. But the magnitude of the Mamba-weight shift is on the order of 10^{-5} , the train-time correlation between Mamba weight and local LMI violation is +0.0043 on aggregate (essentially zero), and per-dataset correlations span both signs (from −0.107 on ETTh1 to +0.054 on ETTm1). We do not interpret this as evidence of a learned energy-aware schedule, and we present it as an appendix-style ablation rather than a positive contribution. The most plausible reason the mechanism is inert is that the router is trained on clean data, where local LMI violation is not a strong predictor of task loss. We expect that retraining with injected perturbations, so that violation magnitude becomes correlated with task loss, could produce a more consistent energy-aware schedule, but we leave this as future work.

8.6 Limitations

Three limitations of the empirical evaluation deserve explicit statement, and they should also be read as the limits of what the present paper claims.

Mamba core, not full network. The discrete LMI regularizer equation 19 is a per-token, per-channel constraint on the Mamba selective-scan recurrence equation 16. Sections 8.2–8.4 consequently report Mamba-core diagnostics, not passivity certificates for the full SST network. The mismatch between strong improvements in Mamba-core diagnostics and the absence of robust-MSE improvement (Section 8.4) is the empirical face of this scoping. Bridging this gap would require either a network-level passivity criterion (a research direction the parametric LMI of Section 5 suggests but does not yet provide) or an explicit cascade argument.

Sampled criterion, not exact discrete passivity. Proposition 6.1 gives an exact one-step dissipativity inequality at the Mamba selective-scan core; what we enforce during training is a soft penalty on the largest positive eigenvalue. Empirically, $\lambda_{\max}(\mathcal{L}_k)$ is reduced uniformly but not driven below zero. The numerical results should therefore be read as “substantial reduction in violation” rather than “certified sampled passivity”. We do not claim the latter.

Single architecture family. All experiments are on SST as the host architecture, a routed Mamba/short-range mixture. We chose this because it gives a clean Mamba branch on which to evaluate the criterion, and because the router lets us also probe the energy-aware routing question of Section 8.5. The criterion of Section 6 is, however, generic: it depends only on the per-token matrices (F_k, G_k, C_k, D_k) extracted from any selective-scan core. Validation on a wider set of selective architectures (HGRN, GLA, vanilla Mamba without the short-range branch, hybrid selective Transformers) is left as future work.

In summary, the discrete LMI regularizer is empirically clean-MSE safe, substantially reduces sampled Mamba-core LMI violations on every clean and perturbed run we tested, and improves Mamba-core stability diagnostics under input corruption. It does not, on its own, produce a robust-MSE improvement, and we explicitly do not claim it does. Section 9 discusses how these scoped empirical results sharpen the theoretical claims of Sections 4–7 and what they imply for selective-architecture design.

9 Conclusion

This paper has done two things. On the theoretical side, we developed a passivity- and ISS-based account of continuous-time selective State-Space Models with discontinuous gating, paying particular attention to the structural regularity of energy storage under the kind of L_{loc}^p coefficient functions that gating actually produces. On the practical side, we translated that account into a sampled, differentiable criterion on the Mamba selective-scan core and evaluated it as a training-time regularizer on a real Mamba-based forecasting architecture across 28 dataset–horizon pairs. We close by pulling together what was proved, what was empirically supported, and what was deliberately not claimed.

By separating the selection signal $x(\cdot)$ from the driving (port) input $u(\cdot)$, we obtained four structural results for system equation 2. State-strict dissipativity, together with quadratic bounds on a storage functional, yields exponential forgetting of homogeneous trajectories uniformly over admissible selection schedules (Theorem 4.1). Freezing the selection produces a passive linear time-varying input–output subsystem whose minimal available storage is necessarily quadratic with AUC_{loc} regularity (Theorem 4.2), giving a canonical intrinsic energy structure that is robust to discontinuous gating. Under the strong but informative hypothesis that a single quadratic storage certifies passivity uniformly over admissible selection schedules, we derived a parametric LMI in x together with universal kernel/unobservability and rank-monotonicity conditions, formalizing what we called *irreversible forgetting* (Theorems 5.1–5.4). Finally, a uniform-in- x matrix inequality on the homogeneous dynamics yields global ISS with respect to the port input (Theorem 7.1).

Section 6 translated the parametric LMI of Theorem 5.2 into a one-step dissipativity inequality on the Mamba selective-scan recurrence equation 16, with a sampled block matrix $\mathcal{L}_k(P, \beta)$ whose structure mirrors the continuous-time LMI exactly (Proposition 6.1). The largest positive eigenvalue of this matrix is differentiable in the model parameters that the Mamba selective scan produces, so it can be used directly as a soft training-time constraint equation 19. We are explicit that this criterion is a property of the Mamba core, not a passivity certificate for the surrounding deep network.

On clean test data across seven standard time-series datasets and four horizons, the discrete LMI regularizer reduces mean Mamba-core LMI violation by approximately 92% relative to the unregularized baseline (28/28 pairs), while clean MSE differs by less than 0.018% (Section 8.2). Replacing $P = I$ with a learnable diagonal $P = \text{diag}(q)$ produces a strict (if small) Pareto improvement in 28/28 paired runs at no measurable accuracy cost (Section 8.3), confirming that the framework is not overly tied to the simplest storage choice. Under five injected input perturbations, the regularizer produces a uniform $\sim 60\%$ reduction in mean Mamba-core LMI violation and a $\sim 3.4\%$ reduction in mean max SSM state norm, holding in 35/35 dataset–mode pairs and on every perturbation type (Section 8.4).

We have been deliberately precise about what the empirics do not say. Under perturbed inputs, the regularizer improves Mamba-core stability and passivity diagnostics but does not produce a robust forecasting MSE improvement on its own (Section 8.4); we report this as an honest negative on the practical task metric. The sampled LMI is reduced but not certified negative semidefinite, so we say “substantial reduction in violation” rather than “certified sampled passivity”. The router-stability ablation preserves clean MSE and consistently nudges Mamba weights downward but does not produce a learned energy-aware routing schedule, with train-time correlations between Mamba weight and LMI violation that are essentially zero on aggregate and inconsistent in sign across datasets (Section 8.5). The universal-storage hypothesis used in Section 5 is a strong sufficient condition that motivates the practical regularizer; it is not a claim that every Mamba in the wild satisfies it.

Read together, the theory and the experiments suggest three concrete design lessons for selective SSMs. First, principled initializations such as HiPPO (Gu et al., 2020) are best understood as the design of a

baseline energy landscape on which input-driven selection then operates: Theorem 4.2 explains why such initializations admit a well-conditioned quadratic energy function, and Appendix B.1 shows empirically that initializations failing this admit no valid energy function at all and diverge. Second, the sampled criterion of Section 6 is a low-cost, differentiable Mamba-core diagnostic that can be added to existing training pipelines with negligible overhead and provably (in the empirical sense of 28/28 paired runs) reduces sampled passivity violations without harming accuracy: it is, at minimum, a useful structural regularizer even if its end-to-end MSE benefit is not a forced consequence of passivity. Third, the structural framework of Sections 4–5 gives a precise control-theoretic vocabulary — forgetting, kernel inertness, irreversibility — in which to phrase questions about expressivity–stability tradeoffs in selective architectures, going beyond the Lyapunov-exponent diagnostics that have so far dominated the literature on Mamba-style stability (Halloran et al., 2024).

The honest scoping in Section 8.6 identifies three concrete next steps. (i) *From Mamba core to deep network.* The most informative gap revealed by our experiments is between Mamba-core diagnostic improvements and absent end-to-end MSE improvements under perturbation. Closing this gap requires either a deep-network passivity criterion (a research direction the parametric LMI of Theorem 5.2 suggests but does not yet resolve) or an explicit cascade argument bridging core dissipativity and end-to-end robustness. (ii) *Perturbation-aware energy routing.* The router-stability ablation suggests, plainly, that an energy-aware routing schedule does not emerge from training on clean inputs, because local LMI violation is not a strong predictor of clean-data task loss. We expect that retraining with injected perturbations, so that violation magnitude becomes correlated with task loss, is a path to a genuine learned routing rule of the kind a reader of Section 7 would write. (iii) *Discrete-time theorems.* The continuous-time results of Sections 4–7 have direct discrete-time analogues whose proofs we have left informal in Section 6; making the rank-monotonicity and irreversibility statements rigorous in the sampled regime, and connecting them to the empirically observed monotonic-reduction behaviour in Section 8, is the natural rigorous extension of the present work.

Selective SSMs are mathematically awkward: they are neither LTI, nor classical LTV, nor smooth nonlinear systems. We have argued that, because they are nonetheless port-Hamiltonian-style systems with a clear supply rate, they are amenable to a passivity-based analysis whose continuous-time conclusions transfer cleanly to the discrete Mamba selective scan as a per-token training-time criterion. Within the boundaries we have drawn, the criterion does what the theory predicts and reveals where the theory currently stops. We hope this combination of structural results and honest empirical scoping is a useful step toward a control-theoretic foundation for selective architectures in deep learning, and toward concrete tools that practitioners can drop into existing training loops without paying for them in accuracy.

References

- Brian DO Anderson and Sumeth Vongpanitlerd. *Network analysis and synthesis: a modern systems theory approach*. Courier Corporation, 2013.
- Michael S Branicky. Multiple lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on automatic control*, 43(4):475–482, 1998.
- Nicola Muca Cirone, Antonio Orvieto, Benjamin Walker, Cristopher Salvi, and Terry Lyons. Theoretical foundations of deep selective state-space models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=3SsrqwupUx>.
- Earl A Coddington, Norman Levinson, and T Teichmann. *Theory of ordinary differential equations*, 1956.
- AF Filippov. Existence and general properties of solutions of discontinuous systems. In *Differential Equations with Discontinuous Righthand Sides*, pp. 48–122. Springer, 1988.
- Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=tEYskw1VY2>.
- Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *Advances in Neural Information Processing Systems*, 33, 2020.
- Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. In *The International Conference on Learning Representations (ICLR)*, 2022.
- John T Halloran, Manbir Gulati, and Paul F Roysdon. Mamba state-space models are lyapunov-stable learners. *arXiv preprint arXiv:2406.00209*, 2024.
- Zhong-Ping Jiang, Eduardo D. Sontag, and Yuan Wang. Input-to-state stability for discrete-time nonlinear systems. *IFAC Proceedings Volumes*, 32:2403–2408, 1999. URL <https://api.semanticscholar.org/CorpusID:12463317>.
- Rudolf E Kalman. Lyapunov functions for the problem of lur’e in automatic control. *Proceedings of the national academy of sciences*, 49(2):201–205, 1963.
- Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural controlled differential equations for irregular time series. *Advances in neural information processing systems*, 33:6696–6707, 2020.
- Daniel Liberzon. *Switching in systems and control*, volume 190. Springer, 2003.
- Terry Lyons. Rough paths, signatures and the modelling of functions on streams. *arXiv preprint arXiv:1405.4537*, 2014.
- Riccardo Morandin and Dorothea Hinsin. Dissipative energy functionals of passive linear time-varying systems. In *ArXiv*, 2024. URL <https://api.semanticscholar.org/CorpusID:274981788>.
- Vasile-Mihai Popov. Absolute stability of nonlinear systems of automatic control. *Avtomatika i tele-mekhanika*, 22(8):961–979, 1961.
- Zhen Qin, Songlin Yang, Weixuan Sun, Xuyang Shen, Dong Li, Weigao Sun, and Yiran Zhong. Hgrn2: Gated linear rnns with state expansion. In *Proceedings of COLM*, 2024.
- Jimmy T.H. Smith, Andrew Warrington, and Scott Linderman. Simplified state space layers for sequence modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Ai8Hw3AXqks>.
- Eduardo Sontag and Yuan Wang. On characterizations of the input-to-state stability property. *Systems & Control Letters*, 24:351–359, 1995. URL <https://api.semanticscholar.org/CorpusID:13893551>.
- Eduardo D Sontag et al. Smooth stabilization implies coprime factorization. *IEEE transactions on automatic control*, 34(4):435–443, 1989.

- Taylan Soydan, Nikola Zubic, Nico Messikommer, Siddhartha Mishra, and Davide Scaramuzza. S7: Selective and simplified state space layers for sequence modeling. *ArXiv*, abs/2410.03464, 2024. URL <https://api.semanticscholar.org/CorpusID:273162247>.
- Arjan Van der Schaft. *L2-gain and passivity techniques in nonlinear control*. Springer, 2000.
- Jan C. Willems. Dissipative dynamical systems part i: General theory. *Archive for Rational Mechanics and Analysis*, 45:321–351, 1972. URL <https://api.semanticscholar.org/CorpusID:123076101>.
- Xiong Xiao Xu, Canyu Chen, Yueqing Liang, Baixiang Huang, Guangji Bai, Liang Zhao, and Kai Shu. Sst: Multi-scale hybrid mamba-transformer experts for time series forecasting. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pp. 3655–3665, 2025.
- Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. In *Proceedings of ICML*, 2024.
- Nikola Zubic, Federico Soldà, Aurelio Sulser, and Davide Scaramuzza. Limits of deep learning: Sequence modeling through the lens of complexity theory. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=DhdqML3FdM>.
- Nikola Zubić, Qian Li, Yuyi Wang, and Davide Scaramuzza. On the expressive power and limitations of multi-layer ssms. *arXiv preprint arXiv:2604.14501*, 2026.

A Theoretical Proofs

A.1 Proof of Theorem 4.1

We are given the existence of a storage function $V(t, h)$ satisfying the strict dissipativity inequality equation 5 and the quadratic bounds equation 6. The regularity assumption allows us to consider the differential version of the inequality. Along any trajectory of the system equation 2, the time derivative of V satisfies:

$$\frac{d}{dt}V(t, h(t)) \leq \operatorname{Re} \langle u(t), y(t) \rangle - \beta \|h(t)\|^2 \quad (22)$$

almost everywhere. This follows from differentiating the integral inequality or directly from the definition of dissipativity via differential supply rates. Now, consider the homogeneous (zero port-input) case, where $u(t) \equiv 0$ for all $t \geq 0$. Then the supply term vanishes:

$$\operatorname{Re} \langle u(t), y(t) \rangle = \operatorname{Re} \langle 0, y(t) \rangle = 0. \quad (23)$$

(As a special case, under frozen selection $x(t) \equiv 0$ this homogeneous trajectory coincides with setting $u \equiv 0$ in equation 10, i.e., $\dot{h} = A_0(t)h$ and $y_0 = C_0(t)h$.)

Substituting $u(t) = 0$ into the differential inequality equation 22, we get:

$$\frac{d}{dt}V(t, h(t)) \leq -\beta \|h(t)\|^2 \quad (24)$$

for trajectories in the homogeneous (zero port-input) case.

We can relate $\|h(t)\|^2$ back to $V(t, h(t))$ using the upper quadratic bound from equation 6: $\|h(t)\|^2 \geq \frac{1}{k_2}V(t, h(t))$. Substituting this into equation 24:

$$\frac{d}{dt}V(t, h(t)) \leq -\frac{\beta}{k_2}V(t, h(t)). \quad (25)$$

Let $\Phi(t) = V(t, h(t))$. This is a scalar non-negative function satisfying the differential inequality $\dot{\Phi}(t) \leq -\gamma'\Phi(t)$, where $\gamma' = \beta/k_2 > 0$. By the Comparison Lemma (a consequence of Grönwall's inequality), this implies:

$$\Phi(t) = V(t, h(t)) \leq V(0, h(0)) e^{-\gamma' t} \quad \text{for all } t \geq 0. \quad (26)$$

Now, we use both quadratic bounds from equation 6:

- Lower bound on $V(t, h(t))$: $k_1 \|h(t)\|^2 \leq V(t, h(t))$.
- Upper bound on $V(0, h(0))$: $V(0, h(0)) \leq k_2 \|h(0)\|^2$.

Combining these with the decay inequality equation 26:

$$k_1 \|h(t)\|^2 \leq V(t, h(t)) \leq V(0, h(0)) e^{-\gamma' t} \leq k_2 \|h(0)\|^2 e^{-\gamma' t}. \quad (27)$$

Dividing by $k_1 > 0$:

$$\|h(t)\|^2 \leq \frac{k_2}{k_1} \|h(0)\|^2 e^{-\gamma' t}. \quad (28)$$

Taking the square root of both sides:

$$\|h(t)\| \leq \sqrt{\frac{k_2}{k_1}} \|h(0)\| e^{-(\gamma'/2)t}. \quad (29)$$

Defining $C = \sqrt{k_2/k_1} \geq 1$ (since $k_2 \geq k_1 > 0$) and $\gamma = \gamma'/2 = \beta/(2k_2) > 0$, we obtain the desired exponential decay:

$$\|h(t)\| \leq C e^{-\gamma t} \|h(0)\|. \quad (30)$$

This completes the proof.

Summary: This theorem establishes that intrinsic energy dissipation ($\beta > 0$), when coupled with a storage function V that is quadratically comparable to the state norm, guarantees exponential stability of the homogeneous system ($u \equiv 0$), and in particular of the frozen-selection homogeneous dynamics ($u \equiv 0$, $x \equiv 0$). This signifies a fundamental "forgetting" capability, ensuring that the influence of the initial state diminishes exponentially over time in the absence of external port input, regardless of the specific (potentially non-quadratic) nature of V .

A.2 Available storage for passive LTV systems

We use the notion of *available storage* for LTV systems as defined in Morandin & Hinsén (2024, Def. 4.1): for an LTV system of the form $\dot{h} = A(t)h + B(t)u$, $y = C(t)h + D(t)u$, the available storage is

$$V_a(t_0, h_0) = \sup_{\substack{t_1 \geq t_0, u \in L^2_{\text{loc}} \\ h(t_0) = h_0}} \left(- \int_{t_0}^{t_1} \text{Re} \langle u(t), y(t) \rangle dt \right).$$

For passive LTV systems (in the sense of Morandin & Hinsén (2024, Def. 1.1)), the available storage is finite and coincides with a *minimal quadratic* storage $V_a = V_Q$ for some $Q \in \text{AUC}_{\text{loc}}$ (Morandin & Hinsén, 2024, Thm. 4.2, Cor. 4.3).

A.3 Proof of Theorem 4.2

We prove the three items (a)–(c).

(a) Passivity of the frozen-selection LTV system. Fix the selection signal $x(t) \equiv 0$ and define

$$A_0(t) := A(\Delta(t), 0), \quad B_0(t) := B(\Delta(t), 0), \quad C_0(t) := C(\Delta(t), 0), \quad D_0(t) \equiv 0.$$

By the regularity assumptions of the main text, $A_0 \in L^1_{\text{loc}}$, $B_0, C_0 \in L^2_{\text{loc}}$, and $D_0 \in L^\infty_{\text{loc}}$, hence the LTV system

$$\dot{h}(t) = A_0(t)h(t) + B_0(t)u(t), \quad y_0(t) = C_0(t)h(t) + D_0(t)u(t)$$

is of the form Morandin & Hinsén (2024, Eq. (1)). Let $(h(\cdot), u(\cdot), y_0(\cdot))$ be any state–input–output solution of this system on $[t_0, T]$. It is also a trajectory of the full selective SSM equation 2 when restricting to $x(\cdot) \equiv 0$. Therefore, applying equation 9 yields

$$V(T, h(T)) - V(t_0, h(t_0)) \leq \int_{t_0}^T \text{Re} \langle u(\tau), y_0(\tau) \rangle d\tau - \beta \int_{t_0}^T \|h(\tau)\|^2 d\tau \leq \int_{t_0}^T \text{Re} \langle u(\tau), y_0(\tau) \rangle d\tau.$$

Thus V satisfies the passivity dissipation inequality of Morandin & Hinsén (2024, Def. 1.1) for the frozen-selection LTV system, i.e., the LTV system is passive.

(b) Existence of minimal quadratic available storage. Let $V_{a,0}$ be the available storage of the frozen-selection LTV system in the sense of Morandin & Hinsén (2024, Def. 4.1). Since the system is passive, Morandin & Hinsén (2024, Cor. 4.3) implies that $V_{a,0}$ is finite and coincides with a quadratic form:

$$V_{a,0}(t, h) = \frac{1}{2} h^H Q_0(t) h$$

for some $Q_0 : \mathbb{T} \rightarrow \mathbf{S}_+^N(\mathbb{C})$, and that $V_{a,0}$ is minimal among all storage functions.

(c) AUC_{loc} regularity of $Q_0(t)$. The same result Morandin & Hinsén (2024, Cor. 4.3) yields that Q_0 can be chosen such that $Q_0 \in \text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$.

Connection to exponential decay. The homogeneous (unforced) trajectories correspond to setting $u \equiv 0$ in the frozen-selection LTV system, i.e., $\dot{h} = A_0(t)h$ (Morandin & Hinsén, 2024, Sec. 2.1). If V furthermore satisfies the hypotheses of Theorem 4.1, then Theorem 4.1 gives exponential decay along these homogeneous trajectories.

A.4 Proof of Theorem 5.1

(a) and (b): The premise is that $V_Q(t, h)$ satisfies the passivity inequality equation 11 for all admissible trajectories of the selective SSM equation 2. As argued in the proof of Theorem 4.2(a), any such trajectory includes the trajectories of the frozen-selection LTV system equation 10 obtained by setting $x(t) \equiv 0$. Therefore, V_Q must also be a valid storage function for the frozen-selection LTV system equation 10: indeed, since equation 11 holds for all admissible pairs $(x(\cdot), u(\cdot))$, it holds in particular for $x(\cdot) \equiv 0$ and for every admissible $u(\cdot)$, yielding the dissipation inequality for all trajectories of equation 10. Dropping the nonpositive term $-\beta \int_{t_0}^T \|h(\tau)\|^2 d\tau$ yields the passivity inequality of Morandin & Hinsén (2024, Def. 1.1).

The framework of Morandin & Hinsén (Morandin & Hinsén, 2024) analyzes quadratic storage functions for passive LTV systems under the assumed L_{loc}^p regularity. Specifically, (Morandin & Hinsén, 2024, Theorem 3.2) establishes necessary conditions for a quadratic form $V_Q(t, h) = \frac{1}{2}h^H Q(t)h$ to be a storage function. It implies that Q must be absolutely upper semi-continuous (which corresponds to AUC_{loc} in their terminology). By Morandin & Hinsén (2024, Thm. 3.2), any quadratic storage function $V_Q(t, h) = \frac{1}{2}h^H Q(t)h$ for a passive LTV system must be induced by a matrix function $Q \in \text{AUC}_{\text{loc}}(\mathbb{T}, \mathbf{S}_+^N(\mathbb{C}))$. Then Morandin & Hinsén (2024, Thm. 5.4(1)) yields that $t \mapsto \text{rank}(Q(t))$ is weakly decreasing on $\mathbb{T}$.

Summary: This theorem reveals that the strong requirement of universal passivity guaranteed by a single V_Q immediately imposes significant structure on $Q(t)$ itself. It must possess AUC_{loc} regularity, accommodating potential discontinuities but ensuring they behave in a controlled manner (e.g., jumps cannot increase energy in the Loewner sense). Crucially, the rank monotonicity implies that the dimension of the subspace captured by the energy function (its image) cannot increase over time. This lays the groundwork for understanding irreversible effects, as the "energy-less" subspace (the kernel) can only grow or stay the same.

A.5 Proof of Theorem 5.2

(a) **Parametric LMI Condition:** The dissipativity inequality equation 12 holding for all trajectories implies the following a.e. differential inequality along trajectories (at times where $\dot{Q}(t)$ exists): $\frac{d}{dt}V_Q(t, h(t)) \leq \text{Re} \langle u(t), y(t) \rangle - \beta \|h(t)\|^2$. Substituting $V_Q = \frac{1}{2}h^H Q h$, $\dot{h} = A(t, x)h + B(t, x)u$, and $y = C(t, x)h$ leads to a quadratic inequality in (h, u) , with the selection value x acting as a parameter:

$$\frac{1}{2}h^H \dot{Q} h + \text{Re}(h^H Q \dot{h}) \leq \text{Re}(u^H C(t, x)h) - \beta \|h\|^2. \quad (31)$$

$$\frac{1}{2}h^H \dot{Q} h + \text{Re}(h^H Q (A(t, x)h + B(t, x)u)) \leq \text{Re}(u^H C(t, x)h) - \beta \|h\|^2. \quad (32)$$

Rearranging terms:

$$\frac{1}{2}h^H \dot{Q} h + \text{Re}(h^H Q A(t, x)h) + \text{Re}(h^H Q B(t, x)u) - \text{Re}((C(t, x)^H u)^H h) + \beta h^H h \leq 0. \quad (33)$$

$$\frac{1}{2}h^H (\dot{Q} + Q A(t, x) + A(t, x)^H Q + 2\beta I)h + \text{Re}(h^H (Q B(t, x) - C(t, x)^H)u) \leq 0. \quad (34)$$

For each fixed admissible selection value x and for almost every time t , this inequality must hold for all $h \in \mathbb{C}^N$ and all port inputs $u \in \mathbb{C}^{d_{\text{in}}}$. This is precisely the condition encoded by the negative semidefiniteness of the LMI matrix $\mathcal{L}(t, x)$ defined in equation 13. Since V_Q must work for any input trajectory, the LMI must hold parametrically for all input values x that can occur at time t . This leads directly to the LMI formulation (setting $D = 0$):

$$\begin{bmatrix} \dot{Q}(t) + Q(t)A(t, x) + A(t, x)^H Q(t) + 2\beta I & Q(t)B(t, x) - C(t, x)^H \\ B(t, x)^H Q(t) - C(t, x) & 0 \end{bmatrix} \preceq 0 \quad (35)$$

(b) **Universal Kernel Condition for C :** Let $v \in \text{Ker}(Q(t))$ at a time t where the LMI equation 13 holds. Consider the augmented vector $z = \begin{bmatrix} v \\ w \end{bmatrix}$ for any $w \in \mathbb{C}^{d_{\text{in}}}$. The LMI implies $z^H \mathcal{L}(t, x) z \leq 0$.

$$\begin{aligned} z^H \mathcal{L}(t, x) z &= \begin{bmatrix} v^H & w^H \end{bmatrix} \begin{bmatrix} \dot{Q} + QA + A^H Q + 2\beta I & QB - C^H \\ B^H Q - C & 0 \end{bmatrix} \begin{bmatrix} v \\ w \end{bmatrix} \\ &= v^H (\dot{Q} + QA + A^H Q + 2\beta I) v + v^H (QB - C^H) w + w^H (B^H Q - C) v + w^H (0) w \end{aligned}$$

Since $v \in \text{Ker}(Q(t))$, we have $Q(t)v = 0$ and $v^H Q(t) = 0$. The expression simplifies to:

$$\begin{aligned} &= v^H (\dot{Q} + 0 + 0 + 2\beta I) v + v^H (0 - C^H) w + w^H (0 - C) v \\ &= v^H \dot{Q} v + 2\beta \|v\|^2 - v^H C^H w - w^H C v \\ &= v^H \dot{Q} v + 2\beta \|v\|^2 - 2 \text{Re}(w^H C v) \end{aligned}$$

This quadratic form in w must be ≤ 0 . Considering the 2×2 projection onto v and w space: $\begin{bmatrix} v^H (\dot{Q} + QA + A^H Q + 2\beta I) v & v^H (QB - C^H) w \\ w^H (B^H Q - C) v & 0 \end{bmatrix} = \begin{bmatrix} v^H \dot{Q} v + 2\beta \|v\|^2 & -v^H C(t, x)^H w \\ -w^H C(t, x) v & 0 \end{bmatrix} \preceq 0$. For this 2×2 matrix to be negative semidefinite, the diagonal elements must be non-positive (taking $w = 0$ shows $v^H (\dot{Q} + QA + A^H Q + 2\beta I) v \leq 0$, and if $v \in \text{Ker} Q(t)$ this reduces to $v^H \dot{Q}(t) v + 2\beta \|v\|^2 \leq 0$), and the determinant must be non-negative. The determinant is $(v^H \dot{Q} v + 2\beta \|v\|^2)(0) - |-w^H C(t, x) v|^2 = -|w^H C(t, x) v|^2$. So we need $-|w^H C(t, x) v|^2 \geq 0$. This can only hold if $|w^H C(t, x) v|^2 = 0$ for all $w \in \mathbb{C}^{d_{\text{in}}}$. This implies $C(t, x)v = 0$. Since this must hold for the $C(t, x)$ corresponding to any admissible x at time t , we conclude $\forall v \in \text{Ker}(Q(t)), C(\Delta(t), x)v = 0$ for all admissible x , a.e. t .

(c) **Implicit Constraints on A and B :** The validity of the LMI equation 13 for all x directly imposes constraints on $A(t, x)$ via the $(1, 1)$ block and couples $B(t, x)$ to $C(t, x)$ (which is already constrained by part (b)) via the off-diagonal blocks and $Q(t)$. These ensure that the dynamics generated by any input x remain compatible with the energy storage/dissipation defined by V_Q .

Summary: This theorem translates the abstract requirement of universal passivity into a concrete parametric LMI equation 13. This LMI must hold not just for a specific input or for the unforced system, but simultaneously for all possible selection values x that the gating mechanism might encounter at any given time t . A striking consequence is the universal kernel condition equation 14: any state direction v that is considered "energy-less" by the storage function (i.e., $v \in \text{Ker}(Q(t))$) must be rendered unobservable at the output ($C(t, x)v = 0$), irrespective of the specific input x driving the gating. This imposes a strong limitation on the gating mechanism: it cannot arbitrarily change the output matrix C in response to input x in a way that would make these energy-less states visible. Passivity demands consistency between the energy accounting (Q) and observability (C).

A.6 Proof of Theorem 5.4

1. **Kernel Non-Shrinking:** This follows directly from Theorem 5.1(b), which established that $\text{rank}(Q(t))$ is weakly monotonically non-increasing. Since $\dim(\text{Ker}(Q(t))) = N - \text{rank}(Q(t))$, a non-increasing rank implies a non-decreasing kernel dimension. Therefore $t \mapsto \dim \text{Ker}(Q(t))$ is weakly non-decreasing.

2. **Energy Consistency within the Kernel:** This condition was derived within the proof of Theorem 5.2(b) by evaluating the $(1, 1)$ block of the parametric LMI equation 13 for a vector $h(t) \in \text{Ker}(Q(t))$. The $(1, 1)$ block inequality is $\dot{Q}(t) + Q(t)A(t, x) + A(t, x)^H Q(t) + 2\beta I \preceq 0$. Pre- and post-multiplying by $h(t)^H$ and $h(t)$ respectively, and using $Q(t)h(t) = 0$, yields $h(t)^H \dot{Q}(t) h(t) + 2\beta \|h(t)\|^2 \leq 0$. This inequality must hold independently of x because it follows from the LMI which holds parametrically.

3. **Constraint on Dynamics Violating Kernel Structure:** If dynamics induced by some $x^*(t)$ were fundamentally incompatible with the passivity condition guaranteed by V_Q (e.g., by attempting to move a state out of the kernel in an "energy-creating" way relative to V_Q), it would necessarily cause the parametric LMI equation 13 to fail for $x = x^*(t)$. This contradicts the core assumption that V_Q ensures passivity.

universally. Therefore, all admissible dynamics under the universal V_Q assumption must inherently respect the energy balance constraints, including those within the kernel.

Summary: This theorem formalizes the notion of "irreversible forgetting" within the framework of a universal quadratic storage function. The non-shrinking kernel (Property 1) implies that once a state direction is deemed irrelevant or forgotten from an energy perspective (i.e., enters $\text{Ker}(Q(t))$), it structurally remains so according to the fixed energy measure $Q(t)$. The gating mechanism cannot manipulate $Q(t)$ to "un-forget" this mode. Furthermore, Property 2 ensures that the system dynamics and the evolution of $Q(t)$ itself must maintain energy consistency within this forgotten subspace, respecting the required dissipation rate β . Any gating strategy $x(t)$ must induce dynamics $(A(t, x), B(t, x), C(t, x))$ compatible with these constraints (Property 3). This provides a lens for analyzing robust memory properties: modes in $\text{Ker}(Q(t))$ are stably forgotten. It also offers potential insights into phenomena like catastrophic forgetting in learning systems; if learning adapts a $Q(t)$ -like structure, changes that violate kernel inertness or energy consistency might be necessary to learn new conflicting information, potentially disrupting the established passive structure.

A.7 Proof of Theorem 7.1

We analyze the time derivative of the Lyapunov function candidate $V_Q(t, h(t)) = \frac{1}{2}h(t)^H Q(t)h(t)$ along the trajectories of the full system equation 2. Since $Q \in W_{\text{loc}}^{1,1}$, its derivative $\dot{Q}(t)$ exists a.e. and lies in L_{loc}^1 . Using the chain rule and substituting the dynamics $\dot{h}(t) = A(\Delta(t), x(t))h(t) + B(\Delta(t), x(t))u(t)$, we obtain

$$\begin{aligned} \frac{d}{dt}V_Q(t, h(t)) &= \frac{1}{2}h^H \dot{Q}h + \frac{1}{2}\dot{h}^H Qh + \frac{1}{2}h^H Q\dot{h} \\ &= \frac{1}{2}h^H \dot{Q}h + \text{Re}(h^H Q\dot{h}) \\ &= \frac{1}{2}h^H \dot{Q}h + \text{Re}(h^H Q(A(t, x(t))h + B(t, x(t))u(t))) \\ &= \underbrace{\frac{1}{2}h^H (\dot{Q}(t) + Q(t)A(t, x(t)) + A(t, x(t))^H Q(t))h}_{\text{Homogeneous Part}} + \underbrace{\text{Re}(h^H Q(t)B(t, x(t))u(t))}_{\text{Input Part}}. \end{aligned}$$

Here, $A(t, x(t))$ stands for $A(\Delta(t), x(t))$, and similarly for B .

1. Homogeneous part. By the uniform dissipativity assumption equation 20, which holds for the specific selection value $x = x(t)$ occurring at time t , we have

$$\frac{1}{2}h^H (\dot{Q}(t) + Q(t)A(t, x(t)) + A(t, x(t))^H Q(t))h \leq \frac{1}{2}h^H (-2\delta Q(t))h = -\delta h^H Q(t)h = -2\delta V_Q(t, h(t)).$$

2. Input part. We bound the input term using $|\Re(z)| \leq |z|$ and Cauchy-Schwarz:

$$\begin{aligned} |\text{Re}(h^H Q(t)B(t, x(t))u(t))| &\leq |h^H Q(t)B(t, x(t))u(t)| \\ &\leq \|h(t)\| \|Q(t)\|_2 \|B(t, x(t))\|_2 \|u(t)\| \\ &\leq \|h(t)\| k_2 M_B \|u(t)\| \quad (\text{since } Q(t) \preceq k_2 I \text{ and } \|B(\Delta(t), x)\|_2 \leq M_B). \end{aligned}$$

Using $Q(t) \succeq k_1 I$, we have

$$V_Q(t, h(t)) = \frac{1}{2}h^H Q(t)h \geq \frac{k_1}{2}\|h(t)\|^2 \Rightarrow \|h(t)\| \leq \sqrt{\frac{2}{k_1}} \sqrt{V_Q(t, h(t))}.$$

Therefore,

$$|\text{Re}(h^H Q(t)B(t, x(t))u(t))| \leq \left(\sqrt{\frac{2}{k_1}} k_2 M_B\right) \|u(t)\| \sqrt{V_Q(t, h(t))}. \quad (36)$$

3. Differential inequality for V_Q . Combining the homogeneous and input bounds yields, for almost every t ,

$$\frac{d}{dt} V_Q(t, h(t)) \leq -2\delta V_Q(t, h(t)) + \left(\sqrt{\frac{2}{k_1}} k_2 M_B \right) \|u(t)\| \sqrt{V_Q(t, h(t))}. \quad (37)$$

Let $\Phi(t) = V_Q(t, h(t)) \geq 0$ and define $K := \sqrt{\frac{2}{k_1}} k_2 M_B$. Then

$$\dot{\Phi}(t) \leq -2\delta\Phi(t) + K\|u(t)\|\sqrt{\Phi(t)} \quad \text{for a.e. } t. \quad (38)$$

4. Comparison inequality via $\Psi = \sqrt{\Phi}$. Define $\Psi(t) = \sqrt{\Phi(t)}$. For almost every t such that $\Phi(t) > 0$,

$$\dot{\Psi}(t) = \frac{\dot{\Phi}(t)}{2\sqrt{\Phi(t)}} \leq \frac{-2\delta\Phi(t) + K\|u(t)\|\sqrt{\Phi(t)}}{2\sqrt{\Phi(t)}} = -\delta\Psi(t) + \frac{K}{2}\|u(t)\|.$$

Hence,

$$\dot{\Psi}(t) \leq -\delta\Psi(t) + \frac{K}{2}\|u(t)\| \quad \text{for a.e. } t. \quad (39)$$

By the comparison principle (integrating factor argument),

$$\Psi(t) \leq e^{-\delta(t-t_0)}\Psi(t_0) + \int_{t_0}^t e^{-\delta(t-\tau)} \frac{K}{2}\|u(\tau)\| d\tau. \quad (40)$$

Let

$$\|u\|_{[t_0, t], \infty} := \operatorname{ess\,sup}_{t_0 \leq \tau \leq t} \|u(\tau)\|.$$

Then

$$\begin{aligned} \int_{t_0}^t e^{-\delta(t-\tau)} \frac{K}{2}\|u(\tau)\| d\tau &\leq \frac{K}{2} \|u\|_{[t_0, t], \infty} \int_{t_0}^t e^{-\delta(t-\tau)} d\tau \\ &= \frac{K}{2} \|u\|_{[t_0, t], \infty} \cdot \frac{1 - e^{-\delta(t-t_0)}}{\delta} \\ &\leq \frac{K}{2\delta} \|u\|_{[t_0, t], \infty}. \end{aligned}$$

Therefore,

$$\Psi(t) \leq e^{-\delta(t-t_0)}\Psi(t_0) + \frac{K}{2\delta} \|u\|_{[t_0, t], \infty}. \quad (41)$$

5. Convert back to $\|h(t)\|$. Substituting $\Psi = \sqrt{V_Q}$ and using $k_1 I \preceq Q(t) \preceq k_2 I$ (for all t),

$$\sqrt{V_Q(t, h(t))} \geq \sqrt{\frac{k_1}{2}} \|h(t)\|, \quad \sqrt{V_Q(t_0, h(t_0))} \leq \sqrt{\frac{k_2}{2}} \|h(t_0)\|.$$

Thus,

$$\begin{aligned} \sqrt{\frac{k_1}{2}} \|h(t)\| &\leq \sqrt{V_Q(t, h(t))} \\ &\leq e^{-\delta(t-t_0)} \sqrt{V_Q(t_0, h(t_0))} + \frac{K}{2\delta} \|u\|_{[t_0, t], \infty} \\ &\leq e^{-\delta(t-t_0)} \sqrt{\frac{k_2}{2}} \|h(t_0)\| + \frac{K}{2\delta} \|u\|_{[t_0, t], \infty}. \end{aligned}$$

Dividing by $\sqrt{k_1/2}$ gives

$$\|h(t)\| \leq \sqrt{\frac{k_2}{k_1}} e^{-\delta(t-t_0)} \|h(t_0)\| + \underbrace{\frac{k_2 M_B}{\delta k_1}}_{=: K'} \|u\|_{[t_0, t], \infty}. \quad (42)$$

This is exactly the ISS estimate equation 21 with $\tilde{C} = \sqrt{k_2/k_1}$, $\tilde{\gamma} = \delta$, and $\sigma(s) = K's$.

Summary: Under a common quadratic Lyapunov function V_Q with uniform bounds and a uniform decay rate equation 20 holding for all admissible selection values x (hence for any selection schedule $x(\cdot)$), the selective SSM is ISS with respect to the port input $u(\cdot)$, with a linear gain depending on k_1, k_2, δ and the uniform bound M_B on $B(\Delta(t), x)$.

A.8 Explanation of Comparison Principle and derivation of Eq. (40)

We start from the differential inequality derived for $\Psi(t)$:

$$\dot{\Psi}(t) \leq -\delta\Psi(t) + \frac{K}{2}\|u(t)\| \quad (43)$$

which can be rewritten as

$$\dot{\Psi}(t) + \delta\Psi(t) \leq \frac{K}{2}\|u(t)\| \quad (44)$$

where $\delta > 0$. This is a standard first-order linear differential inequality. Its solution bound can be obtained either by the explicit method of integrating factors or by invoking a suitable comparison lemma. Both approaches lead to the same upper bound, namely Eq. equation 40.

Method 1: Integrating factor (standard ODE technique). Multiply both sides of equation 44 by the integrating factor $e^{\delta t}$:

$$e^{\delta t}\dot{\Psi}(t) + \delta e^{\delta t}\Psi(t) \leq e^{\delta t}\frac{K}{2}\|u(t)\|. \quad (45)$$

Recognize the left-hand side as the derivative of a product:

$$\frac{d}{dt}(e^{\delta t}\Psi(t)) = e^{\delta t}\dot{\Psi}(t) + \delta e^{\delta t}\Psi(t). \quad (46)$$

Hence,

$$\frac{d}{dt}(e^{\delta t}\Psi(t)) \leq e^{\delta t}\frac{K}{2}\|u(t)\|. \quad (47)$$

Integrating both sides from t_0 to t (using τ as the integration variable) gives

$$\int_{t_0}^t \frac{d}{d\tau}(e^{\delta\tau}\Psi(\tau)) d\tau \leq \int_{t_0}^t e^{\delta\tau}\frac{K}{2}\|u(\tau)\| d\tau. \quad (48)$$

By the Fundamental Theorem of Calculus,

$$e^{\delta t}\Psi(t) - e^{\delta t_0}\Psi(t_0) \leq \int_{t_0}^t e^{\delta\tau}\frac{K}{2}\|u(\tau)\| d\tau. \quad (49)$$

Solving for $\Psi(t)$ yields

$$\begin{aligned} \Psi(t) &\leq e^{-\delta(t-t_0)}\Psi(t_0) + e^{-\delta t} \int_{t_0}^t e^{\delta\tau}\frac{K}{2}\|u(\tau)\| d\tau \\ &= e^{-\delta(t-t_0)}\Psi(t_0) + \int_{t_0}^t e^{-\delta(t-\tau)}\frac{K}{2}\|u(\tau)\| d\tau, \end{aligned} \quad (50)$$

which is exactly Eq. equation 40.

Method 2: Comparison lemma (a Grönwall-type argument). A common comparison lemma states the following: if $\eta(\cdot)$ is absolutely continuous and satisfies

$$\dot{\eta}(t) \leq a(t)\eta(t) + b(t) \quad \text{a.e. on } [t_0, t],$$

and if $\zeta(\cdot)$ solves the corresponding equality

$$\dot{\zeta}(t) = a(t)\zeta(t) + b(t), \quad \zeta(t_0) = \eta(t_0),$$

then $\eta(t) \leq \zeta(t)$ for all $t \geq t_0$.

In our case, equation 43 has $a(t) = -\delta$ and $b(t) = \frac{K}{2}\|u(t)\|$. The solution of the equality $\dot{\zeta} = -\delta\zeta + \frac{K}{2}\|u(t)\|$ with $\zeta(t_0) = \Psi(t_0)$ is given by variation of constants:

$$\zeta(t) = e^{-\delta(t-t_0)}\Psi(t_0) + \int_{t_0}^t e^{-\delta(t-\tau)} \frac{K}{2}\|u(\tau)\| d\tau. \quad (51)$$

By the comparison lemma, $\Psi(t) \leq \zeta(t)$, which again yields Eq. equation 40.

B Toy Simulations Illustrating the Continuous-Time Theory

This appendix presents three small simulation studies that illustrate the structural results of Sections 4–5. The runs here are deliberately small-scale: they use synthetic dynamics rather than real Mamba selective scans, and their purpose is pedagogical, not benchmark-setting. The full empirical evaluation of the discrete LMI regularizer on a real Mamba-based forecasting architecture appears in Section 8, with extended experimental details in Appendix C.

B.1 Initialization’s Decisive Impact on the Energy Landscape

Objective. This experiment provides a concrete demonstration of the core claims in Section 4: that a model’s initialization directly governs its intrinsic stability and memory-vs-stability trade-off. It illustrates Theorem 4.2 by exhibiting cases where a quadratic energy function $V_{a,0}(h) = \frac{1}{2}h^T Q_0 h$ exists and behaves as the theory predicts, and a control case where it does not.

Setup. We compared three initialization strategies for the frozen-selection state matrix $A_0(t) = A(\Delta(t), 0)$ of an 8-dimensional SSM ($N = 8$).

1. **HiPPO-style initialization (A_H):** Eigenvalues are generated to be stable (negative real parts) but clustered near the imaginary axis, with the rightmost (slowest) eigenvalue at $\text{Re}(\lambda) \approx -0.1$.
2. **Random Stable initialization (A_S):** Eigenvalues are randomly generated but constrained to be in the left half-plane, resulting in a wider spectral spread. The rightmost eigenvalue is at $\text{Re}(\lambda) \approx -0.52$.
3. **Random Unstable initialization (A_U):** Eigenvalues are drawn from the same distribution as A_S but without the stability constraint, resulting in a dominant eigenvalue at $\text{Re}(\lambda) \approx +0.79$.

For each A_0 , we first attempted to find a corresponding quadratic energy matrix Q_0 by solving the continuous-time Lyapunov equation $A_0^T Q_0 + Q_0 A_0 = -I$. We then simulated the frozen-selection system ($x(t) \equiv 0$) by feeding it a white-noise *port input* $u(t)$ for 5 seconds to populate its state, after which the port input was turned off ($u(t) \equiv 0$) to observe the free energy decay.

Results & Interpretation. Our findings, summarized in Figure 3, provide strong empirical validation for our theory:

- **Existence of a Valid Energy Function (Q_0):** Both the *HiPPO-style* and *Random Stable* initializations admitted a unique, positive-definite solution Q_0 to the Lyapunov equation (with condition numbers of ≈ 21 and ≈ 13 , respectively). The *Random Unstable* initialization failed to produce a valid solution, yielding an indefinite matrix, confirming that no quadratic energy function exists for an unstable system.

- **Spectral Gap Governs Memory:** The slowest decay rate is dictated by the rightmost eigenvalue. The spectral gap for HiPPO ($\approx -0.1 \text{ s}^{-1}$) was five times smaller than for Random Stable ($\approx -0.52 \text{ s}^{-1}$). Consequently, after the input was removed, the state norm of the HiPPO-initialized model took ≈ 23 seconds to decay by two orders of magnitude, while the Random Stable model did so in just ≈ 4 seconds. The Unstable model diverged exponentially.
- **Energy Trajectories Confirm Theory:** The energy $V_{a,0}(t) = \frac{1}{2}h(t)^T Q_0 h(t)$ for the HiPPO model decayed almost perfectly linearly on a log-scale with a slope of ≈ -0.1 , illustrating a long memory horizon. The Random Stable model's energy decayed five times faster.

This experiment confirms that merely being stable is insufficient. The precise spectral placement of A_0 , as achieved by principled initializations like HiPPO, directly governs the memory-vs-stability trade-off, exactly as our energy-based theoretical framework predicts.




Figure 3: Energy decay $V_{a,0}(t)$ on a log-scale for an 8-D SSM with three different initializations after the port input $u(t)$ is removed (set to zero) at $t = 5\text{s}$. The HiPPO-style initialization yields a valid energy function and a slow decay rate (long memory). The Random Stable initialization also has a valid energy function but decays much faster (short memory). The Unstable initialization has no valid energy function and diverges.

B.2 Visualizing Irreversible Forgetting via Rank-Deficient Gating

Objective. This experiment is designed to provide a concrete, visual demonstration of the "irreversible forgetting" concept introduced in Section 5. It illustrates Theorem 5.1 (rank monotonicity of any universal quadratic storage matrix) and Theorem 5.4 (kernel inertness) on a small synthetic switched system.

Setup. We construct a 3-dimensional linear system $\dot{h} = A(t)h$ that can be switched between two operating modes via a gating signal. Each mode is defined by a state matrix A_i and an associated minimal storage matrix Q_i . The storage matrices are the unique positive semidefinite solutions to the continuous-time Lyapunov equation $A_i^T Q_i + Q_i A_i = -C_i^T C_i$, which defines the energy landscape for an observable system. Modes of the system are:

- **Mode 1 (Full-Rank Storage):** The system dynamics are governed by $A_1 = \text{diag}[-0.2, -0.3, -0.4]$. We choose $C_1 = I_3$, representing full observability of the state. The re-

sulting storage matrix Q_1 is full-rank (rank=3), meaning the system can store energy in all three state dimensions.

- **Mode 2 (Rank-Deficient Storage):** The dynamics are governed by $A_2 = \text{diag}[-0.2, -0.3, -15]$. The large negative eigenvalue is designed to rapidly dissipate the third state component. We choose $C_2 = \text{diag}[1, 1, 0]$, making the third state dimension unobservable. The resulting storage matrix Q_2 is rank-deficient (rank=2), with the third dimension lying in its kernel ($\text{Ker}(Q_2)$). In this mode, the system loses the capacity to store energy in the third dimension.

Simulation Protocol. We simulate the homogeneous system (no input channel) from a random initial state with non-zero components in all dimensions. The gating signal switches the system's mode according to the following timeline:

- **0s $\leq t < 5$ s:** The system operates in Mode 1 (full-rank energy storage).
- **5s $\leq t < 10$ s:** The gating switches the system to Mode 2 (rank-deficient storage).
- **$t \geq 10$ s:** The gating attempts to switch the system back to Mode 1.

Theoretical Prediction & Results. Our theory predicts that if a single, universal quadratic storage function $V_Q(t)$ governs the entire trajectory, its defining matrix $Q(t)$ must have a non-increasing rank. More precisely, in this case:

- The switch at $t = 5$ s is permissible, as the storage rank can decrease from 3 to 2.
- The attempted switch back at $t = 10$ s, however, would necessitate an increase in the storage rank from 2 to 3. This violates the rank monotonicity property inherent to AUC_{loc} functions.

This implies that no single $V_Q(t)$ can certify passivity for the entire trajectory. The "forgetting" of the energy storage capacity in the third dimension is, from the perspective of a single passive structure, irreversible.

The simulation results, visualized in Figure 4, confirm this prediction empirically. During Mode 2, the third state component (h_3) rapidly collapses to zero and, crucially, does not recover even after the system dynamics are switched back to Mode 1. The system becomes permanently confined to the two-dimensional subspace. This provides a clear, empirical illustration of our theoretical claims: once the system's energy storage rank collapses due to gating, it cannot be increased again without violating the fundamental constraints required for a stable, passive system.

B.3 LMI Regularization for Improved Training Robustness on a Toy System

Objective. This experiment validates, on a small synthetic system, the central practical claim that motivated the discrete-time bridge of Section 6: that the continuous-time LMI condition from Theorem 5.2 can be used as a regularizer to train more robust selective SSMs. The full evaluation of the discrete-time analogue of this regularizer on a real Mamba-based architecture is in Section 8.

Setup.

- **Task:** We designed a challenging 1D tracking task where the SSM's output $y(t)$ must track a "spiky" reference signal $r(t)$ composed of steps and sharp impulses, designed to provoke large internal state excursions.
- **Model:** We use a 2-dimensional ($N = 2$) selective SSM. We use the coupled case $u(t) \equiv x(t)$, so the same signal both selects parameters and drives the dynamics. The state matrix is explicitly selection-dependent to ensure selectivity: $A(x) = A_{\text{base}} + \tanh(x) \cdot A_{\text{sel}}$, where A_{base} and A_{sel} are learned 2×2 matrices. The input matrix B and output matrix C are also learned.




Figure 4: Illustration of irreversible forgetting. The rank of the underlying minimal storage matrix is plotted over time (or implied by state trajectories). The system switches from a full-rank mode (rank=3) to a rank-deficient mode (rank=2) at $t = 5$ s, causing the third state component to collapse. Our theory proves that any attempt to switch back in a way that increases the rank of a universal storage function (e.g., the dashed red line) is incompatible with maintaining a single, passive energy structure, demonstrating the irreversibility principle from Theorem 5.1.

- **LMI Regularizer:** We employ the LMI from Theorem 5.2. For simplicity and to demonstrate the core principle, we fix the storage matrix $Q = I$. The regularization loss is defined as the magnitude of the LMI violation: $\mathcal{L}_{\text{LMI}} = \max(0, \lambda_{\max}(\mathcal{L}(x)))$, where $\mathcal{L}(x)$ is the LMI matrix evaluated for the selection value x (equivalently for the port input since $u \equiv x$ in this experiment).
- **Two Training Conditions:** We compare two models: 1. **Baseline Model:** Trained only with the task loss, $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}}$, where $\mathcal{L}_{\text{task}} = \text{MSE}(y(t), r(t))$. 2. **LMI-Regularized Model:** Trained with the combined loss, $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \gamma \cdot \mathcal{L}_{\text{LMI}}$, using a small regularization weight $\gamma = 0.01$.

Results and Interpretation. The results of the experiment, evaluated on a held-out test set, are summarized in Table 7.

Table 7: Comparison of Baseline vs. LMI-Regularized SSM on a challenging tracking task.

Model	Task MSE (Test)	Max State Norm $\ h\ _{\infty}$	Max LMI Violation
Baseline Model	0.073	18.4	3.51
LMI-Regularized Model	0.081	1.9	0.003

The results lead to three clear conclusions:

1. **Robustness is Dramatically Improved:** The most interesting result is in the maximum state norm. The baseline model, while achieving a slightly better task MSE, does so by allowing its

internal state to reach a very large norm (18.4). This is indicative of a model operating on the edge of instability, vulnerable to exploding states. In contrast, the **LMI-regularized model's state norm is an order of magnitude smaller (1.9)**. It has learned to perform the task while keeping its internal dynamics constrained and provably more stable.

2. **The Regularizer Works as Intended:** The "Max LMI Violation" column shows that the regularizer was highly effective. The baseline model frequently and significantly violates the passivity condition (max violation of 3.51). The regularized model has learned parameters that keep the LMI matrix negative semidefinite (max violation is near zero), successfully enforcing the theoretical condition for stability derived in our work.
3. **Minimal Impact on Task Performance:** Crucially, this significant gain in robustness comes at a negligible cost to task performance. The MSE of the regularized model is only marginally higher than the baseline, demonstrating that a stable, well-behaved solution exists that is also effective for the task.

To conclude, this experiment demonstrates on a small toy system that our theoretical condition can be directly translated into a practical tool for improving training robustness. The discrete-time analogue of this regularizer, applied to a real Mamba selective-scan core, is evaluated systematically in Section 8.

C Extended Mamba/SST Experimental Details

This appendix collects the experimental details that did not fit in Section 8. We give the full training and architecture configuration in Appendix C.1, additional per-dataset and per-run breakdowns of the headline discrete LMI result in Appendix C.2, the milder LMI-weight regime that complements the conservativeness ablation in Appendix C.3, the full per-mode breakdown of the robustness study in Appendix C.4, and the per-dataset breakdown of the router-stability ablation in Appendix C.5. Aggregate analysis CSVs and additional plots are released alongside the paper.

C.1 Architecture, Datasets, and Training Configuration

Architecture. We use SST (Xu et al., 2025) as the host architecture. SST routes a Mamba long-range expert and a short-range branch via a learned router; we apply our discrete LMI regularizer only to the Mamba branch. For the smaller datasets (ETTh1, ETTh2, ETTm1, ETTm2, weather), Mamba is configured with model dimension $d_m = 16$, two heads, $d_f = 64$. For the larger datasets (electricity, traffic), $d_m = 32$, four heads, $d_f = 128$. All other architectural details follow the SST defaults.

Datasets and horizons. We use the seven public time-series benchmarks ETTh1, ETTh2, ETTm1, ETTm2, weather, electricity, and traffic, with the standard train/val/test splits. The look-back window is `seq_len = 672`, the label length is `label_len = 336`, and the prediction horizons are `pred_len` $\in \{96, 192, 336, 720\}$, giving the $7 \times 4 = 28$ dataset-horizon pairs reported throughout Section 8.

Training. All variants use the SST default optimizer, learning-rate schedule, and number of epochs. The discrete LMI regularizer of equation 19 is added to the standard task loss with weight $\lambda_{\text{LMI}} \in \{10^{-3}, 10^{-2}\}$ depending on the regime: $\lambda_{\text{LMI}} = 10^{-3}$ is the milder weight studied in Appendix C.3; $\lambda_{\text{LMI}} = 10^{-2}$ is the headline weight used in Sections 8.2–8.3, with $\beta = 10^{-3}$ and $P \in \{I, \text{diag}(q)\}$. The optional projection step that re-PSDifies P after each update was effectively non-binding in our runs (yielding outputs visually identical to the non-projection variant on every diagnostic), so we report results for it only briefly in Appendix C.4 and do not present it as a central component. Each variant is trained with three seeds; the aggregate numbers in Section 8 are means over seeds.

C.2 Discrete LMI Regularization: Per-Run Breakdown

The aggregate result of Section 8.2 is that the discrete LMI regularizer reduces mean Mamba-core LMI violation by approximately 92% (28/28 pairs) while clean MSE differs by less than 0.018% per run. The

per-run picture, in the form of paired MSE deltas vs. the unregularized baseline, is shown in Figure 5. The per-dataset breakdown of how strongly $\lambda_{\max}(\mathcal{L}_k)$ is reduced is shown in Figure 6, complementing the mean-violation plot in the main paper.




Figure 5: Per-run clean-test MSE delta vs. baseline of the LMI-regularized SST, sorted by dataset and horizon. Across 28/28 paired runs, the MSE delta is below 0.018%, with most deltas under 0.005%. The largest delta appears at **traffic**, $\text{pred_len} = 720$, the regime where the regularizer pushes hardest on the Mamba core.




Figure 6: Per-dataset reduction in $\lambda_{\max}(\mathcal{L}_k)$ between baseline SST and the LMI-regularized variant, averaged across the four prediction horizons. The reduction is strict on every dataset; the residual positive value of $\lambda_{\max}$ is the basis for the “substantial reduction in violation, not certified passivity” framing in Section 8.6.

C.3 Mild LMI-Weight Regime

The conservativeness ablation in Section 8.3 reported the strong LMI-weight regime $\lambda_{\text{LMI}} = 10^{-2}$. Here we report the same comparison with the milder weight $\lambda_{\text{LMI}} = 10^{-3}$, which probes how much of the conservative-

ness story depends on the strength of the regularizer. Aggregate results across the same 28 dataset–horizon pairs are summarized in Table 8.

Variant	Clean MSE	Mean LMI viol.	q_{mean}	$q_{\text{log-std}}$
SST baseline	0.385804	0.035824	–	–
SST + discrete LMI ($P = I$, mild)	0.385809	0.007541	–	–
SST + discrete LMI (diag- Q , mild)	0.385809	0.007537	1.0002	3.8×10^{-4}

Table 8: Mild-weight ($\lambda_{\text{LMI}} = 10^{-3}$) conservativeness ablation. The regularizer reduces mean Mamba-core LMI violation by approximately 79% in this regime, and diag- Q remains a strict (if smaller) Pareto improvement over $P = I$ (28/28). The learned q stays much closer to identity than under the strong weight, a useful sanity check that the diagonal mechanism is responsive to the regularizer strength.

The conservativeness conclusion of Section 8.3 carries through. Under the milder regime, the relative gap between $P = I$ and diag- Q is smaller (paired-run violation reductions average -0.05% vs. the -0.22% of the strong regime) and the learned q moves visibly less from identity. This is the expected sensitivity to the regularizer strength: when the LMI penalty is small, the optimizer does not need diagonal storage flexibility to satisfy it, and accordingly q stays close to its initialization.

C.4 Robustness Diagnostics: Per-Mode Breakdown

The aggregate robustness result of Section 8.4 averages across five injected perturbation modes. The corresponding state-norm picture is in Figure 7. The robust MSE picture, with all six variants superimposed, is in Figure 8.




Figure 7: Mean max SSM state norm per dataset and per perturbation mode. The state-norm reduction is small in absolute terms (mean -3.4% , peaking at -11% on the worst-state-norm runs) but holds in 35/35 pairs, providing a complementary internal diagnostic to the LMI-violation picture.

On the projection variant. The projection-only variant (no LMI regularizer, just the post-update projection) is statistically indistinguishable from the unregularized baseline across all robustness diagnostics in our runs. The LMI-with-projection variants (both $P = I$ and diag- Q) are statistically indistinguishable from their non-projection counterparts. We interpret this as evidence that the projection step is not binding in our setting and accordingly do not promote it as a meaningful design choice.




Figure 8: Robust forecasting MSE per dataset and per perturbation mode. All six variants (baseline, projection-only, $P = I$ LMI, $P = I + \text{projection}$, diag- Q , diag- $Q + \text{projection}$) overlap visibly on every cell. This is the visual face of the “no robust-MSE improvement” finding discussed in Section 8.4: the regularizer does not change end-to-end task error under perturbation, even though it changes Mamba-core internal diagnostics consistently.

C.5 Router-Stability Ablation: Per-Dataset Detail

The aggregate result of Section 8.5 is that adding the router-stability penalty preserves clean MSE in 28/28 runs and consistently nudges the Mamba branch downward in 28/28 runs, but the train-time correlation between Mamba weight and LMI violation is essentially zero on aggregate (+0.0043). The per-dataset breakdown of this train-time correlation is in Figure 9; per-run deltas in m_{weight} vs. the LMI-only variant are in Figure 10.




Figure 9: Per-dataset train-time correlation between the router’s Mamba weight and the local LMI violation, under the energy-aware router penalty. The signs are inconsistent across datasets, ranging from -0.107 on ETTh1 (negative correlation, the desired energy-aware direction) to $+0.054$ on ETTm1 (positive correlation, the wrong direction). This dataset-level inconsistency is the precise reason we do not interpret the aggregate ($+0.0043$) as evidence of a learned energy-aware schedule.

What the per-dataset breakdown adds. On ETTh1 specifically, the train-time correlation is negative (-0.107) and consistent across all four horizons of that dataset. This shows the energy-aware mechanism




Figure 10: Per-run m_{weight} delta of the router-energy variant relative to the LMI-only variant. Every bar is negative (28/28 runs nudge the Mamba branch downward), confirming the directional finding of Section 8.5, but the magnitudes are uniformly in the 10^{-5} to 10^{-4} range.

is in principle capable of producing the desired sign of correlation, but it does not generalize to a uniform negative-sign pattern across the other six datasets. We expect that the underlying issue is that the router is trained on clean inputs, where local LMI violations are not strongly correlated with task loss, so the router has no incentive (in the task-loss gradient) to actively avoid the Mamba branch when the Mamba core is locally near a violation. The most direct path to a genuinely learned, energy-aware schedule, as discussed in Section 9, is to retrain with injected perturbations so that violation magnitude becomes correlated with task loss; we leave this as future work.