

Horizon-Uniform Sensitivity and Decay of Terminal Reward Perturbations in Discrete-Time Pontryagin Systems [★]

Pyuyi Chufeng Huang ^aZikang Song ^a

^a*School of Mathematics, Sichuan University, Chengdu, Sichuan, China*

Abstract

We study local stationary solutions of finite-horizon discrete-time Pontryagin systems near a steady extremal. Suppose that the stationarity equation for the control is regular, the reduced state–costate map is hyperbolic, and the endpoint conditions satisfy a scaled transversality condition with respect to the stable and unstable subspaces. Then the linearized boundary-value problem admits an inverse whose Green estimate is uniform in the horizon. The Green kernel separates interior decay from the two reflections induced by the endpoint conditions. For $x_0 = x_{\text{in}}$ and $p_T = r_x(x_T, y)$, a contraction argument in a weighted norm proves existence and uniqueness in a neighborhood independent of T , together with uniform Lipschitz estimates and a pointwise quadratic remainder. We also derive an explicit admissible data radius and an a posteriori criterion for existence and local uniqueness near an approximate trajectory. For these graph boundary conditions, a one-sided Green estimate shows that a perturbation of the terminal reward changes the initial control and the gradient with respect to the initial state of the stationary objective by $O(e^{-\alpha_{\text{ter}}T})$ for every α_{ter} below the dichotomy rate. For linear-quadratic systems with invertible A , stabilizable (A, B) , $Q \succ 0$, $R \succ 0$, and a nonpositive terminal Hessian, a symplectic graph condition verifies the assumptions, and the finite-horizon Riccati matrix and initial feedback gain converge at rate $O(e^{-2\gamma T})$. Numerical experiments verify the certificates and the predicted decay rates.

Key words: Optimal control; Boundary-value problems; Exponential dichotomies; Dynamic programming; Sensitivity analysis; Symplectic methods

1 Introduction

Finite-horizon optimal control methods repeatedly solve two-point Pontryagin equations, for example in shooting, continuation, and receding-horizon computations (Bertsekas, 2017; Grüne and Pannek, 2017). A fixed-horizon implicit function theorem gives a local stationary branch, but does not control the inverse uniformly as the horizon grows.

Dichotomy theory relates boundary transversality to well-conditioned linear boundary-value problems (Coppel, 1978). Implicit-function arguments in weighted sequence spaces yield horizon-uniform nonlinear sensitivity once a uniformly bounded linear inverse is available (Grüne, Schaller, and Schiela, 2021). Hyperbolic Hamiltonian equilibria and transverse stable and unstable manifolds also yield two-sided turnpike estimates. For optimal solutions and KKT systems, uniform expo-

nential sensitivity has been proved under second-order sufficient conditions, constraint qualifications, or controllability assumptions. Dynamic programming gives the complementary value-function description of the same optimal trajectories; its differentiable characteristics are the Hamiltonian state–costate equations (Bertsekas, 2017; Sassano, 2025). The correspondence is classical; here it determines when the stationary quantities certified below coincide with value gradients and feedback laws.

First, Na and Anitescu (2020) prove decay of directional sensitivities for locally optimal discrete-time programs under uniform second-order sufficiency, controllability, and derivative bounds, while Shin and Zavala (2021) relate second-order sufficiency and constraint regularity to controllability and observability; we instead use a verifiable transversality condition for a hyperbolic state–costate map to construct a local stationary Pontryagin branch without assuming optimality. Second, Shin, Anitescu, and Zavala (2022) prove that primal–dual sensitivity in graph-structured nonlinear programs decays with graph distance under a strong second-order sufficient condition and the linear independence constraint

[★] Corresponding author: Pyuyi Chufeng Huang.

Email addresses: pyuyi233@gmail.com (Pyuyi Chufeng Huang), 1289094751@qq.com (Zikang Song).

qualification; for the temporal chain, we separate the contributions of both boundaries and use a scaled boundary matrix to bound the influence of the terminal condition on the initial control and initial costate. Third, Breiten and Pfeiffer (2020) and Kunisch and Pfeiffer (2020) prove exponential receding-horizon estimates for infinite-dimensional linear-quadratic (LQ) problems and stabilization problems with terminal penalties, respectively; we instead bound terminal reward perturbations in a discrete-time nonlinear Pontryagin boundary problem, with optimality required only for its interpretation in optimal control. Fourth, Sakamoto and Zuazua (2021) derive nonlinear turnpike estimates from stable and unstable manifolds, while Xu and Anitescu (2018) prove exponential accuracy of an overlapping temporal decomposition in the overlap size; our framework combines the scaled boundary matrix with a horizon-independent nonlinear existence radius and pointwise quadratic remainder, adds a separate a posteriori existence test, and uses the LQ Möbius representation to obtain the $O(e^{-2\gamma T})$ bound on sensitivity to the terminal Hessian.

We seek a finite-horizon matrix criterion that controls both endpoint reflections and remains valid for the nonlinear terminal condition $p_T = r_x(x_T, y)$. Our contributions are as follows. First, we derive a Green estimate from the scaled matrices $\hat{\Gamma}_T$; its kernel separates interior decay from two endpoint reflections. Second, we prove the existence of a horizon-uniform local stationary branch, a pointwise remainder of order $O(w_T(t)^2 \|\Delta\|^2)$, and a uniform Lipschitz estimate for the initial control. We also derive explicit a priori and a posteriori radii for existence and local uniqueness. Third, for graph boundary conditions, we isolate the contribution of the terminal boundary and prove that perturbations of the terminal reward have exponentially small effects on controls near the initial time and on the gradient of the stationary objective with respect to the initial state. Fourth, for LQ systems with invertible A , stabilizable (A, B) , $Q \succ 0$, $R \succ 0$, and $S \preceq 0$, we reduce the matrix criterion to symplectic graph conditions, prove invertibility for every horizon, and obtain an explicit Möbius formula and $O(e^{-2\gamma T})$ convergence of the Riccati matrices and feedback gains.

The main results establish stationary extremals. An additional second-order sufficient condition is required to identify them as local optima. Section 2 gives the reduction, Section 3 proves the linear estimate, Sections 4–5 give the nonlinear results, decay under terminal reward perturbations, and computable certificates, Section 6 treats the LQ verification, and Section 7 reports the numerical tests.

2 Model, reduction, and standing assumptions

For a horizon $T \in \mathbb{N}$, the state is in $\mathbb{R}^d$, the control set $U \subset \mathbb{R}^m$ is compact and convex, and

$$x_{t+1} = F(x_t, u_t), \quad 0 \leq t < T. \quad (2.1)$$

For y in an open set $Y \subset \mathbb{R}^q$, the maximization objective and Hamiltonian are

$$J_T(u) = r(x_T, y) - \lambda \sum_{t=0}^{T-1} \ell(x_t, u_t), \quad \lambda > 0, \quad (2.2)$$

$$H(x, u, p) = -\lambda \ell(x, u) + \langle p, F(x, u) \rangle. \quad (2.3)$$

All local branch estimates use fixed neighborhoods, independent of T . The maps F and ℓ are C^3 near the reference, and terminal rewards $r(\cdot, y)$ are C^3 in x when nonlinear terminal conditions are used. The required uniform derivative bounds are stated below.

An interior stationary reference satisfies

$$\begin{aligned} x_\infty &= F(x_\infty, u_\infty), & p_\infty &= H_x(x_\infty, u_\infty, p_\infty), \\ H_u(x_\infty, u_\infty, p_\infty) &= 0, \end{aligned} \quad (2.4)$$

with $u_\infty \in \text{int } U$.

2.1 Pontryagin equations and sign convention

We use the maximization convention for (2.2) and the forward Hamiltonian (2.3); the costate paired with the step from t to $t+1$ is p_{t+1} . Consequently

$$\begin{aligned} H_x(x, u, p_+) &= \partial_x F(x, u)^\top p_+ - \lambda \partial_x \ell(x, u), \\ H_u(x, u, p_+) &= \partial_u F(x, u)^\top p_+ - \lambda \partial_u \ell(x, u). \end{aligned}$$

If a trajectory is locally optimal for (2.2), one-sided variations of a single control value in the convex set U give

$$\begin{aligned} x_{t+1} &= F(x_t, u_t), & p_t &= H_x(x_t, u_t, p_{t+1}), \\ p_T &= \partial_x r(x_T, y), \\ \langle H_u(x_t, u_t, p_{t+1}), v - u_t \rangle &\leq 0, \\ 0 \leq t < T, & & v \in U. \end{aligned} \quad (2.5)$$

These are the standard discrete-time Pontryagin conditions (Halkin, 1966; Bertsekas, 2017); the signs reflect maximization of terminal reward minus accumulated cost.

All sufficiently small branches considered below remain in $\text{int } U$. On such branches the variational inequality is equivalent to $H_u = 0$. The local stationarity graph constructed next therefore eliminates u_t as $u_t = \nu(x_t, p_{t+1})$

and turns the first two equations in (2.5) into the reduced implicit equation (2.11). The endpoint rows are retained rather than eliminated. Thus a stationary Pontryagin branch in this paper means a local branch of (2.5) after interior stationarity reduction.

2.2 Relation to Bellman recursion

For an optimal branch, define

$$V_{t,T}(x; y) = \max_{u_t, \dots, u_{T-1}} \left\{ r(x_T, y) - \lambda \sum_{k=t}^{T-1} \ell(x_k, u_k) \right\},$$

$$Q_{t,T}(x, u; y) = -\lambda \ell(x, u) + V_{t+1,T}(F(x, u); y),$$

with $V_{T,T} = r$. If the relevant value functions are differentiable and the Bellman maximizers are interior, then, along the optimal trajectory,

$$p_{t+1} = D_x V_{t+1,T}(x_{t+1}; y),$$

$$D_u Q_{t,T} = H_u = 0, \quad D_x V_{t,T} = H_x = p_t. \quad (2.6)$$

These identities follow from the chain rule and the envelope theorem. Thus the costate equals the state gradient of the future value, and the Pontryagin equations are the first-order characteristic equations associated with the differentiable Bellman recursion (Bertsekas, 2017; Sassano, 2025).

The Hamiltonian and the Bellman Q -function are nevertheless different: in general $H(x, u, p_+) \neq Q_{t,T}(x, u; y)$. If $V_{t+1,T}$ is twice differentiable, set $x_+ = F(x, u)$ and $p_+ = D_x V_{t+1,T}(x_+; y)$. Then

$$D_{uu}^2 Q_{t,T} = H_{uu} + F_u^\top D_{xx}^2 V_{t+1,T} F_u. \quad (2.7)$$

Consequently, (2.8) below guarantees only that the stationarity equation can be solved locally for u at fixed p_+ . It does not imply that u maximizes the Bellman Q -function or that the resulting branch is optimal.

Local control elimination. Assume that F and ℓ are C^3 near the reference and

$$H_u(x_\infty, u_\infty, p_\infty) = 0, \quad \det H_{uu}(x_\infty, u_\infty, p_\infty) \neq 0. \quad (2.8)$$

Put $\zeta_\infty = (x_\infty, p_\infty)$. The finite-dimensional implicit function theorem gives neighborhoods Ω of ζ_∞ and $V \Subset \text{int } U$ of u_∞ , and a unique C^2 map $\nu : \Omega \rightarrow V$ such that

$$H_u(x, \nu(x, p_+), p_+) = 0, \quad \nu(x_\infty, p_\infty) = u_\infty. \quad (2.9)$$

After shrinking the neighborhoods, $D\nu$ and $D^2\nu$ are bounded and

$$D\nu = -H_{uu}^{-1}(H_{ux}, H_{up}). \quad (2.10)$$

These are standard consequences of the implicit function theorem; no local maximum property is used.

With $z_t = (x_t, p_t)$, the reduced state-costate step is

$$\mathcal{F}(z_t, z_{t+1}) = 0, \quad (2.11)$$

where

$$\mathcal{F}(z_t, z_{t+1}) = \begin{pmatrix} x_{t+1} - F(x_t, \nu(x_t, p_{t+1})) \\ p_t - H_x(x_t, \nu(x_t, p_{t+1}), p_{t+1}) \end{pmatrix}. \quad (2.12)$$

Let $\nu_x = D_x \nu(x_\infty, p_\infty)$ and $\nu_p = D_{p_+} \nu(x_\infty, p_\infty)$. In (2.13), F_x, F_u are the derivatives of the original dynamics F at (x_∞, u_∞) and all Hamiltonian derivatives are evaluated at $(x_\infty, u_\infty, p_\infty)$ and H_{xp_+} denotes the derivative of H_x with respect to the forward costate argument. For $h = (\xi, \pi)$ and $h^+ = (\xi^+, \pi^+)$, the linearization is $L_0 h + L_1 h^+$ with

$$L_0 h = \begin{pmatrix} -(F_x + F_u \nu_x) \xi \\ \pi - (H_{xx} + H_{xu} \nu_x) \xi \end{pmatrix},$$

$$L_1 h^+ = \begin{pmatrix} \xi^+ - F_u \nu_p \pi^+ \\ -(H_{xp_+} + H_{xu} \nu_p) \pi^+ \end{pmatrix}. \quad (2.13)$$

Uniform local assumptions. Given a product neighborhood $Z = Z_x \times Z_p \subset \mathbb{R}^d \times \mathbb{R}^d$ of 0 contained in $\Omega - \zeta_\infty$, set

$$\tilde{\mathcal{F}}(z, z^+) := \mathcal{F}((x_\infty, p_\infty) + z, (x_\infty, p_\infty) + z^+), \quad (2.14)$$

$$z, z^+ \in Z.$$

The reduced state-costate step (2.11) on $(x_\infty, p_\infty) + Z$ is $\tilde{\mathcal{F}}(z, z^+) = 0$. We fix Z , a convex $U_0 \Subset \text{int } U$, and compact product balls $B_x \Subset Z_x, B_p \Subset Z_p$ such that ν maps all mixed arguments $(x_\infty + \xi, p_\infty + \pi^+)$ into U_0 and, for some constants $M_{\text{int}}, M_{\text{ter}}$,

$$\sup_{(x_\infty, p_\infty) + Z} \|D\nu\| + \sup_{(x_\infty, p_\infty) + Z} \|D^2\nu\| \leq M_{\text{int}},$$

$$\sup_{(z, z^+) \in (B_x \times B_p) \times (B_x \times B_p)} \|D^2 \tilde{\mathcal{F}}(z, z^+)\| \leq M_{\text{int}}.$$

For a reference $y_\infty \in Y$, assume that a neighborhood $Y_0 \Subset Y$ satisfies

$$\sup_{x \in x_\infty + B_x, y \in Y_0, 1 \leq j \leq 3} \|D_x^j r(x, y)\| \leq M_{\text{ter}}.$$

The maps $y \mapsto r_x(x, y)$ and $y \mapsto r_{xx}(x, y)$ are assumed continuous uniformly in this tube, and locally Lipschitz uniformly in x when joint Lipschitz dependence on (x_{in}, y) is asserted. All radii below are chosen inside this fixed tube. Finally, direct verification from problem data uses

$$\det L_1 \neq 0, \quad M := -L_1^{-1} L_0. \quad (2.15)$$

Thus the linearized interior equation is $h_{t+1} = Mh_t + L_1^{-1}f_t$. The alternative abstract formulation only requires the Green property stated next.

3 Linear endpoint estimates

The nonlinear argument requires the following uniform inverse estimate.

Definition 3.1 (Uniform Green property). Let $L : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ be linear and let $B_0, B_T \in \mathbb{R}^{n \times n}$ be fixed boundary-row matrices, not necessarily invertible individually. Given constants $C_G \geq 1$, $\gamma > 0$, and $T_0 \geq 1$, the triple (L, B_0, B_T) has the uniform Green property with these constants if, for every $T \geq T_0$, every interior forcing $f = (f_0, \dots, f_{T-1})$, and every boundary datum $b \in \mathbb{R}^n$, the following two conditions hold:

(WP) The linear problem

$$\begin{aligned} L[h_t, h_{t+1}] &= f_t, & 0 \leq t < T, \\ B_0 h_0 + B_T h_T &= b. \end{aligned} \quad (3.1)$$

has a unique solution $h = (h_0, \dots, h_T)$.

(G) With fixed finite-dimensional norms and

$$\begin{aligned} \kappa_T(t, s) &= e^{-\gamma|t-s|} + e^{-\gamma(t+T-s)} + e^{-\gamma(T-t+s)}, \\ &0 \leq s < T, \end{aligned} \quad (3.2)$$

the solution from (WP) satisfies

$$\begin{aligned} \|h_t\| &\leq C_G \left[(e^{-\gamma t} + e^{-\gamma(T-t)}) \|b\| \right. \\ &\quad \left. + \sum_{s=0}^{T-1} \kappa_T(t, s) \|f_s\| \right] \end{aligned} \quad (3.3)$$

for all $0 \leq t \leq T$.

When only the property is named, the constants are independent of T .

The last two terms in (3.2) are the finite-horizon boundary reflections. We write $\mathcal{L}_T^{\text{bd}} b$ and $\mathcal{L}_T^{\text{for}} f$ for the boundary and forcing solutions.

A concrete verification comes from the explicit transition (2.15) and a uniformly invertible scaled boundary matrix.

Theorem 3.2 (Finite-horizon Green certificate). Consider, for a horizon N ,

$$h_{t+1} = Mh_t + g_t, \quad 0 \leq t < N, \quad B_0 h_0 + B_T h_N = b, \quad (3.4)$$

on $\mathbb{R}^n$. Suppose M is hyperbolic and let the columns of V_s and V_u be bases of its stable and unstable subspaces,

with $\dim E_s + \dim E_u = n$; write $MV_s = V_s D_s$ and $MV_u = V_u D_u$, and assume that, for some $K \geq 1$ and $\gamma > 0$, $\|V_s D_s^k\| \leq K e^{-\gamma k}$ and $\|V_u D_u^{-k}\| \leq K e^{-\gamma k}$ for all $k \geq 0$. For each horizon define the scaled boundary matrix

$$\widehat{\Gamma}_N = \begin{bmatrix} B_0 V_s + B_T V_s D_s^N & B_0 V_u D_u^{-N} + B_T V_u \end{bmatrix}, \quad (3.5)$$

where the superscripts are powers, not transposes. If $\widehat{\Gamma}_N$ is invertible for every $N \geq T_0$ and $\sup_{N \geq T_0} \|\widehat{\Gamma}_N^{-1}\| < \infty$, then the uniform Green property of Definition 3.1 holds for (3.4), with constants depending only on the dichotomy bounds, the spectral-projector norms, the fixed boundary matrices, and the displayed uniform bound for $\widehat{\Gamma}_N^{-1}$. Moreover, the uniform inverse bound follows if

$$\widehat{\Gamma}_\infty := \begin{bmatrix} B_0 V_s & B_T V_u \end{bmatrix} \quad (3.6)$$

is invertible and, for some cutoff N_* large enough for the Neumann argument, each matrix $\widehat{\Gamma}_N$ with $T_0 \leq N < N_*$ is invertible.

We say that *scaled boundary transversality* holds if $\widehat{\Gamma}_N$ in (3.5) is invertible for every $N \geq T_0$ with $\sup_{N \geq T_0} \|\widehat{\Gamma}_N^{-1}\| < \infty$. The bases are fixed once and for all: replacing them by a fixed pair $V_s A_s, V_u A_u$ right-multiplies $\widehat{\Gamma}_N$ by $\text{diag}(A_s, A_u)$, so the property is basis independent up to condition numbers, while no N -dependent rescaling is allowed.

Proof. Let $V = [V_s \ V_u]$, write $V^{-1} = \begin{pmatrix} W_s \\ W_u \end{pmatrix}$, and set $\Pi_s = V_s W_s$, $\Pi_u = V_u W_u$. Every homogeneous solution has the unique form

$$q_t = V_s D_s^t a + V_u D_u^{t-N} \eta.$$

Its boundary equation is $\widehat{\Gamma}_N \begin{pmatrix} a \\ \eta \end{pmatrix} = b$, so uniform invertibility gives

$$\|q_t\| \leq C(e^{-\gamma t} + e^{-\gamma(N-t)}) \|b\|. \quad (3.7)$$

For forcing, define

$$p_t = \sum_{j < t} M^{t-1-j} \Pi_s g_j - \sum_{j \geq t} (M|_{E_u})^{t-1-j} \Pi_u g_j.$$

Then $p_{t+1} - M p_t = g_t$ and

$$\|p_t\| \leq C \sum_j e^{-\gamma|t-j|} \|g_j\|, \quad (3.8)$$

$$\|p_0\| + \|p_N\| \leq C \sum_j (e^{-\gamma j} + e^{-\gamma(N-j)}) \|g_j\|. \quad (3.9)$$

Correct the boundary residual $b - B_0 p_0 - B_T p_N$ by a homogeneous solution. Multiplying (3.9) by the two factors in (3.7) produces four terms. The products $e^{-\gamma(t+j)}$ and $e^{-\gamma(2N-t-j)}$ are bounded by $e^{-\gamma|t-j|}$; the other two are $e^{-\gamma(t+N-j)}$ and $e^{-\gamma(N-t+j)}$. Together with (3.8), these are exactly the terms in $\kappa_N(t, j)$. This proves (3.3). With zero forcing and boundary datum, invertibility of $\hat{\Gamma}_N$ gives uniqueness.

Finally, $\|\hat{\Gamma}_N - \hat{\Gamma}_\infty\| \leq C e^{-\gamma N}$. Neumann's lemma covers all sufficiently large N , and the finitely many remaining inverses give one uniform bound. $\square$

This scaled boundary transversality property is also open: under sufficiently small continuous perturbations of M, B_0, B_T , the Riesz projectors and local bases vary continuously, a common spectral gap persists, and $\hat{\Gamma}_N \rightarrow \hat{\Gamma}_\infty$ uniformly. Neumann's lemma covers large N , while continuity of the finitely many remaining least singular values covers short horizons. Any exponential rate below the stable and unstable spectral separation is admissible after increasing K .

The two-sided kernel above is needed for general endpoint data. For Pontryagin graph rows, terminal-only data admit a sharper, one-sided estimate. Write $P_x = (I \ 0)$ and $R_S = (-S \ I)$ on $\mathbb{R}^d \times \mathbb{R}^d$.

Lemma 3.3 (One-sided terminal Green estimate). In Theorem 3.2, let $n = 2d$, $\dim E_s = \dim E_u = d$, and write $V_s = \begin{pmatrix} X_s \\ P_s \end{pmatrix}$ and $V_u = \begin{pmatrix} X_u \\ P_u \end{pmatrix}$. Let $\mathcal{S}$ be a bounded set of matrices for which X_s is nonsingular and

$$\hat{\Gamma}_T(S) = \begin{pmatrix} X_s & X_u D_u^{-T} \\ (P_s - S X_s) D_s^T & P_u - S X_u \end{pmatrix} \quad (3.10)$$

has a uniformly bounded inverse for $S \in \mathcal{S}$ and $T \geq T_0$. For $0 < \alpha_{\text{ter}} < \gamma$, put $v_T(t) = e^{-\alpha_{\text{ter}}(T-t)}$. The solution of

$$h_{t+1} = M h_t + g_t, \quad P_x h_0 = 0, \quad R_S h_T = b_T \quad (3.11)$$

satisfies

$$\max_{0 \leq t \leq T} \frac{\|h_t\|}{v_T(t)} \leq C \rightarrow \left(\|b_T\| + \max_{0 \leq s < T} \frac{\|g_s\|}{v_T(s+1)} \right), \quad (3.12)$$

where $C_\rightarrow$ is independent of (S, T) .

Proof. In the dichotomy representation, the boundary system is $(3.10)(a, \eta)^\top = (b_L, b_R)^\top$, where, for $G =$

$$\max_j \|g_j\|/v_T(j+1),$$

$$b_L = P_x V_u \sum_{j < T} D_u^{-1-j} W_u g_j, \\ b_R = b_T - R_S V_s \sum_{j < T} D_s^{T-1-j} W_s g_j.$$

Geometric sums give $b_L = O(v_T(0)G)$ and $b_R = O(\|b_T\| + G)$. Uniform inversion bounds η ; the first row and X_s^{-1} then give the extra estimate $a = O(v_T(0)(\|b_T\| + G))$. Substitution uses the summable ratios $e^{-(\gamma+\alpha_{\text{ter}})k}$ and $e^{-(\gamma-\alpha_{\text{ter}})k}$, proving the claim. $\square$

The following quantitative row-perturbation result will be used for the parameter-dependent terminal Hessian in Section 5.

Lemma 3.4 (Quantitative perturbation of endpoint rows). Suppose that (L, B_0, B_T) has the uniform Green property with constants (C_G, γ, T_0) . All induced operator norms below use the fixed finite-dimensional norms of Definition 3.1. For endpoint-row perturbations E_0, E_T , put

$$\varepsilon = \|E_0\| + \|E_T\|, \quad \vartheta = 2C_G \varepsilon. \quad (3.13)$$

If $\vartheta < 1$, then $(L, B_0 + E_0, B_T + E_T)$ has the same Green exponent and initial horizon, with the common admissible constant

$$C_G^{\text{pert}} = C_G \frac{1 + \vartheta}{1 - \vartheta}. \quad (3.14)$$

The conclusion is uniform whenever ϑ is bounded above by a fixed number smaller than one.

Proof. Use the unperturbed inverse to write the perturbed problem as

$$h = \mathcal{L}_T^{\text{for}} f + \mathcal{L}_T^{\text{bd}}(b - E_0 h_0 - E_T h_T). \quad (3.15)$$

Because $\max_{0 \leq t \leq T} \|(\mathcal{L}_T^{\text{bd}} c)_t\| \leq 2C_G \|c\|$, its right-hand side is a contraction with factor ϑ , proving existence. Set

$$m = \max\{\|h_0\|, \|h_T\|\}, \\ F_\vartheta(f) = \sum_{s=0}^{T-1} (e^{-\gamma s} + e^{-\gamma(T-s)}) \|f_s\|.$$

At $t = 0, T$, the boundary factor and the Green kernel are bounded by twice the corresponding terms in $\|b\| + F_\vartheta(f)$. Hence

$$m \leq 2C_G(\|b\| + F_\vartheta(f) + \varepsilon m) \leq \frac{2C_G}{1 - \vartheta}(\|b\| + F_\vartheta(f)). \quad (3.16)$$

Writing $d_T(t) = e^{-\gamma t} + e^{-\gamma(T-t)}$, direct multiplication gives

$$d_T(t)(e^{-\gamma s} + e^{-\gamma(T-s)}) \leq 2\kappa_T(t, s). \quad (3.17)$$

Indeed, the cross-products are the two reflected kernels, whereas each of the other products is bounded by $e^{-\gamma|t-s|}$. Substitution of (3.16) into the original point-wise Green estimate yields

$$\begin{aligned} \|h_t\| &\leq \frac{C_G}{1-\vartheta} d_T(t) \|b\| \\ &\quad + \frac{C_G(1+\vartheta)}{1-\vartheta} \sum_{s=0}^{T-1} \kappa_T(t, s) \|f_s\|, \end{aligned}$$

and hence Definition 3.1 with (3.14). For homogeneous data, (3.16) gives $h_0 = h_T = 0$, and unperturbed well-posedness then gives $h = 0$; thus the perturbed problem is well posed. $\square$

4 Nonlinear reconstruction in weighted spaces

The nonlinear fixed-point argument uses endpoint weights matched to the Green kernel. For $0 < \alpha < \gamma/2$, set

$$w_T(t) = e^{-\alpha t} + e^{-\alpha(T-t)}. \quad (4.1)$$

For each horizon define the weighted space

$$\begin{aligned} X_T &= \{z = (z_0, \dots, z_T) : \|z\|_{w,T} < \infty\}, \\ \|z\|_{w,T} &= \sup_{0 \leq t \leq T} \frac{\|z_t\|}{w_T(t)}. \end{aligned}$$

All differentiability statements below concern the family of maps into these horizon-dependent spaces uniformly in T .

Weighted convolution. If $0 < \alpha < \gamma/2$, then there is a constant C_{conv} , independent of T and t , such that

$$\sum_{s=0}^{T-1} \kappa_T(t, s) w_T(s)^2 \leq C_{\text{conv}} w_T(t)^2.$$

Moreover $e^{-\gamma t} + e^{-\gamma(T-t)} \leq C_{\text{conv}} w_T(t)^2$, with κ_T from (3.2). For example, with $q_- = e^{-(\gamma-2\alpha)}$ and $q_+ = e^{-(\gamma+2\alpha)}$, one may use the explicit value

$$C_{\text{conv}} = 4 \left(\frac{1}{1-q_-} + \frac{q_+}{1-q_+} \right). \quad (4.2)$$

To verify this value, put $\chi_s^- = e^{-\alpha s}$, $\chi_s^+ = e^{-\alpha(T-s)}$, and define

$$A_{\gamma,\alpha} = \frac{1}{1-q_-} + \frac{q_+}{1-q_+}.$$

Set $\kappa_T^{\text{ref}}(t, s) = e^{-\gamma(t+T-s)} + e^{-\gamma(T-t+s)}$. Since $w_T(s)^2 \leq 2((\chi_s^-)^2 + (\chi_s^+)^2)$, splitting the interior sum at $s = t$ gives the two geometric ratios q_- and q_+ and hence

$$\sum_{s=0}^{T-1} e^{-\gamma|t-s|} w_T(s)^2 \leq 2A_{\gamma,\alpha} ((\chi_t^-)^2 + (\chi_t^+)^2), \quad (4.3)$$

$$\sum_{s=0}^{T-1} \kappa_T^{\text{ref}}(t, s) w_T(s)^2 \leq 2A_{\gamma,\alpha} ((\chi_t^-)^2 + (\chi_t^+)^2). \quad (4.4)$$

For the reflected kernels, the coefficients of $(\chi_t^-)^2$ are $q_+/(1-q_+)$ and $1/(1-q_-)$; those of $(\chi_t^+)^2$ are $q_-/(1-q_-)$ and $1/(1-q_+)$. Both pairs sum to $A_{\gamma,\alpha}$. Because $(\chi_t^-)^2 + (\chi_t^+)^2 \leq w_T(t)^2$, the two estimates prove (4.2). Moreover, $\gamma > 2\alpha$ implies

$$e^{-\gamma t} + e^{-\gamma(T-t)} \leq (\chi_t^-)^2 + (\chi_t^+)^2 \leq w_T(t)^2.$$

Thus the same squared weight controls the interior and endpoint remainders.

Quadratic residual. Let $\mathcal{F}$ be C^2 near $(0,0)$, assume $\mathcal{F}(0,0) = 0$, and let $L = D\mathcal{F}(0,0)$. Define

$$N(a, b) = \mathcal{F}(a, b) - L(a, b).$$

After shrinking the neighborhood, there is a constant C_N such that

$$\|N(a, b)\| \leq C_N (\|a\| + \|b\|)^2$$

and

$$\begin{aligned} \|N(a, b) - N(\tilde{a}, \tilde{b})\| &\leq C_N (\|a\| + \|b\| + \|\tilde{a}\| + \|\tilde{b}\|) \\ &\quad \times (\|a - \tilde{a}\| + \|b - \tilde{b}\|). \end{aligned}$$

Both estimates follow directly from Taylor's formula with integral remainder on a fixed convex neighborhood.

We now combine the inverse estimate, the convolution bound, and the residual estimate. Affine endpoint rows are included as the case with zero remainder.

Theorem 4.1 (Uniform reconstruction with smooth endpoint rows). Let $\mathcal{F} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ be C^2 near $(0,0)$, satisfy $\mathcal{F}(0,0) = 0$, and have the quadratic residual bounds above. Assume that its linearized equation with rows $(\mathbf{B}_0, \mathbf{B}_T)$ satisfies Definition 3.1 with exponent γ . Fix $0 < \alpha < \gamma/2$. Let $\mathcal{B} : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ be C^2 near the origin, with bounded second derivative, and satisfy $\mathcal{B}(0,0) = 0$ and $D\mathcal{B}(0,0)[a, b] = \mathbf{B}_0 a + \mathbf{B}_T b$ (the affine row is allowed). Put $R_B(a, b) = \mathcal{B}(a, b) - \mathbf{B}_0 a - \mathbf{B}_T b$ and choose C_B such that

$$\|R_B(a, b)\| \leq C_B (\|a\| + \|b\|)^2, \quad (4.5)$$

$$\|R_B(a, b) - R_B(\tilde{a}, \tilde{b})\| \leq C_B \Sigma(\|a - \tilde{a}\| + \|b - \tilde{b}\|), \quad (4.6)$$

where $\Sigma = \|a\| + \|b\| + \|\tilde{a}\| + \|\tilde{b}\|$.

Let C_N be the interior residual constant above, set $c = C_G$, and let

$$a = 2C_G C_{\text{conv}} C_N (1 + e^\alpha)^2 + 16C_G C_B. \quad (4.7)$$

Choose any $R_* > 0$ such that the weighted ball of radius R_* lies in the fixed smoothness tube and $2aR_* < 1$, and define

$$\delta_* := \frac{R_* - aR_*^2}{c}. \quad (4.8)$$

Then, for $T \geq T_0$ and $d := \|\Delta\| < \delta_*$, the problem

$$\mathcal{F}(z_t, z_{t+1}) = 0, \quad \mathcal{B}(z_0, z_T) = \Delta \quad (4.9)$$

has a solution in the explicitly computable ball

$$\|z\|_{w,T} \leq R_-(d) := \frac{1 - \sqrt{1 - 4acd}}{2a} < R_*, \quad (4.10)$$

where, if $a = 0$, the continuous convention is $R_-(d) = cd$. It is the only solution in $\|z\|_{w,T} \leq R_*$. With $z^{\text{lin}}(\Delta) := \mathcal{L}_T^{\text{bd}} \Delta$ and $E_t(\Delta) := z_t(\Delta) - z_t^{\text{lin}}(\Delta)$, uniformly in T ,

$$\|z_t(\Delta)\| \leq Cw_T(t)\|\Delta\| + Cw_T(t)^2\|\Delta\|^2, \quad (4.11)$$

$$\|E_t(\Delta)\| \leq Cw_T(t)^2\|\Delta\|^2. \quad (4.12)$$

$$\|z_t(\Delta_1) - z_t(\Delta_2)\| \leq Cw_T(t)\|\Delta_1 - \Delta_2\|. \quad (4.13)$$

Moreover,

$$\|E_t(\Delta_1) - E_t(\Delta_2)\| \leq Cw_T(t)^2(\|\Delta_1\| + \|\Delta_2\|) \times \|\Delta_1 - \Delta_2\|. \quad (4.14)$$

Consequently the solution family is Fréchet differentiable at 0 uniformly in T , with derivative $\mathcal{L}_T^{\text{bd}}$.

Proof. Let $\mathcal{A}_T$ denote the linear interior and boundary operator, and let $\mathcal{N}_T(z)$ collect $N(z_t, z_{t+1})$ and $R_B(z_0, z_T)$. The Green estimate, $w_T(t+1) \leq e^\alpha w_T(t)$, and the weighted convolution give

$$\|\mathcal{A}_T^{-1} \mathcal{N}_T(z)\|_{w,T} \leq a\|z\|_{w,T}^2, \quad (4.15)$$

$$\|\mathcal{A}_T^{-1}(\mathcal{N}_T(z) - \mathcal{N}_T(\tilde{z}))\|_{w,T} \leq a(\|z\|_{w,T} + \|\tilde{z}\|_{w,T}) \times \|z - \tilde{z}\|_{w,T}. \quad (4.16)$$

Indeed, the interior contribution is bounded by the first term in (4.7). At the endpoints $w_T(0), w_T(T) \leq 2$, so the two factors in (4.6) are at most $4(\|z\|_{w,T} + \|\tilde{z}\|_{w,T})$ and $4\|z - \tilde{z}\|_{w,T}$; this gives the second term in (4.7).

Define

$$\Phi_T(z) = \mathcal{L}_T^{\text{bd}} \Delta - \mathcal{A}_T^{-1} \mathcal{N}_T(z).$$

Since $\|\mathcal{L}_T^{\text{bd}} \Delta\|_{w,T} \leq cd$, (4.15) shows that the ball of radius $R_-(d)$ is invariant: $cd + aR_-^2 = R_-$. Equation (4.16) gives contraction factor $2aR_- < 1$. Banach's theorem proves existence. If $z, \tilde{z}$ are any two solutions in the ball of radius R_* , then

$$\|z - \tilde{z}\|_{w,T} \leq 2aR_* \|z - \tilde{z}\|_{w,T},$$

so they coincide. This proves the computable radius and uniqueness claims.

Subtracting z^{lin} from the fixed-point equation leaves the quadratic interior and endpoint residuals. The sharper convolution estimate for w_T^2 and $e^{-\gamma t} + e^{-\gamma(T-t)} \leq Cw_T(t)^2$ give (4.12) and then (4.11). For two data Δ_1, Δ_2 with $d_i := \|\Delta_i\| < \delta_*$, write

$$R_i = R_-(d_i), \quad q_{12} = a(R_1 + R_2) < 1.$$

Subtracting the fixed-point equations and applying (4.16) gives

$$\|z_1 - z_2\|_{w,T} \leq \frac{c}{1 - q_{12}} \|\Delta_1 - \Delta_2\|. \quad (4.17)$$

For the pointwise remainder difference, set

$$\bar{a} = C_G C_{\text{conv}} C_N (1 + e^\alpha)^2 + 16C_G C_B \leq a. \quad (4.18)$$

The interior residual difference is at most

$$C_N (1 + e^\alpha)^2 w_T(s)^2 (R_1 + R_2) \|z_1 - z_2\|_{w,T},$$

and (4.6), together with $w_T(0), w_T(T) \leq 2$, gives the boundary bound $16C_B(R_1 + R_2)\|z_1 - z_2\|_{w,T}$. The estimates for w_T^2 therefore yield

$$\|E_t(\Delta_1) - E_t(\Delta_2)\| \leq \frac{\bar{a}c(R_1 + R_2)}{1 - q_{12}} w_T(t)^2 \|\Delta_1 - \Delta_2\|. \quad (4.19)$$

Since $R_i = cd_i + aR_i^2$ and $aR_i < aR_* < 1/2$, $R_i \leq 2cd_i$ and $q_{12} \leq 2aR_*$. Consequently (4.13) and (4.14) hold with

$$C_{\text{Lip}} = \frac{c}{1 - 2aR_*}, \quad C_{\text{rem}} = \frac{2\bar{a}c^2}{1 - 2aR_*}. \quad (4.20)$$

All constants depend only on the displayed uniform bounds, not on T . $\square$

The restriction $\alpha < \gamma/2$ is needed for the remainder with weight w_T^2 . Existence, size, and Lipschitz continuity alone hold for every $0 < \alpha < \gamma$ by the corresponding convolution estimate for w_T .

5 Application to Pontryagin boundary conditions

We now apply Theorem 4.1 directly to the reduced Pontryagin equation and its nonlinear terminal condition.

Theorem 5.1 (Uniform local Pontryagin branch with a nonlinear terminal condition). Assume the local hypotheses of Section 2, including (2.8), and let $y_\infty \in Y$ satisfy $p_\infty = r_x(x_\infty, y_\infty)$. For y near y_∞ define

$$\mathcal{B}_y(z_0, z_T) = \begin{pmatrix} \xi_0 \\ \pi_T - (r_x(x_\infty + \xi_T, y) - r_x(x_\infty, y)) \end{pmatrix},$$

$$\Delta_{\text{oc}}(x_{\text{in}}, y) = \begin{pmatrix} x_{\text{in}} - x_\infty \\ r_x(x_\infty, y) - p_\infty \end{pmatrix},$$

where $z = (\xi, \pi) = (x - x_\infty, p - p_\infty) \in \mathbb{R}^{2d}$. Then $\mathcal{B}_y(z_0, z_T) = \Delta_{\text{oc}}(x_{\text{in}}, y)$ is equivalent to

$$x_0 = x_{\text{in}}, \quad p_T = r_x(x_T, y).$$

If the rows $D\mathcal{B}_{y_\infty}(0, 0)$ satisfy Definition 3.1 with constants (C_G, γ, T_0) , choose a relatively compact neighborhood $Y_1 \subseteq Y_0$. With

$$E_T(y) = \begin{pmatrix} 0 & 0 \\ -[r_{xx}(x_\infty, y) - r_{xx}(x_\infty, y_\infty)] & 0 \end{pmatrix},$$

shrink Y_1 so that

$$\bar{\vartheta} := 2C_G \sup_{y \in Y_1} \|E_T(y)\| < 1. \quad (5.1)$$

Then Lemma 3.4 gives the common Green constant

$$\bar{C}_G = C_G \frac{1 + \bar{\vartheta}}{1 - \bar{\vartheta}}. \quad (5.2)$$

Set $C_B = M_{\text{ter}}$ in (4.5)–(4.6), define the common constants a, c from (4.7) using $\bar{C}_G$, and choose one pair R_*, δ_* as in (4.8). Then, for every $T \geq T_0$ and every (x_{in}, y) satisfying $d := \|\Delta_{\text{oc}}(x_{\text{in}}, y)\| < \delta_*$, the problem admits a unique local stationary Pontryagin branch in the fixed ball $\|z\|_{w,T} \leq R_*$. Its radius is bounded by (4.10), uniformly in T and y .

For fixed y , let $z_{T,y}^{\text{lin}} = (\xi^{\text{lin}}, \pi^{\text{lin}})$ be the linear branch formed with $D\mathcal{B}_y(0, 0)$, and put

$$u_t^{\text{lin}} = D\nu(x_\infty, p_\infty)[\xi_t^{\text{lin}}, \pi_{t+1}^{\text{lin}}].$$

Uniformly for $0 \leq t < T$,

$$\|z_t\| + \|u_t - u_\infty\| \leq Cw_T(t)d, \quad (5.3)$$

$$\|z_t - z_{T,y,t}^{\text{lin}}\| + \|(u_t - u_\infty) - u_t^{\text{lin}}\| \leq Cw_T(t)^2 d^2. \quad (5.4)$$

For two branches with the same y ,

$$\|z_t^1 - z_t^2\| + \|u_t^1 - u_t^2\| \leq Cw_T(t)\|\Delta_{\text{oc}}^1 - \Delta_{\text{oc}}^2\|. \quad (5.5)$$

The z -estimates also hold at $t = T$. The expansion is with respect to the combined boundary datum Δ_{oc} ; at fixed y it is an expansion at $x_{\text{in}} = x_\infty$ only when $r_x(x_\infty, y) = p_\infty$.

If, in addition, $y \mapsto r_x(x, y)$ and $y \mapsto r_{xx}(x, y)$ are locally Lipschitz uniformly in x on the terminal tube, then the branch is jointly locally Lipschitz in (x_{in}, y) on this neighborhood of the endpoint data.

The Green hypothesis is verified directly from problem data if (2.15) holds, M is hyperbolic, and the matrices $\hat{\Gamma}_T$ for $D\mathcal{B}_{y_\infty}(0, 0)$ satisfy the finite-horizon certificate of Theorem 3.2.

Proof. The boundary identity is immediate from the definitions. Its derivative is

$$D\mathcal{B}_y(0, 0)[z_0, z_T] = \begin{pmatrix} \xi_0 \\ \pi_T - r_{xx}(x_\infty, y)\xi_T \end{pmatrix}.$$

The terminal remainder is

$$-\left(\int_0^1 [r_{xx}(x_\infty + \theta\xi_T, y) - r_{xx}(x_\infty, y)]\xi_T d\theta\right),$$

so the uniform $D_x^3 r$ bound gives (4.5)–(4.6) with $C_B = M_{\text{ter}}$. Continuity of r_{xx} in y permits (5.1). The only endpoint-row perturbation is $E_T(y)$; in the standard Euclidean product norm, its norm is the norm of the displayed difference between terminal Hessians. Hence Lemma 3.4 gives (5.2) uniformly on Y_1 .

Theorem 4.1 now gives the explicit radius, uniqueness, and the estimates for z . Set $u_t = \nu(x_\infty + \xi_t, p_\infty + \pi_{t+1})$. The fixed tube keeps these controls in $\text{int } U$, and (2.9) recovers the full stationary Pontryagin equations. Taylor's formula for ν , together with $w_T(t+1) \leq e^\alpha w_T(t)$, transfers the size, remainder, and Lipschitz estimates to u .

For two parameters y_1, y_2 , subtract the interior and boundary equations. The additional terms contain $r_x(\cdot, y_1) - r_x(\cdot, y_2)$ and $r_{xx}(\cdot, y_1) - r_{xx}(\cdot, y_2)$. On an endpoint ball of radius ε , the uniform Lipschitz bounds and the common Green estimate give

$$\|z^1 - z^2\|_{w,T} \leq C\|x_{\text{in}}^1 - x_{\text{in}}^2\| + C\|y_1 - y_2\| + C\varepsilon\|z^1 - z^2\|_{w,T}. \quad (5.6)$$

The constant is horizon independent by (5.2). Shrinking the fixed neighborhood of the endpoint data until

$C\varepsilon < 1/2$ absorbs the last term and proves the joint assertion. The final statement in terms of the problem data is Theorem 3.2 applied to $M = -L_1^{-1}L_0$. $\square$

Proposition 5.2 (A posteriori existence and local uniqueness certificate). Let $\mathcal{R}_{T,\Delta}$ be the full interior and boundary residual of (4.9), obtained by stacking the T interior residuals and the boundary residual in $\mathbb{R}^{n(T+1)}$. Equip the trajectory space X_T with a chosen norm $\|\cdot\|_X$; this may be $\|\cdot\|_{w,T}$, while Section 7 uses the ordinary vector infinity norm. Let $\bar{z}$ be an approximate solution and put $J = D\mathcal{R}_{T,\Delta}(\bar{z})$. Suppose J is invertible and define

$$\beta = \|J^{-1}\mathcal{R}_{T,\Delta}(\bar{z})\|_X. \quad (5.7)$$

Assume that, whenever z and $\bar{z}$ are in $\|\cdot - \bar{z}\|_X \leq R_{\text{tube}}$,

$$\begin{aligned} \|J^{-1}(D\mathcal{R}_{T,\Delta}(z) - D\mathcal{R}_{T,\Delta}(\bar{z}))\|_{\mathcal{L}(X_T)} \\ \leq \omega \|z - \bar{z}\|_X. \end{aligned} \quad (5.8)$$

If $2\omega\beta < 1$ and

$$r_- := \frac{1 - \sqrt{1 - 2\omega\beta}}{\omega} < R_{\text{tube}}, \quad (5.9)$$

then an exact solution exists in $\|z - \bar{z}\|_X \leq r_-$. For every

$$r_- < R < \min\{R_{\text{tube}}, 1/\omega\},$$

it is the unique zero in $\|z - \bar{z}\|_X < R$. For $\omega = 0$, use the conventions $r_- = \beta$ and $1/\omega = +\infty$.

Proof. The frozen-Newton map $\Psi(z) = z - J^{-1}\mathcal{R}_{T,\Delta}(z)$ satisfies

$$\begin{aligned} \|\Psi(\bar{z} + h) - \bar{z}\|_X &\leq \beta + \frac{1}{2}\omega\|h\|_X^2, \\ \|D\Psi(\bar{z} + h)\|_{\mathcal{L}(X_T)} &\leq \omega\|h\|_X. \end{aligned}$$

Equation (5.9) is the smaller root of $\beta + \omega r^2/2 = r$ and obeys $\omega r_- < 1$. Banach's theorem proves existence and uniqueness in the smaller ball. On any larger ball with $\omega r < 1$, two zeros would be two fixed points of a strict contraction; hence the computed zero is the only one there. $\square$

All quantities in Proposition 5.2 are finite-dimensional and computable. Double-precision evaluation is a diagnostic; a formal computer-assisted certificate requires outward-rounded interval bounds for β and ω .

Corollary 5.3 (Uniform sensitivity of the initial control). Under the hypotheses of Theorem 5.1, define

$$\mathbf{K}_T(x_{\text{in}}, y) := u_0(x_{\text{in}}, y),$$

the first control of the local stationary branch. After possibly shrinking the neighborhood of the endpoint data, there is a constant L independent of (T, y) such that, for every fixed admissible y , every $T \geq T_0$, and every two admissible initial states x_1, x_2 ,

$$\|\mathbf{K}_T(x_1, y) - \mathbf{K}_T(x_2, y)\| \leq L\|x_1 - x_2\|.$$

If the final Lipschitz assumption of Theorem 5.1 holds, then for arbitrary admissible (x_1, y_1) and (x_2, y_2) ,

$$\begin{aligned} \|\mathbf{K}_T(x_1, y_1) - \mathbf{K}_T(x_2, y_2)\| &\leq L\|x_1 - x_2\| \\ &\quad + L\|y_1 - y_2\|. \end{aligned}$$

Moreover, for each fixed y , set $\Delta_{\text{oc}} = \Delta_{\text{oc}}(x_{\text{in}}, y)$. With

$$z_{T,y}^{\text{lin}}(\Delta_{\text{oc}}) = (\xi^{\text{lin}}, \pi^{\text{lin}})$$

denoting the linearized branch formed with $D\mathcal{B}_y(0, 0)$,

$$\begin{aligned} \mathbf{K}_T(x_{\text{in}}, y) &= u_\infty + D\nu(x_\infty, p_\infty)[\xi_0^{\text{lin}}, \pi_1^{\text{lin}}] \\ &\quad + O(\|\Delta_{\text{oc}}(x_{\text{in}}, y)\|^2), \end{aligned}$$

with a constant independent of (T, y) . This expansion treats the initial and terminal boundary data jointly, subject to the qualification stated after (5.5).

Proof. Take $t = 0$ in (5.5) and (5.4), and use $w_T(0) \leq 2$. The joint estimate follows from the last part of Theorem 5.1. $\square$

Proposition 5.4 (Stationary objective and gradient with respect to the initial state). For the branch of Theorem 5.1, define the objective evaluated along the stationary branch by

$$\mathcal{J}_T(x, y) = r(x_T(x, y), y) - \lambda \sum_{t=0}^{T-1} \ell(x_t(x, y), u_t(x, y)). \quad (5.10)$$

After shrinking the data neighborhood, the branch is C^1 in x for each fixed admissible y , and

$$D_x \mathcal{J}_T(x, y) = p_0(x, y), \quad \text{Lip}(D_x \mathcal{J}_T) \leq C, \quad (5.11)$$

with C independent of T for fixed admissible y . If the stationary branch is the unique maximizer in the specified tube, then $\mathcal{J}_T$ is the corresponding local value function and $\mathbf{K}_T$ is its initial optimal feedback map. Otherwise, they describe only the objective value and initial control associated with a stationary branch.

Proof. Finite-dimensional implicit differentiation applies for each T ; the common inverse bounds give the

stated uniform estimates. For a variation of x_0 , stationarity, the costate equation, and $p_T = r_x$ give

$$\begin{aligned} D_x \mathcal{S}_T h &= p_T^\top \delta x_T - \lambda \sum_t (\ell_x^\top \delta x_t + \ell_u^\top \delta u_t) \\ &= p_T^\top \delta x_T + \sum_t (p_t^\top \delta x_t - p_{t+1}^\top \delta x_{t+1}) = p_0^\top h. \end{aligned}$$

Equation (5.5) at $t = 0$ proves the $C^{1,1}$ bound. At the reference point $(x, y) = (x_\infty, y_\infty)$, (5.4) also yields a horizon-uniform quadratic expansion of p_0 , and hence a cubic remainder for $\mathcal{S}_T$. $\square$

Theorem 5.5 (Exponential decay of sensitivity to the terminal reward). Assume the explicit hyperbolic transition (2.15) and the graph certificate of Lemma 3.3 at a reference terminal Hessian $S_\star$. Let $\mathcal{X}$ be a convex terminal-state neighborhood of x_∞ and let $\mathcal{R}$ be a class of terminal rewards with a uniform C^3 bound. Fix $0 < \alpha_{\text{ter}} < \gamma$ and assume

$$\sup_{\rho \in \mathcal{R}} \sup_{x \in \mathcal{X}} \|D^2 \rho(x) - S_\star\| \leq \eta. \quad (5.12)$$

Put

$$\begin{aligned} G_\Gamma &= \sup_{T \geq T_0} \|\widehat{\Gamma}_T(S_\star)^{-1}\|, \\ C_X &= \sup_{T \geq T_0} \| [X_s D_s^T X_u] \|, \quad G_\Gamma C_X \eta < 1. \end{aligned}$$

Let $C_\rightarrow$ be the resulting common constant in Lemma 3.3, and let C_N be a common residual constant in Theorem 4.1. Assume that common branch and data radii R and $\delta_\star$ have been chosen so that

$$\theta := 8C_\rightarrow \|L_1^{-1}\| C_N (1 + e^{-\alpha_{\text{ter}}}) R < 1. \quad (5.13)$$

Assume also that the initial-state neighborhood and $\mathcal{R}$ satisfy $d_\rho(x) = \|(x - x_\infty, D\rho(x_\infty) - p_\infty)\| < \delta_\star$ and the resulting terminal states lie in $\mathcal{X}$. Then Theorem 4.1, applied to each endpoint map, gives local branches on the same neighborhood. For the same admissible initial state, every $\rho, \tilde{\rho} \in \mathcal{R}$, $T \geq T_0$, and $0 \leq t < T$,

$$\begin{aligned} \|z_t^\rho - z_t^{\tilde{\rho}}\| + \|u_t^\rho - u_t^{\tilde{\rho}}\| \\ \leq C_{\text{ter}} e^{-\alpha_{\text{ter}}(T-t)} \|D(\rho - \tilde{\rho})\|_{L^\infty(\mathcal{X})}. \end{aligned} \quad (5.14)$$

In particular,

$$\boxed{\|K_T^\rho(x) - K_T^{\tilde{\rho}}(x)\| \leq C_{\text{ter}} e^{-\alpha_{\text{ter}} T} \|\rho - \tilde{\rho}\|_{C^2(\mathcal{X})}}. \quad (5.15)$$

The gradients with respect to the initial state in (5.11) obey the same $e^{-\alpha_{\text{ter}} T}$ bound. The constants are horizon independent for every fixed $0 < \alpha_{\text{ter}} < \gamma$.

Proof. The perturbation of (3.10) caused by $S - S_\star$ is confined to its lower row and has norm at most $C_X \eta$. Neumann's lemma therefore gives the uniform inverse bound $G_\Gamma / (1 - G_\Gamma C_X \eta)$ and hence one common $C_\rightarrow$.

Set $h = z^\rho - z^{\tilde{\rho}}$ and

$$\Sigma_T = \int_0^1 D^2 \rho(x_T^\rho + s(x_T^\rho - x_T^{\tilde{\rho}})) ds.$$

Subtracting the boundary conditions gives

$$P_x h_0 = 0, \quad R_{\Sigma_T} h_T = D(\rho - \tilde{\rho})(x_T^\rho) =: b_T. \quad (5.16)$$

Writing $\tilde{\mathcal{F}} = L + N$, the interior difference becomes $h_{t+1} = M h_t + g_t$, where

$$g_t = -L_1^{-1} \{N(z_t^\rho, z_{t+1}^\rho) - N(z_t^{\tilde{\rho}}, z_{t+1}^{\tilde{\rho}})\}.$$

If $H = \max_t \|h_t\| / v_T(t)$, the residual Lipschitz bound, $w_T \leq 2$, and $v_T(t) = e^{-\alpha_{\text{ter}} t} v_T(t+1)$ imply

$$\max_t \frac{\|g_t\|}{v_T(t+1)} \leq 8 \|L_1^{-1}\| C_N (1 + e^{-\alpha_{\text{ter}}}) R H.$$

Lemma 3.3 and (5.13) now give $H \leq C_\rightarrow \|D(\rho - \tilde{\rho})\|_\infty + \theta H$. Thus $C_z = C_\rightarrow / (1 - \theta)$ controls the state-costate difference, and $C_{\text{ter}} = C_z [1 + L_\nu (1 + e^{\alpha_{\text{ter}}})]$ is admissible, where $L_\nu = \sup \|D\nu\|$. Indeed, the control estimate follows from $u_t = \nu(x_t, p_{t+1})$; at $t = 0$, the common initial state and $v_T(1) = e^{\alpha_{\text{ter}}} e^{-\alpha_{\text{ter}} T}$ give (5.15). The gradient estimate follows from (5.11). $\square$

6 Verification for definite linear-quadratic systems

After absorbing the positive multiplier λ into Q and R , consider the LQ dynamics $x_{t+1} = Ax_t + Bu_t$ with running term $-\frac{1}{2}x_t^\top Qx_t - \frac{1}{2}u_t^\top Ru_t$, $Q \succeq 0$, $R \succ 0$. The Hamiltonian stationarity equation is $u_t = R^{-1}B^\top p_{t+1}$ and the state-costate equations are

$$\begin{aligned} x_{t+1} &= Ax_t + BR^{-1}B^\top p_{t+1}, \\ p_t &= A^\top p_{t+1} - Qx_t. \end{aligned} \quad (6.1)$$

Assume A is invertible, put $G = BR^{-1}B^\top$, and eliminate p_{t+1} to obtain

$$M(A, B, Q, R) = \begin{pmatrix} A + GA^{-\top}Q & GA^{-\top} \\ A^{-\top}Q & A^{-\top} \end{pmatrix}. \quad (6.2)$$

Then $\begin{pmatrix} x_{t+1} \\ p_{t+1} \end{pmatrix} = M(A, B, Q, R) \begin{pmatrix} x_t \\ p_t \end{pmatrix}$, and let $\mathbb{J} = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}$. The factorization

$$M = \begin{pmatrix} I & G \\ 0 & I \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & A^{-\top} \end{pmatrix} \begin{pmatrix} I & 0 \\ Q & I \end{pmatrix}$$

consists of symplectic factors, hence $M^\top \mathbb{J} M = \mathbb{J}$ and the spectrum is reciprocal (Bittanti, Laub, and Willems, 1991).

Proposition 6.1 (Symplectic graph transversality). Let $M \in \text{Sp}(2d, \mathbb{R})$ be hyperbolic. Write bases of its stable and unstable subspaces as

$$\begin{aligned} V_s &= \begin{pmatrix} X_s \\ P_s \end{pmatrix}, & V_u &= \begin{pmatrix} X_u \\ P_u \end{pmatrix}, \\ MV_s &= V_s D_s, & MV_u &= V_u D_u. \end{aligned}$$

For $S = S^\top$, consider the endpoint rows

$$B_0 = \begin{pmatrix} I & 0 \\ 0 & 0 \end{pmatrix}, \quad B_T = \begin{pmatrix} 0 & 0 \\ -S & I \end{pmatrix}, \quad (6.3)$$

which impose $x_0 = \xi$ and $p_T - Sx_T = \pi$. The subspaces E_s, E_u are Lagrangian, and

$$\hat{\Gamma}_\infty = [B_0 V_s \quad B_T V_u] = \begin{pmatrix} X_s & 0 \\ 0 & P_u - SX_u \end{pmatrix}. \quad (6.4)$$

The two blocks are nonsingular exactly when

$$\begin{aligned} E_s \cap (\{0\} \times \mathbb{R}^d) &= \{0\}, \\ E_u \cap \{(x, Sx) : x \in \mathbb{R}^d\} &= \{0\}. \end{aligned} \quad (6.5)$$

Under these graph conditions, the uniform Green property holds for all sufficiently large horizons. For a prescribed T_0 , it holds for every $T \geq T_0$ if the finitely many scaled matrices $\hat{\Gamma}_T$ below a large-horizon cutoff are also nonsingular. The conclusions are open under small perturbations of (M, S) within the hyperbolic symplectic region. The large-horizon conclusion is uniform on compact subsets satisfying the two graph conditions. Uniformity from a prescribed T_0 additionally requires the finitely many short-horizon matrices to remain nonsingular on the compact set.

Proof. For $v, w \in E_s$, symplecticity gives

$$v^\top \mathbb{J} w = (M^k v)^\top \mathbb{J} (M^k w) \longrightarrow 0 \quad (k \rightarrow \infty).$$

Thus E_s is isotropic. Reciprocal spectral symmetry gives $\dim E_s = \dim E_u = d$, so E_s is Lagrangian; applying the same argument to M^{-k} on E_u proves the unstable

assertion. Direct multiplication by (6.3) gives (6.4), and the two kernel conditions are precisely (6.5).

Moreover, $\|\hat{\Gamma}_T - \hat{\Gamma}_\infty\| \leq Ce^{-\gamma T}$. If the diagonal blocks in (6.4) are nonsingular, Neumann's lemma gives a uniform inverse for all sufficiently large T ; Theorem 3.2 then gives the Green property. The finitely many shorter horizons are covered by their assumed nonsingularity. Continuity of hyperbolic spectral projectors and the two graph blocks proves openness and uniformity on compact subsets for the large-horizon conclusion. For fixed T_0 , continuity of the finitely many remaining singular values gives the corresponding assertion for every $T \geq T_0$. $\square$

Recall that (A, B) is stabilizable if $\text{rank}[A - \lambda I \quad B] = d$ for every $|\lambda| \geq 1$. Along every complex solution of (6.1), direct substitution gives the energy identity

$$\begin{aligned} p_{t+1}^* x_{t+1} - p_t^* x_t &= p_{t+1}^* G p_{t+1} + x_t^* Q x_t, \\ G &= BR^{-1}B^\top. \end{aligned} \quad (6.6)$$

Theorem 6.2 (Stabilizable definite LQ data). Let A be invertible, $Q = Q^\top \succ 0$, $R = R^\top \succ 0$, and let (A, B) be stabilizable. Then $M(A, B, Q, R)$ is hyperbolic, X_s is invertible, and for every $S = S^\top \preceq 0$ the block $P_u - SX_u$ is invertible. The rows

$$x_0 = \xi, \quad p_T - Sx_T = \pi$$

satisfy the uniform Green property for every $T \geq 1$. The constants can be chosen uniformly on compact subsets of this data class.

Proof. If $Mz = \lambda z$ with $|\lambda| = 1$, the left side of (6.6) vanishes along $\lambda^t z$. Hence $x = 0$ and $B^\top p = 0$. The second block equation gives $A^\top p = \lambda^{-1} p$, contradicting the PBH condition unless $p = 0$. Thus M is hyperbolic.

If $(0, p) \in E_s$, sum (6.6) forward to infinity. Exponential decay makes both endpoint pairings vanish, hence $x_t = 0$ and $B^\top p_t = 0$. Here $p_{t+1} = A^{-\top} p_t$. If $p \neq 0$, then $V = \text{span}\{p_t : t \geq 1\}$ is nonzero and invariant under $A^{-\top}$. Since $A^{-\top}$ is injective, $A^{-\top} V = V$, so V is also $A^\top$ -invariant. The powers of $A^{-\top}|_V$ converge to zero, hence $A^\top|_V$ has an eigenvalue of modulus greater than one. Its eigenvector is annihilated by $B^\top$, contradicting the PBH condition. Therefore X_s is invertible.

Let $(x_0, Sx_0) \in E_u$ with $S \preceq 0$. Summing the energy identity from $-\infty$ to -1 gives a nonnegative sum equal to $x_0^\top Sx_0 \leq 0$. Thus $x_t = 0$ for $t \leq -1$ and $Gp_0 = 0$; the state equation gives $x_0 = 0$ and then $p_0 = 0$. Hence $P_u - SX_u$ is invertible. By (6.4) and Theorem 3.2, the Green property holds for all sufficiently large horizons.

For a finite homogeneous problem, $x_0 = 0$ and $p_T = Sx_T$. Summing (6.6) gives

$$\sum_{t=0}^{T-1} (p_{t+1}^* G p_{t+1} + x_t^* Q x_t) = x_T^* S x_T \leq 0.$$

The sum is nonnegative, so both sides vanish. Consequently $x_t = 0$, $Gp_T = 0$, then $x_T = 0$, $p_T = 0$, and backward recursion gives $p_t = 0$. Thus every finite-horizon boundary matrix is invertible. The finitely many matrices below the large-horizon cutoff are absorbed into the uniform bound. Continuous dependence of invariant subspaces and finite-horizon matrices gives compact-subset uniformity. $\square$

Corollary 6.3 (Riccati representation and sensitivity to the terminal Hessian). Under Theorem 6.2, let the terminal reward be $r_S(x) = \frac{1}{2}x^\top Sx$ with $S \preceq 0$. For the unconstrained problem (or where its optimizer is interior), strict concavity makes the stationary branch the unique optimal solution of the Bellman problem, with

$$V_T^S(x) = \frac{1}{2}x^\top \mathcal{P}_T(S)x, \quad K_T^S(x) = K_T(S)x.$$

Fix $K \geq 1$ and $\gamma > 0$ such that, for every $k \geq 0$, $\|V_s D_s^k\| \leq K e^{-\gamma k}$ and $\|V_u D_u^{-k}\| \leq K e^{-\gamma k}$. Put

$$\begin{aligned} C_s &= P_s - S X_s, & C_u &= P_u - S X_u, \\ E_T^{\text{Ric}} &= D_u^{-T} C_u^{-1} C_s D_s^T. \end{aligned} \quad (6.7)$$

Then

$$\mathcal{P}_T(S) = (P_s - P_u E_T^{\text{Ric}})(X_s - X_u E_T^{\text{Ric}})^{-1}, \quad (6.8)$$

$$K_T(S) = R^{-1} B^\top A^{-\top} (\mathcal{P}_T(S) + Q). \quad (6.9)$$

For $\mathcal{P}_\infty = P_s X_s^{-1}$ and $K_\infty = R^{-1} B^\top A^{-\top} (\mathcal{P}_\infty + Q)$,

$$\|\mathcal{P}_T(S) - \mathcal{P}_\infty\| + \|K_T(S) - K_\infty\| \leq C e^{-2\gamma T}. \quad (6.10)$$

Uniformly for $S, \tilde{S}$ in compact subsets of $S \preceq 0$,

$$\begin{aligned} \|\mathcal{P}_T(S) - \mathcal{P}_T(\tilde{S})\| + \|K_T(S) - K_T(\tilde{S})\| \\ \leq C e^{-2\gamma T} \|S - \tilde{S}\|. \end{aligned} \quad (6.11)$$

Thus, when two terminal rewards differ only through their Hessians, the induced Riccati matrices and initial gains differ by $O(e^{-2\gamma T})$. Theorem 5.5 gives the one-sided rate for general terminal-gradient perturbations. The Riccati convergence itself is classical (Caines and Mayne, 1970); here it is derived explicitly from the same scaled graph certificate.

Proof. Strict concavity follows from $R \succ 0$, $Q \succeq 0$, and $S \preceq 0$. Write $z_0 = V_s a + V_u D_u^{-T} \eta$ and

$z_T = V_s D_s^T a + V_u \eta$. The terminal graph gives $C_s D_s^T a + C_u \eta = 0$, hence $\eta = -C_u^{-1} C_s D_s^T a$; substitution at $t = 0$ proves (6.8). The costate identity $p_1 = A^{-\top}(p_0 + Qx_0)$ proves (6.9). The dichotomy gives $\|E_T^{\text{Ric}}\| \leq C e^{-2\gamma T}$. Moreover, $(X_s - X_u E_T^{\text{Ric}})^{-1}$ is uniformly bounded: for large T this follows from a Neumann argument around X_s , and the finitely many remaining horizons follow from boundary-value uniqueness (uniformly on the stated compact sets). Therefore

$$\mathcal{P}_T - \mathcal{P}_\infty = (\mathcal{P}_\infty X_u - P_u) E_T^{\text{Ric}} (X_s - X_u E_T^{\text{Ric}})^{-1},$$

which proves (6.10). The map $S \mapsto C_u^{-1} C_s$ is uniformly Lipschitz on the stated compact sets, giving (6.11). $\square$

Corollary 6.4 (Nonlinear terminal condition with a nonpositive reference Hessian). Assume the dynamics and running cost are the LQ data of Theorem 6.2, translated to the stationary reference, and choose a compact U with $0 \in \text{int } U$. If the terminal reward satisfies the hypotheses of Theorem 5.1, with $p_\infty = r_x(x_\infty, y_\infty)$ and

$$r_{xx}(x_\infty, y_\infty) \preceq 0,$$

then the original Pontryagin problem

$$x_0 = x_{\text{in}}, \quad p_T = r_x(x_T, y)$$

has the branch, explicit radius, and sensitivity estimates of Theorem 5.1, with $T_0 = 1$.

Proof. Apply Theorem 6.2 with $S = r_{xx}(x_\infty, y_\infty)$, then apply Theorem 5.1. $\square$

7 Numerical illustration

The finite-dimensional computations below evaluate the stated certificates and nonlinear checks. All coordinates in the example are nondimensional. We use the coupled stabilizable data

$$A = \begin{pmatrix} 1.5 & 0.4 \\ 0 & 0.6 \end{pmatrix}, \quad B = \begin{pmatrix} 1 \\ 0.5 \end{pmatrix}, \quad (7.1)$$

$$Q = \begin{pmatrix} 1 & 0.3 \\ 0.3 & 0.6 \end{pmatrix}, \quad R = 1,$$

with rectangular input. Here $\lambda_{\min}(Q) \approx 0.439$ and $\|[A, Q]\|_2 \approx 0.334$. Since 1.5 is the only eigenvalue of A with modulus at least one and $\sigma_{\min}([A - 1.5I \ B]) \approx 0.980$, the PBH test gives stabilizability, so Theorem 6.2 applies. The matrix (6.2) has eigenvalue moduli $\{0.367, 0.556, 1.798, 2.723\}$, spectral gap $\gamma \approx 0.587$, and, for unit-norm eigenvector bases, $\sigma_{\min}(X_s) \approx 0.188$ and

$\sigma_{\min}(P_u) \approx 0.325$. Thus the canonical graph transversality condition holds with an explicit numerical margin. For shifted rows we take

$$S_{\text{term}} = - \begin{pmatrix} 0.8 & 0.2 \\ 0.2 & 0.5 \end{pmatrix} \preceq 0,$$

so Theorem 6.2 applies. The nonlinear terminal reward is

$$r_{\text{nl}}(x) = \frac{1}{2}x^\top S_{\text{term}}x - (v_{\text{term}}^\top x)^3, \\ v_{\text{term}} = 2^{-1/2}(1, 1)^\top.$$

Direct evaluation gives $\lambda_{\max}(r_{\text{nl},xx}) \leq -0.106$ on $\|x\|_\infty \leq 0.08$. For the nonlinear test we perturb the dynamics to $F(x, u) = Ax + Bu + 0.1\varphi(x)$, where $\varphi(x) = (x_2^2, x_1x_2)$, and solve in the unit direction proportional to $(0, 0, -0.95, -0.30)$. Table 1 reports the scaled Green test, quadratic remainder slopes, boundary conditioning, and the explicit neighborhood and posterior tests. The normalized remainder is

$$\hat{C}(T, s) = \sup_{0 \leq t \leq T} \frac{\|z_t - z_t^{\text{lin}}\|}{w_T(t)^2 s^2}, \quad \alpha = 0.05 < \gamma/2.$$

Across $T = 20, 40, 80, 160$, the scaled boundary condition number is at most 7.725, whereas a reliable column-norm lower bound for the unscaled matrix reaches 7.56×10^{69} at $T = 160$. Fits over $s \in \{3 \times 10^{-4}, 10^{-3}, 3 \times 10^{-3}, 10^{-2}\}$ give remainder slopes between 1.99987 and 2.00713 with $R^2 \geq 0.999997$. The a posteriori test additionally uses $s = 3 \times 10^{-2}$ and $s = 10^{-1}$. To test the decay caused by a terminal reward perturbation, we set $r_s(x) = r_{\text{nl}}(x) + sd^\top x$, with $d = (-0.95, -0.30)^\top / \|(-0.95, -0.30)\|$, and evaluate $|u_0^s - u_0^0|/s$. Fits over $T = 20, 25, 30, 35$ for $s \in \{0.003, 0.01, 0.03\}$ give rates 0.58706–0.58731 with $R^2 > 0.999999$, consistent with $\gamma = 0.58664$. The exact LQ graph calculation for changing the terminal Hessian from 0 to S_{term} gives rate 1.17187 with $R^2 > 0.999999$, consistent with $2\gamma = 1.17328$.

For the finite-dimensional residual $F_T(z; \Delta) = L_T z - E\Delta + N_T(z)$, where $N_T(0) = DN_T(0) = 0$, we also evaluate an explicit direction-independent Newton–Kantorovich data radius. All operator norms in the following calculation are ordinary trajectory infinity norms. With $\mu_T = \|L_T^{-1}\|_\infty$, $c_T = \|L_T^{-1}E\|_\infty$, and $\|DN_T(z) - DN_T(w)\|_\infty \leq K\|z - w\|_\infty$, the condition

$$\|\Delta\|_\infty < \rho_T := \frac{1}{2\mu_T K c_T} \quad (7.2)$$

places the linear predictor in the Kantorovich regime. Indeed, if $d_T^{\text{lin}} = c_T \|\Delta\|_\infty$ and $a_T = \mu_T K d_T^{\text{lin}} < 1/2$, then a Neumann bound at the linear predictor gives

$$\beta_T \leq \frac{\mu_T K (d_T^{\text{lin}})^2}{2(1 - a_T)}, \quad \omega_T \leq \frac{\mu_T K}{1 - a_T},$$

$$\beta_T \omega_T \leq \frac{1}{2} \left(\frac{a_T}{1 - a_T} \right)^2 < \frac{1}{2}.$$

Thus Proposition 5.2, in the ordinary infinity norm, applies. Using 80% of this open threshold gives uniform evaluated radii 2.10×10^{-2} for the canonical row and 1.53×10^{-3} for the nonlinear shifted row over the four tested horizons. These radii are valid for every endpoint direction, unlike the directional sampling. The a posteriori test of Proposition 5.2 succeeds in 44 of 48 runs: all canonical runs and all shifted runs with $s \leq 0.03$ pass; the four shifted $s = 0.1$ runs fail the sufficient inequality and lie outside the stated concavity box; they are included only to examine the sufficient test outside the stated concavity range. All shifted runs with $s \leq 0.03$ remain inside that box. Failure of the test does not imply nonexistence. At $s = 10^{-2}$, the shifted posterior reported admissible uniqueness radius is approximately 6.14×10^{-3} . These are floating-point contraction checks using inverse-norm bounds with a Neumann correction, not interval-verified computer-assisted proofs.

Table 1 reports the numerical inputs behind the reported radii. Here K is a global infinity-norm Lipschitz bound for the Jacobian of the polynomial residual, so Proposition 5.2 allows $R_{\text{tube}} = +\infty$. Also,

$$\rho_T^\circ = \rho_T \text{ in (7.2),} \quad \rho_T^{\text{eval}} = 0.8\rho_T^\circ.$$

Panel (b) reports the directly evaluated posterior quantities $\beta_T^{\text{post}}, \omega_T^{\text{post}}$. There $r_{-,T}$ bounds the distance from the linear predictor to an exact zero, while

$$R_T^{\text{rep}} = \frac{1}{2}(r_{-,T} + 1/\omega_T^{\text{post}}) \quad (7.3)$$

is one concrete radius strictly inside the admissible uniqueness interval, not a maximal radius. The inverse residual in the displayed rows is below 2.1×10^{-15} . The directional posterior test can succeed outside the smaller direction-independent a priori ball because it is centered at the computed predictor for that particular direction.

The additional sweep $S(\theta) = S_{\text{term}} + \theta I$ detects the value at which limiting transversality fails, near $\theta_* \approx 1.23347$; the least singular values there are 4.67×10^{-16} for $\hat{\Gamma}_\infty$ and 3.86×10^{-12} for $T = 20$. The terminal perturbation data, remainder slopes, radii, and pass counts are summarized above.

8 Concluding remarks

The scaled boundary matrix gives a Green estimate that separates both endpoint effects. Its one-sided terminal estimate shows that an error in the terminal reward gradient changes the initial stationary control and the gradient with respect to the initial state by $O(e^{-\alpha_{\text{ter}} T})$, with an explicit local constant. The same analysis gives

Table 1

Representative numerical quantities. Panel (a) evaluates the direction-independent a priori bound at zero. Panel (b) evaluates the a posteriori bound at the linear predictor for $\Delta = s\hat{d}$, $s = 10^{-2}$, $\|\hat{d}\|_2 = 1$, and $\|\hat{d}\|_\infty = 0.95358$. Certificate norms are infinity norms; only $\hat{d}$ is normalized in the Euclidean norm. Entries are upper estimates obtained from a Neumann-series bound.

(a) Direction-independent a priori bound evaluation						
T	row	μ_T	c_T	K	ρ_T^o	ρ_T^{eval}
20	can.	9.6705	3.2807	0.6000	2.6267×10^{-2}	2.1013×10^{-2}
20	shift.	9.6705	3.1803	8.4853	1.9160×10^{-3}	1.5328×10^{-3}
160	can.	9.6706	3.2807	0.6000	2.6266×10^{-2}	2.1013×10^{-2}
160	shift.	9.6706	3.1803	8.4853	1.9160×10^{-3}	1.5328×10^{-3}

(b) A posteriori bound evaluation at $s = 10^{-2}$						
T	row	β_T^{post}	ω_T^{post}	$h_T^{\text{post}} = \beta_T^{\text{post}} \omega_T^{\text{post}}$	$r_{-,T}$	R_T^{rep}
20	can.	6.5069×10^{-6}	5.8023	3.7755×10^{-5}	6.5070×10^{-6}	8.6176×10^{-2}
20	shift.	9.7658×10^{-5}	82.0572	8.0136×10^{-3}	9.8053×10^{-5}	6.1423×10^{-3}
160	can.	6.5069×10^{-6}	5.8024	3.7755×10^{-5}	6.5070×10^{-6}	8.6175×10^{-2}
160	shift.	9.7658×10^{-5}	82.0578	8.0136×10^{-3}	9.8053×10^{-5}	6.1423×10^{-3}

(c1) Finite-horizon boundary diagnostics				(c2) Nonlinear diagnostics		
T	$\ \hat{\Gamma}_T - \hat{\Gamma}_\infty\ _2$	$\sigma_{\min}$ (can.)	$\max \kappa_2$	diagnostic	canonical	shifted
20	5.4×10^{-6}	0.1879	7.72	fitted remainder slope	1.99987	2.00713
40	4.3×10^{-11}	0.1879	7.72	a posteriori test satisfied	24/24	20/24
80	2.8×10^{-21}	0.1879	7.72			
160	1.2×10^{-41}	0.1879	7.72			

horizon-uniform branches, quadratic remainders, and a priori and a posteriori certificates. For LQ systems with invertible A , stabilizable (A, B) , $Q \succ 0$, $R \succ 0$, and a nonpositive terminal Hessian, the certificate holds from $T = 1$, and sensitivity to that Hessian decays at rate $O(e^{-2\gamma T})$. Under strict concavity these stationary solutions are Bellman optimal. Establishing horizon-uniform second-order sufficient conditions for the nonlinear problem remains open.

Data and code availability

No external data were used; all numerical values follow from the equations and parameters in this article.

Declaration of generative AI and AI-assisted technologies in the writing process

The authors used OpenAI Codex in ChatGPT Work for language editing and to draft parts of the numerical code. They independently verified the analytic Jacobian and numerical results and take full responsibility for this article.

References

- D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 4th edition, 2017.
- S. Bittanti, A. J. Laub, and J. C. Willems, editors. *The Riccati Equation*. Springer, 1991. doi: 10.1007/978-3-642-58223-3.
- T. Breiten and L. Pfeiffer. On the turnpike property and the receding-horizon method for linear-quadratic optimal control

- problems. *SIAM Journal on Control and Optimization*, 58(2): 1077–1102, 2020. doi: 10.1137/18M1225811.
- P. E. Caines and D. Q. Mayne. On the discrete time matrix Riccati equation of optimal control. *International Journal of Control*, 12(5):785–794, 1970. doi: 10.1080/00207177008931892.
- W. A. Coppel. *Dichotomies in Stability Theory*. Springer, 1978. doi: 10.1007/BFb0067780.
- L. Grüne and J. Pannek. *Nonlinear Model Predictive Control: Theory and Algorithms*. Springer, 2nd edition, 2017. doi: 10.1007/978-3-319-46024-6.
- L. Grüne, M. Schaller, and A. Schiela. Abstract nonlinear sensitivity and turnpike analysis and an application to semilinear parabolic PDEs. *ESAIM: Control, Optimisation and Calculus of Variations*, 27:56, 2021. doi: 10.1051/cocv/2021030.
- H. Halkin. A maximum principle of the Pontryagin type for systems described by nonlinear difference equations. *SIAM Journal on Control*, 4(1):90–111, 1966.
- K. Kunisch and L. Pfeiffer. The effect of the terminal penalty in receding horizon control for a class of stabilization problems. *ESAIM: Control, Optimisation and Calculus of Variations*, 26: 58, 2020. doi: 10.1051/cocv/2019037.
- S. Na and M. Anitescu. Exponential decay in the sensitivity analysis of nonlinear dynamic programming. *SIAM Journal on Optimization*, 30(2):1527–1554, 2020. doi: 10.1137/19M1265065.
- N. Sakamoto and E. Zuazua. The turnpike property in nonlinear optimal control—a geometric approach. *Automatica*, 134: 109939, 2021. doi: 10.1016/j.automatica.2021.109939.
- M. Sassano. Infinite-horizon optimal control of nonlinear discrete-time systems: HJB PDE, Hamiltonian dynamics and invariant manifolds. *Automatica*, 179:112441, 2025. doi: 10.1016/j.automatica.2025.112441.
- S. Shin and V. M. Zavala. Controllability and observability imply exponential decay of sensitivity in dynamic optimization. *IFAC-PapersOnLine*, 54(6):179–184, 2021. doi: 10.1016/j.ifacol.2021.08.542.
- S. Shin, M. Anitescu, and V. M. Zavala. Exponential decay of sensitivity in graph-structured nonlinear programs. *SIAM Journal on Optimization*, 32(2):1156–1183, 2022. doi: 10.1137/21M1391079.
- W. Xu and M. Anitescu. Exponentially accurate temporal decomposition for long-horizon linear-quadratic dynamic optimization. *SIAM Journal on Optimization*, 28(3):2541–2573, 2018. doi: 10.1137/16M1081993.