PID, LQR, and Integral State-Feedback for Saturated Inverted Pendulum Stabilization

Basis

Research Paper

Abstract

This paper compares PID, LQR, and integral state-feedback controllers for upright stabilization of an inverted pendulum with actuator saturation and external disturbances. All three controllers are synthesized from the same upright linearization and then evaluated on the same nonlinear cart-pendulum model, which keeps the comparison focused on recoverability, disturbance rejection, control effort, and saturation-limited performance rather than on mismatched modeling assumptions. The resulting benchmark supports scenario-specific selection guidance: LQR is the strongest default near the nominal operating point when clipping is limited, integral state-feedback is most compelling when steady-state disturbance rejection is central, and PID remains a useful low-complexity baseline in the smaller recoverable regime.

Keywords. inverted pendulum, PID, LQR, integral state-feedback, actuator saturation, disturbance rejection

Contents

1 Introduction and scope

1.1 Problem Statement and Motivation

The inverted pendulum on a cart is used here as a controlled benchmark for stabilization near the upright equilibrium under open-loop instability, finite actuator authority, and prescribed exogenous disturbances. The study is comparative rather than encyclopedic. Three controller structures, namely a PID baseline, linear quadratic regulation, and integral state-feedback, are synthesized from the same upright operating point and then evaluated on the same saturated nonlinear plant. This restriction is deliberate. It removes several common sources of ambiguity in controller comparison, including changes in plant model, operating point, sensing assumptions, disturbance semantics, and actuator limits. Under those shared conditions, observed differences are interpreted as differences in controller behavior within the benchmark regime rather than as universal rankings of controller classes.

The central engineering question is operational and narrow: under identical nonlinear simulations with actuator clipping and fixed disturbance definitions, which controller retains recoverability over the largest tested set of initial conditions near the upright state, and how do regulation quality and disturbance response change across those same constrained regimes. Local linear stabilization is treated as a prerequisite, but not as the endpoint of the analysis, because the nonlinear plant and the input limit materially affect the practical operating region. This framing is consistent with the manuscript-wide emphasis on recoverable initial-condition sets under saturation rather than on unconstrained linear performance alone.

The introduction fixes the semantics used throughout the manuscript so that later evidence is read consistently. The manuscript does not use one scalar score as a universal decision rule, does not claim global stabilization, and does not infer generic superiority of any controller family beyond the benchmark studied here. Recommendations are conditional on the disturbance class, saturation regime, and recoverability objective represented in the reported nonlinear simulations. Where later chapters compress several metrics into a scorecard or decision table, those summaries remain subordinate to the underlying recoverability, disturbance, and transient records.

1.2 System Definition and Common Information Set

The plant is a planar cart-pendulum system driven by a single horizontal force applied to the cart. The generalized coordinates are

$$q = \begin{bmatrix} x \\ \theta \end{bmatrix},$$

where x denotes cart displacement and θ denotes pendulum angle measured from the upright equilibrium, with $\theta = 0$ corresponding to the unstable balancing configuration. The state vector is

$$z = \begin{bmatrix} x & \dot{x} & \theta & \dot{\theta} \end{bmatrix}^\top.$$

The nonlinear dynamics are represented abstractly as

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + G(q) + F\,\dot{q} = B_u\,u + B_w\,w,$$

with symmetric amplitude saturation imposed on the commanded cart force u. The expanded equations and numerical parameters are given in the modeling chapter. At introduction level, the relevant point is the shared interpretation: all controllers act on the same nonlinear plant and all comparative claims are downstream of that common model.

All controllers are constructed about the upright equilibrium

$$z_e = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix}^\top,$$

with zero nominal input in the disturbance-free case. The corresponding local model is written as

$$\dot{z} = Az + B_u u + B_w w, \qquad y = Cz.$$

The measurement set is shared across the benchmark: x, $\dot{x}$, θ, and $\dot{\theta}$ are available to every controller. The LQR and integral state-feedback designs use the full state directly. The PID baseline uses the same physical

Table 1: Durable evidence surface referenced by the introduction. These artifacts define the provenance for later comparative claims and decision summaries.

Artifact or figure	Role in manuscript interpretation
controller tradeoff summary table and Table/figure renderings derived from it	Canonical comparative summary of recoverability, disturbance, and transient tradeoffs across the shared benchmark.
controller tradeoff scatter data	Pointwise tradeoff surface used to support benchmark-conditioned comparisons rather than universal rankings.
robust controller comparison summary and corresponding figure files	Compact summary of controller behavior under the shared nonlinear stress tests.

information through regulation errors referenced to $x_{\text{ref}} = 0$ and $\theta_{\text{ref}} = 0$, together with derivative information formed from the measured state. The comparison is therefore not confounded by observer design or unequal sensing assumptions.

The integral controller is defined narrowly because later disturbance statements depend on that definition. The integral state-feedback controller augments the plant with integrators on the benchmark regulation errors associated with cart position and pendulum angle. Any statement about offset reduction or bias rejection is therefore interpreted only with respect to these regulated outputs and to the disturbance channels defined below. No broader claim about arbitrary disturbance rejection is intended.

1.3 Benchmark Scenarios and Disturbance Semantics

The manuscript uses one disturbance formalism throughout. Two disturbance scenarios are included. The first is a constant matched disturbance entering through the cart-force channel as an additive bias. In the local linear description, this disturbance is aligned with the control-input direction. The second is an impulse-like angular disturbance represented as a brief external torque acting on the pendulum coordinate. These scenarios are intentionally limited. They separate persistent bias entering through actuation from transient recovery after a short rotational upset, and they are sufficient for the comparative claims actually supported by the benchmark records.

The common nonlinear evaluation surface combines three ingredients: nonzero initial angular displacement, finite actuator limit, and one of the prescribed disturbance scenarios. Within that surface, the primary empirical object is the recoverable initial-condition envelope, meaning the tested set of initial states from which the closed loop returns to the upright operating region without violating the benchmark failure criteria. That envelope is an empirical object obtained from nonlinear simulation under saturation. It is not presented as a proof of attraction and it should not be read as a global region of attraction.

This choice of primary metric resolves an important interpretive point. A controller can be locally stabilizing in the linearized model and still lose recoverability earlier on the nonlinear saturated plant when demanded force clips for a sustained interval. Conversely, a controller that improves offset behavior under the constant matched disturbance can do so while changing transient behavior or shrinking the recoverable set under tight force limits. For that reason, the manuscript treats recoverability under saturation as the main comparison lens and reads disturbance response and transient quality as coordinated secondary criteria.

1.4 Interpretive Scope and Decision Semantics

Because the benchmark holds the synthesis basis fixed and then validates all controllers on one saturated nonlinear plant, the manuscript should be read as a controlled local comparison with nonlinear stress testing. It does not establish global attraction, does not treat swing-up control, and does not study partial-state measurement, switching logic, observer design, sensor noise, estimator error, or broad parametric uncertainty. These exclusions are substantive rather than incidental. They keep the evidentiary scope aligned with the engineering question actually answered by the simulations. For the same reason, the manuscript does not convert the present benchmark into a claim about deployment robustness under unmodeled sensing or estimation effects.

The same caution applies to statements about the three controller families. The introduction does not assign a generic structural advantage to LQR, PID, or integral state-feedback outside the shared setup. Any comparative statement later in the manuscript is benchmark-specific and derives from the reported simulation summaries, figures, and tables. In particular, score-based summaries are interpreted only as compact descriptions of observed tradeoffs under the chosen benchmark metrics. They do not replace the underlying recoverability, disturbance, and transient records, and they are not treated as autonomous proof of superiority.

The manuscript-wide selection logic is therefore conditional and operational. Let R denote the empirical recoverable set under the relevant actuator limit and disturbance case, let J_{off} denote the residual offset measure under the constant matched disturbance, and let J_{tr} denote the transient-quality indicators reported later. Let $R_{\min}$ be the minimum recoverability requirement for the intended initial-condition range, and let $J_{\mathrm{off,max}}$ and $J_{\mathrm{tr,max}}$ denote acceptable bounds for offset and transient degradation under the same scenario. A controller is admissible for a regime only when its observed benchmark outcomes satisfy

$$R \supseteq R_{\min}, \qquad J_{\mathrm{off}} \leq J_{\mathrm{off,max}}, \qquad J_{\mathrm{tr}} \leq J_{\mathrm{tr,max}}.$$

Within the set of admissible controllers, later comparison tables summarize the observed tradeoff rather than asserting universal dominance. This keeps the decision semantics falsifiable: recoverability remains primary, disturbance response is interpreted with respect to the specified disturbance class, and transient advantages are read only within the same actuator limit and initial-condition envelope.

1.5 Scope of Contribution

The contribution is precise and limited. The paper does not introduce a new control law. It provides a unified benchmark in which three established controller structures are compared on the same nonlinear inverted-pendulum regulation problem, under common sensing assumptions, common disturbance definitions, common actuator limits, and common failure semantics. Within that benchmark, the manuscript supports comparative statements about recoverability near the upright state, offset behavior under a constant matched disturbance, and regime-dependent controller selection under saturation. Claims outside that scope are not made.

The practical output is therefore benchmark-conditioned guidance rather than a universal controller ranking. The manuscript identifies how recoverability, disturbance attenuation for the matched constant bias, and transient behavior align or trade off against one another within the tested regimes summarized by the durable artifacts listed above. This narrower interpretation is the one intended throughout the paper, and the introduction is written to keep later result and recommendation language consistent with that evidentiary scope.

2 Nonlinear model and upright linearization

2.1 Configuration, notation, and modeling assumptions

We use a standard planar cart-pendulum model with horizontal force actuation. The cart position is denoted by x, the pendulum angle by θ, and the control input by u. The angle convention is chosen so that $\theta = 0$ corresponds to the upright equilibrium. A positive constant disturbance force d acts through the same channel as the actuator, which makes the disturbance matched with the control input. This choice is consistent with the disturbance-rejection tests used later in the paper, where the cart is subject to sustained or impulsive horizontal loading while the pendulum must remain stabilized.

The physical parameters are the cart mass $M > 0$, the pendulum mass $m > 0$, the distance from pivot to pendulum center of mass $l > 0$, the pendulum moment of inertia about its center of mass $I \geq 0$, the gravitational constant $g > 0$, and a viscous cart-friction coefficient $b \geq 0$. No Coulomb friction, pivot friction, actuator dynamics, sensor dynamics, or structural flexibility are introduced in this benchmark model. That restriction is deliberate: the purpose of the chapter is to define a single local design model that can be used consistently across all compared controllers, while the later simulations still evaluate the resulting controllers on the full nonlinear plant with actuator saturation.

The measurable state is taken as

$$\xi = \begin{bmatrix} x & \dot{x} & \theta & \dot{\theta} \end{bmatrix}^{\top}.$$

This full-state convention is the common basis for the state-feedback designs. It also provides a neutral reference point for the PID benchmark, since the comparison later in the paper concerns closed-loop behavior on the same plant rather than differences induced by incompatible sensing assumptions. The actuator saturation nonlinearity is not absorbed into the state equations below; instead, the plant dynamics are written for the unconstrained input force, and saturation is applied externally in the simulation loop. Separating the plant from the actuator nonlinearity keeps the local linearization transparent and avoids mixing a smooth model derivation with a piecewise input map.

Two modeling assumptions should be stated explicitly. First, the chapter adopts the usual rigid-body approximation for the cart and pendulum. Second, the derivation treats the rail as horizontal and the motion as strictly planar. These assumptions are standard for the inverted-pendulum benchmark, but they also determine the meaning of the later performance conclusions: the comparison is about controller behavior for this canonical nonlinear plant, not about robustness to unmodeled three-dimensional effects, backlash, or actuator bandwidth limitations.

2.2 Exact nonlinear dynamics

The generalized coordinates are $q = [x, \theta]^\top$. With the above sign convention, the cart equation expresses horizontal force balance and the pendulum equation expresses rotational balance about the pivot. A convenient form of the coupled nonlinear model is

$$(M+m)\ddot{x} + b\dot{x} + ml\cos\theta\,\ddot{\theta} - ml\sin\theta\,\dot{\theta}^2 = u + d, \tag{1}$$

$$ml\cos\theta\,\ddot{x} + (I+ml^2)\ddot{\theta} - mgl\sin\theta = 0. \tag{2}$$

The first equation shows the direct action of the applied horizontal force on the cart and the coupling terms generated by pendulum acceleration and centripetal effects. The second equation shows that gravity enters through $\sin\theta$, which vanishes at the upright equilibrium but produces the well-known local destabilizing term after linearization.

For compactness, the equations can be written in matrix form as

$$D(\theta)\begin{bmatrix}\ddot{x}\\ \ddot{\theta}\end{bmatrix} + \begin{bmatrix}b\dot{x} - ml\sin\theta\,\dot{\theta}^2\\ -mgl\sin\theta\end{bmatrix} = \begin{bmatrix}\nu\\ 0\end{bmatrix}, \qquad \nu := u + d, \tag{3}$$

with

$$D(\theta) = \begin{bmatrix}M+m & ml\cos\theta\\ ml\cos\theta & I+ml^2\end{bmatrix}.$$

The determinant of the inertia matrix is

$$\det D(\theta) = (M+m)(I+ml^2) - m^2l^2\cos^2\theta.$$

Under the physical assumptions $M > 0$, $m > 0$, $l > 0$, and $I \geq 0$, this determinant is strictly positive for all θ of interest in the present study. In particular,

$$\det D(\theta) \geq (M+m)(I+ml^2) - m^2l^2 = MI + Mml^2 + mI > 0,$$

so the acceleration vector is well defined without singular inversion of the mass matrix. This property is useful in simulation because it guarantees that the nonlinear state equation remains explicit after solving for $\ddot{x}$ and $\ddot{\theta}$.

Solving (1) and (2) for the accelerations gives

$$\ddot{x} = \frac{(I+ml^2)\big(\nu - b\dot{x} + ml\sin\theta\,\dot{\theta}^2\big) - m^2l^2g\sin\theta\cos\theta}{\Delta(\theta)}, \tag{4}$$

$$\ddot{\theta} = \frac{-ml\cos\theta\,\big(\nu - b\dot{x} + ml\sin\theta\,\dot{\theta}^2\big) + (M+m)mgl\sin\theta}{\Delta(\theta)}, \tag{5}$$

where

$$\Delta(\theta) = (M+m)(I+ml^2) - m^2l^2\cos^2\theta.$$

Equations (4) and (5) make two structural features explicit. First, the disturbance force enters through exactly the same channel as the control force. This is the reason that constant-force disturbances can be rejected by augmenting the linear model with integral action. Second, the nonlinear coupling contains both a gravitational term proportional to $\sin\theta$ and a velocity-dependent term proportional to $\sin\theta\,\dot{\theta}^2$. The latter disappears in the local model and becomes relevant mainly when trajectories leave the immediate neighborhood of the upright operating point.

The corresponding first-order nonlinear state equation is

$$\dot{\xi} = f(\xi, \nu) = \begin{bmatrix} \dot{x} \\ \ddot{x}(x, \dot{x}, \theta, \dot{\theta}, \nu) \\ \dot{\theta} \\ \ddot{\theta}(x, \dot{x}, \theta, \dot{\theta}, \nu) \end{bmatrix}, \tag{6}$$

with $\ddot{x}$ and $\ddot{\theta}$ given by (4) and (5). This nonlinear model is the plant used for the comparative simulations discussed later. The controllers are not judged only on the local linearized equations; they are synthesized from the same upright linearization and then exercised on the full nonlinear plant with actuator limits.

A comment on the sign convention is useful here. Because $\theta = 0$ denotes the upright position, the gravitational term in the rotational equation appears as $-mgl\sin\theta$ on the left-hand side of (2). After linearization this yields a positive coefficient on θ in the $\ddot{\theta}$ equation, which is the local signature of open-loop instability about the upright equilibrium. This is the expected behavior and provides a direct algebraic check on the consistency of the model.

2.3 Upright equilibrium and local linearization

The common synthesis point for all controllers is the upright resting equilibrium

$$\xi_e = \begin{bmatrix} 0 & 0 & 0 & 0 \end{bmatrix}^\top, \qquad u_e = 0, \qquad d_e = 0.$$

More generally, any constant cart position $x = x_0$ with zero velocities, zero angle, and zero applied force is an equilibrium because the translational coordinate is free. For controller design it is therefore natural to regulate deviations from an arbitrary reference cart position rather than an absolute origin. The local stabilization problem addressed in the paper concerns the neighborhood of one such equilibrium, with x interpreted as displacement from the chosen reference.

Linearization about (ξ_e, u_e, d_e) uses the standard small-angle approximations

$$\sin\theta \approx \theta, \qquad \cos\theta \approx 1, \qquad \theta\dot{\theta}^2 \approx 0,$$

and neglects all terms of second order and higher in the state and input deviations. Define the constant quantity

$$p := (M + m)(I + ml^2) - m^2 l^2 = MI + Mml^2 + mI. \tag{7}$$

The positivity of p follows directly from the physical parameter assumptions and guarantees that the upright linearization is well defined.

Applying the first-order approximation to (1) and (2) gives the coupled linear second-order equations

$$(M + m)\ddot{x} + b\dot{x} + ml\ddot{\theta} = u + d, \tag{8}$$

$$ml\ddot{x} + (I + ml^2)\ddot{\theta} - mgl\theta = 0. \tag{9}$$

Solving this pair for the accelerations yields

$$\ddot{x} = -\frac{(I + ml^2)b}{p}\dot{x} - \frac{m^2 l^2 g}{p}\theta + \frac{I + ml^2}{p}(u + d), \tag{10}$$

$$\ddot{\theta} = \frac{mlb}{p}\dot{x} + \frac{(M + m)mgl}{p}\theta - \frac{ml}{p}(u + d). \tag{11}$$

Several points are worth noting. The translational subsystem contains an unstable coupling through θ, and the rotational subsystem contains the positive gravitational coefficient $\frac{(M+m)mgl}{p}$ that makes the upright

configuration open-loop unstable. Cart friction contributes a damping term to $\ddot{x}$ and, through coupling, a secondary term to $\ddot{\theta}$. The disturbance enters with the same coefficients as the control input in both equations, which means that the linearized disturbance channel is matched as well.

Using the state vector $\xi = [x, \dot{x}, \theta, \dot{\theta}]^\top$, the local dynamics can be written in state-space form as

$$\dot{\xi} = A\xi + Bu + Ed, \tag{12}$$

with

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & -\dfrac{(I+ml^2)b}{p} & -\dfrac{m^2l^2g}{p} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \dfrac{mlb}{p} & \dfrac{(M+m)mgl}{p} & 0 \end{bmatrix},$$

$$B = \begin{bmatrix} 0 \\ \dfrac{I+ml^2}{p} \\ 0 \\ -\dfrac{ml}{p} \end{bmatrix}, \qquad E = \begin{bmatrix} 0 \\ \dfrac{I+ml^2}{p} \\ 0 \\ -\dfrac{ml}{p} \end{bmatrix}.$$

Thus $E = B$, reflecting the matched-force assumption already visible in the nonlinear equations. The measured output used for state-feedback synthesis is the full state itself, while reduced output selections can be extracted from the same model by introducing an output equation $y = C\xi$ as needed for a particular controller structure.

The paper adopts this single upright linearization as the common benchmark design model for all three controllers. The point of the comparison is not to argue that a single linear model is globally accurate. The point is that a fixed local model makes differences in closed-loop behavior attributable to the controller structures rather than to inconsistent synthesis assumptions. In particular, PID, LQR, and integral state-feedback are all tied to the same operating point, the same state definitions, and the same force-actuated plant.

The local validity of (12) is limited by the neglected nonlinear terms. The approximation deteriorates as $|\theta|$ grows, as $|\dot{\theta}|$ grows, or as actuator saturation forces the closed loop into regimes where the nominal linear input cannot be delivered. This limitation is not a defect of the chapter; it is the reason later sections return to the nonlinear equations for performance assessment. The linear model is a synthesis tool, whereas the nonlinear saturated plant is the evaluation environment.

2.4 Augmented model for disturbance rejection

Because the disturbance force is constant or slowly varying in one part of the test matrix, it is useful to state the standard integral augmentation directly at the model level. Let the cart-position tracking error relative to a constant reference r be

$$e_x := x - r.$$

Introduce an integral state

$$\eta(t) = \int_0^t e_x(\tau)\, d\tau, \qquad \dot{\eta} = e_x = x - r.$$

For regulation about a fixed reference, the augmented state may be written as

$$\bar{\xi} = \begin{bmatrix} \eta & x & \dot{x} & \theta & \dot{\theta} \end{bmatrix}^\top.$$

With $r = 0$ for local analysis, the augmented linear model becomes

$$\dot{\bar{\xi}} = \bar{A}\bar{\xi} + \bar{B}u + \bar{E}d, \tag{13}$$

where

$$\bar{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -\dfrac{(I+ml^2)b}{p} & -\dfrac{m^2l^2g}{p} & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & \dfrac{mlb}{p} & \dfrac{(M+m)mgl}{p} & 0 \end{bmatrix},$$

$$\bar{B} = \begin{bmatrix} 0 \\ 0 \\ \dfrac{I+ml^2}{p} \\ 0 \\ -\dfrac{ml}{p} \end{bmatrix}, \qquad \bar{E} = \begin{bmatrix} 0 \\ 0 \\ \dfrac{I+ml^2}{p} \\ 0 \\ -\dfrac{ml}{p} \end{bmatrix}.$$

This augmentation is the natural linear counterpart of rejecting a constant matched disturbance. At steady state with successful regulation, $\dot{\eta} = 0$ implies $x = r$, while the integral channel supplies the constant bias required to offset d. The same principle can also be interpreted from the original equations: in the presence of a constant horizontal load, a purely proportional state feedback generally leaves a nonzero cart-position offset, whereas an added integrator can generate the extra constant force needed to remove that bias.

The role of the augmented model in the paper is narrow and specific. It is not introduced as a general output-regulation framework, and no claim is made here about disturbance rejection beyond the matched constant-force case represented by d. The later comparison only requires the basic observation that integral action is structurally well aligned with this disturbance channel. That observation comes directly from (13) and does not depend on a particular tuning algorithm.

2.5 Interpretation for the common benchmark study

The chapter has defined the modeling layer on which the rest of the comparison is built. Three design choices are especially important for interpreting the later results.

First, all controllers are anchored to the same local operating point. This matters because the inverted pendulum is strongly nonlinear away from the upright equilibrium, and controller comparisons can be distorted when one design is tuned on a different approximation than another. By using a common upright linearization, the study isolates the effect of controller structure more cleanly. Differences observed later under the same nonlinear simulations can therefore be read as differences in how each controller exploits, or fails to exploit, the same local plant information.

Second, actuator limits are treated as part of the evaluation problem rather than the derivation of the nominal model. This separation is useful for two reasons. It keeps the linear design model analytically simple, and it exposes the extent to which each controller retains acceptable behavior when the nominal control law is clipped by the actuator. Since the saturation is external to (12), any degradation observed under tight limits should be interpreted as a closed-loop interaction between the controller and the saturated nonlinear plant, not as a change in the underlying benchmark model.

Third, the disturbance channel has been chosen to be matched and physically interpretable. A horizontal force disturbance is one of the simplest load models for a cart-pendulum, and its matched nature creates a meaningful test of integral action. The comparison later in the paper therefore has a clear local mechanism: state feedback without integration stabilizes the upright equilibrium but does not in general eliminate position bias under a constant matched force, while an integral state can compensate that bias when the closed loop remains within the region where the augmented linear design remains relevant.

These observations also identify the limits of what this chapter establishes. The derivation proves neither global stabilization nor robustness to arbitrary uncertainties. It provides an exact nonlinear plant model, a mathematically consistent upright linearization, and an augmented linear form tailored to the disturbance scenario studied in the experiments. Claims about relative controller performance are therefore deferred to the later nonlinear simulation results. The present section supports those claims by ensuring that the benchmark plant, coordinates, equilibrium, and local state-space model are explicit and internally consistent.

A final remark concerns local versus nonlocal interpretations. The unstable upright mode, the matched disturbance channel, and the structure of the input matrix B are all local facts derived from first principles. They justify the use of PID, LQR, and integral state-feedback as comparable benchmark designs. By contrast, recovery from large initial angles, sensitivity to saturation, and the onset of performance loss outside the linear region are empirical questions for the nonlinear simulations. Maintaining that distinction keeps the paper's later conclusions properly grounded: the linearization explains why the chosen controllers are reasonable candidates, while the nonlinear tests determine how far that local design picture carries under the imposed operating constraints.

For that reason, later result statements should be read together with their recorded computational provenance rather than as consequences of the local model alone. In this project, the principal provenance anchors for the comparison layer are the summary and regime artifacts *controller tradeoff summary table record*, *controller selection regime table record*, and *controller operating region map record*. The role of the present chapter is narrower: it defines the common nonlinear plant and the shared upright design model that those compute-backed comparisons instantiate.

3 Common controller synthesis

The comparison in this study is organized around a single synthesis principle: every controller is designed from the same local model of the upright equilibrium, and every design is then evaluated on the same nonlinear cart–pendulum dynamics with the same actuator saturation and disturbance scenarios. This choice is narrower than a broad survey of possible stabilizers, but it provides a controlled basis for interpretation. Differences observed later in the nonlinear benchmarks can therefore be attributed primarily to controller structure, integral augmentation, and sensitivity to saturation, rather than to differences in the plant model used during design.

A common synthesis basis is particularly important for the inverted pendulum because the plant is simultaneously underactuated, unstable at the operating point of interest, and strongly nonlinear away from that operating point. Direct comparison across unrelated designs would mix several effects: one controller might be advantaged by a more favorable model, another by looser state assumptions, and another by implicit disturbance accommodation embedded in the design stage. The present chapter removes that ambiguity by fixing the operating point, the measured state vector, and the treatment of the actuator as shared design premises. The resulting comparison is local in the synthesis step and nonlinear in the validation step.

Throughout this chapter the state is written as

$$z = \begin{bmatrix} x & \dot{x} & \theta & \dot{\theta} \end{bmatrix}^\top,$$

where x denotes cart position and θ denotes pendulum angle measured relative to the upright equilibrium. The control input is the horizontal cart force u, and the measured output set is taken to be the full state. This full-state measurement assumption is part of the common benchmark. It removes observer design from the comparison so that the paper isolates differences among feedback laws rather than differences among estimators.

3.1 Shared synthesis objectives and benchmark assumptions

The synthesis problem is formulated around local stabilization of the upright equilibrium together with two practical requirements that appear repeatedly in the nonlinear study: bounded actuation and rejection of persistent disturbances. The first requirement reflects the fact that the cart force is not unbounded in the benchmark simulations. The second reflects the fact that a controller that stabilizes the nominal plant may still exhibit a steady bias under constant disturbances. These two requirements separate the three controller families in a meaningful way. The PID law introduces integral action through an output-error channel, the LQR law optimizes a linear quadratic surrogate without inherent disturbance integral compensation, and the integral state-feedback law augments the plant with an explicit integral state before full-state feedback synthesis.

The benchmark is intentionally local at the design stage. Each controller is synthesized from the upright linearization rather than from a global nonlinear optimal control problem. This restriction has two methodological advantages. First, it keeps the information available to each design comparable. Second, it allows the

paper to interpret failure modes in a concrete way: loss of recovery for large initial angles can be attributed to the limited region over which the local design remains effective, and nonzero equilibrium offsets under sustained load can be attributed to the absence of integral compensation in the loop.

This local scope also imposes limits on what the chapter claims. The synthesis constructions below do not establish global asymptotic stabilization of the nonlinear pendulum. They establish a common linear benchmark from which local closed-loop properties can be derived and from which nonlinear performance can be assessed empirically. The supporting computational checks available in the project confirm controllability and feasibility of the augmented linear designs, and they also support local agreement between the nonlinear dynamics and the upright linear model over a limited neighborhood. Accordingly, the design claims in this chapter are restricted to local synthesis validity and to the fairness of the comparison setup.

A further benchmark assumption is that the same actuator saturation map is applied after the controller output is computed. This matters because unconstrained linear feedback can request forces that the simulated actuator cannot deliver. Applying the same saturation nonlinearity across all controllers ensures that the nonlinear comparison later in the paper reflects the same physical force limitation. In particular, it prevents the implicit use of unrealistic actuation from becoming a hidden design advantage.

The chapter therefore uses the following common synthesis interpretation. First, derive a linear model around the upright equilibrium. Second, construct three feedback laws from that model under the same state definitions and the same sign conventions. Third, impose the same force saturation in nonlinear simulation. Fourth, interpret the resulting closed-loop behavior in terms of local linear design structure, not as a claim of universal dominance by any one controller. This interpretation is consistent with the scope of the available evidence.

3.2 Upright linearization as the common design model

The nonlinear cart–pendulum equations are linearized about the equilibrium

$$(x, \dot{x}, \theta, \dot{\theta}, u) = (0, 0, 0, 0, 0),$$

with $\theta = 0$ denoting the upright position. The resulting model has the standard continuous-time state-space form

$$\dot{z} = Az + Bu,$$

where $A \in \mathbb{R}^{4\times4}$ and $B \in \mathbb{R}^{4\times1}$ are obtained from the Jacobian of the nonlinear dynamics at the equilibrium. The exact symbolic derivation of this linearization is handled elsewhere in the manuscript, and the computational sanity checks reported in the project artifacts support the consistency of the symbolic and numerical forms. For the synthesis chapter, the key point is not the individual matrix entry values, but the structural role of the model: it is the sole linear plant used to construct all three controllers.

The linearized model serves three purposes simultaneously. It provides the transfer path used to define the PID design channel. It provides the state-space pair (A, B) used in LQR synthesis. It also provides the base plant for integral augmentation, where an additional state accumulates a regulated output error. Because all three designs inherit the same unstable open-loop dynamics, the comparison is controlled at the level of plant information.

The local nature of the model should be stated explicitly. Near the upright equilibrium, the linearization captures the dominant unstable mode and the coupling between cart motion and pendulum angle. As the angle departs from the equilibrium, neglected higher-order terms become more important, and the linear prediction degrades. This does not invalidate the synthesis strategy, but it constrains the interpretation of the results. When the nonlinear simulations later identify a recoverable initial-angle envelope, that envelope is not an intrinsic property of the plant alone. It is a property of the nonlinear closed loop induced by a controller designed from this local approximation.

For later use it is convenient to distinguish between regulated variables and measured variables. The measured state is the full vector z. The regulated objective is narrower: maintain the pendulum near upright while keeping the cart near its nominal position and rejecting sustained disturbances. In practical terms, the synthesis gives priority to angular stabilization because the upright equilibrium is open-loop unstable, but it does not ignore cart displacement because successful balancing by unbounded cart drift would not satisfy the benchmark objective. This dual role of x and θ motivates both the LQR state weighting philosophy and the augmented integral design.

The state-space model also provides the natural language for comparing controller structure. PID is a dynamic compensator attached to a selected error signal. LQR is static full-state feedback on the four-state benchmark model. Integral state-feedback enlarges the benchmark to a five-state model by appending an error integrator. Framing the designs this way makes clear that the paper is not comparing unrelated algorithmic families under loosely matched assumptions. It is comparing three specific ways of closing the loop around the same local plant.

The common linearization also standardizes sign conventions. A positive angular error, a positive cart displacement, and a positive control force are interpreted consistently across the chapter. This may appear minor, but it is essential for reproducibility of a comparative study, particularly when one controller is synthesized from a transfer-channel intuition and another from matrix optimization. The shared sign convention ensures that differences in transient behavior are not artifacts of inconsistent feedback polarity.

3.3 PID design from the common benchmark

The PID controller is included because it represents a simple and widely used structure that can stabilize an unstable mode in a limited operating region while also offering an intuitive route to disturbance compensation through integral action. In the present benchmark, the PID design is not introduced as a universal solution to the inverted pendulum problem. It is introduced as a structured low-order baseline synthesized from the same upright model used by the state-feedback designs.

To keep the comparison coherent, the PID law is defined around a single regulated signal derived from the upright linearization. Let $e(t)$ denote the selected regulation error. The controller takes the standard form

$$u(t) = K_{\mathrm{P}} e(t) + K_{\mathrm{I}} \int_0^t e(\tau)\, d\tau + K_{\mathrm{D}} \dot{e}(t),$$

and the actuator command is then obtained by the common saturation map,

$$u_{\mathrm{sat}}(t) = \mathrm{sat}(\nu(t)).$$

In this chapter the notation ν is used for the unsaturated controller output and $u = \nu_{sat}$ for the applied plant input. This separation is useful because it makes explicit where nonlinearity enters the loop. The PID synthesis is performed on the linear model, while the applied nonlinear simulations use the saturated signal.

The particular choice of regulated error is part of the benchmark definition. Because the plant objective is upright balance with bounded cart motion, a purely angular PID law would omit information about cart drift, while a purely positional PID law would fail to address the unstable pendulum mode directly. The common benchmark therefore interprets PID as a low-order controller centered on the dominant balancing objective, with cart-position behavior judged in the subsequent nonlinear evaluation rather than encoded through a full multivariable state penalty. This design choice is one reason the PID controller remains a meaningful comparator rather than a hidden approximation to state feedback.

Integral action is the key structural feature of the PID design in this study. A constant matched disturbance enters the force channel as an additive load. In a locally stable closed loop, the proportional and derivative terms shape transient response but do not, by themselves, guarantee zero steady-state regulation error under a constant bias. The integral term addresses exactly this deficiency by integrating persistent error and driving the equilibrium control demand to the level required to offset the disturbance. The project artifacts include a dedicated bias check for the PID case, and the synthesis interpretation used here is consistent with that role: PID is the simplest among the three benchmark controllers that contains a built-in mechanism for eliminating steady bias in its regulated channel.

This advantage comes with familiar qualifications that are especially relevant for an unstable plant with input saturation. Large integral accumulation can continue while the actuator is saturated, and the internal controller state can then drive a delayed or overly aggressive recovery once the commanded force re-enters the admissible range. For that reason, the nonlinear benchmark later in the paper is not interpreted through nominal linear poles alone. Saturation-aware behavior is part of the observed controller character. The common synthesis setup makes this comparison fair because the same saturation acts on all controllers, but it does not erase the structural susceptibility of PID to windup-like effects.

From the standpoint of this chapter, the principal reason to retain PID is not that it is expected to dominate a full-state design in all regimes. The reason is methodological. PID provides a low-order, interpretable

controller with integral disturbance accommodation, synthesized from the same operating point as the more model-based alternatives. When it performs well, the result shows that modest structure can suffice in a local region. When it performs poorly, the failure can be traced to the limited information and limited structural flexibility of the controller rather than to an unfairly disadvantaged benchmark.

3.4 LQR state feedback on the four-state linear model

The LQR controller is synthesized directly from the shared state-space model (A, B). With full-state measurement available, the feedback law takes the form

$$u(t) = -Kz(t),$$

where the gain matrix K is obtained from the continuous-time algebraic Riccati equation associated with positive semidefinite state weighting and positive definite input weighting. The role of the LQR design in the benchmark is to provide a systematic state-feedback baseline that balances stabilization effort against control usage within a standard quadratic surrogate.

The appeal of LQR in this context is twofold. First, it uses the complete local state vector rather than a single regulation error, so it can directly shape the coupled cart and pendulum dynamics. Second, it yields a controller that is internally simple once the gain is computed: online implementation requires only a linear state combination followed by the same saturation map used for the other designs. This makes LQR a natural midpoint in the comparison, more structured and more model-informed than PID, but still grounded in the same local linear benchmark.

The quadratic performance interpretation is useful but should not be overstated. The LQR gain is optimal only for the unconstrained linear model under the chosen quadratic cost. The actual comparison plant is nonlinear and saturating, and the disturbance scenarios include persistent loads. Therefore the later nonlinear results are not evidence of optimality in the original LQR sense. They are evidence of how an LQR-designed local feedback law behaves when deployed on the common nonlinear benchmark. This distinction matters because actuator saturation can substantially alter closed-loop transients and even reduce the practical recovery region relative to what a purely linear analysis suggests.

Within this benchmark, LQR has a predictable structural limitation relative to the integral designs. For a constant disturbance entering the input channel, a stabilizing LQR law generally settles to a shifted equilibrium unless the regulated output happens to be insensitive to that bias or unless an external precompensation mechanism is added. Because the present comparison keeps the synthesis premises uniform and does not add separate disturbance estimators or feedforward terms to the LQR case, the LQR controller is interpreted as a nominal stabilizer rather than as a zero-bias disturbance-rejection design. This limitation is not a defect of the benchmark. It is one of the central structural contrasts that the paper aims to expose.

The computational evidence available in the project supports the feasibility of the LQR synthesis from the common linearization. In particular, controllability and augmented-design feasibility checks indicate that the linearized plant contains sufficient control authority in the local model to support the state-feedback constructions used in the study. That evidence justifies presenting LQR as an executed synthesis component rather than as a merely conceptual comparator.

From a practical perspective, the LQR design also provides a clean reference for interpreting saturation effects. Because its unsaturated control law is linear in the full state, any loss of performance under stricter force limits can be read against a transparent nominal structure. This is useful for the later selection guidance. A controller that performs strongly in the unconstrained local model may still degrade rapidly when the admissible control magnitude is reduced. The LQR baseline helps identify whether such degradation is mild and graceful or abrupt and region-limiting.

3.5 Integral state-feedback via plant augmentation

The integral state-feedback design extends the four-state benchmark model by appending an integrator driven by a regulated output error. Let the regulated output be written as

$$y = C_r z,$$

where C_r selects the variable, or linear combination of variables, whose steady regulation is required in the presence of constant disturbances. The augmented state is

$$z_a = \begin{bmatrix} z \\ \eta \end{bmatrix}, \qquad \dot{\eta} = r - y,$$

with reference $r = 0$ for upright stabilization. The resulting augmented dynamics are

$$\dot{z}_a = A_a z_a + B_a u,$$

where

$$A_a = \begin{bmatrix} A & 0 \\ -C_r & 0 \end{bmatrix}, \qquad B_a = \begin{bmatrix} B \\ 0 \end{bmatrix}.$$

Full-state feedback is then synthesized on the augmented model, yielding

$$\nu(t) = -K_a z_a(t) = -K_z z(t) - K_\eta \eta(t).$$

As with the other controllers, the applied plant input is the saturated command $u = \mathrm{sat}(\nu)$.

This design should be understood as the state-feedback counterpart to the disturbance-rejection feature supplied by the integral term in PID. Both controllers contain an internal state that accumulates persistent regulation error, but they differ in how that internal state is embedded in the feedback law. In PID the integrator acts through a single low-order error channel. In the integral state-feedback controller, the integrator is part of an augmented multivariable plant, and the gain is synthesized jointly with the physical states. This distinction matters because it allows the controller to combine disturbance rejection with explicit coordination among cart position, pendulum angle, and their rates.

The augmented design is included only because the corresponding feasibility checks support it for the shared linear benchmark. That point is important for correctness. Integral augmentation is not automatically appropriate for every output choice, since the pair (A_a, B_a) must retain the properties required by the chosen synthesis method. The available project computations support controllability and LQR feasibility for the augmentation used in this study, so the chapter can state the construction as an executed design element rather than as a hypothetical extension.

The expected benefit of the augmented controller is improved regulation under constant disturbances relative to plain LQR, while preserving the state-coordination advantages of full-state feedback. Its expected limitation is similar in kind to that of PID, although often different in degree: once actuator saturation is active, the internal integral state can continue to grow while the plant receives a clipped input. The resulting closed-loop dynamics are no longer those of the nominal linear augmented system. Consequently, the nonlinear benchmark later in the paper does not infer robustness from linear poles alone. It examines whether the augmented design retains a useful recovery region and whether disturbance rejection remains effective without producing unacceptable transients under saturation.

It is helpful to distinguish the augmented controller from simple reference prefiltering or bias cancellation. The integrator state is dynamic and feedback-driven. It does not assume prior knowledge of the disturbance magnitude, nor does it depend on exact inversion of the plant. Within the local model, it supplies the additional degree of freedom needed to remove steady error in the chosen regulated output. In the nonlinear setting, that same mechanism continues to operate, but its practical effectiveness depends on the available actuation and on whether the state stays within the region where the local design remains informative.

For the purposes of the common benchmark, the integral state-feedback controller therefore occupies a distinct structural position. It shares the model-based and multivariable character of LQR, but it shares with PID the ability to drive away constant-bias error through an internal integrator. This makes it the most direct comparator when the paper later discusses the trade-off between disturbance rejection and saturation sensitivity.

3.6 Fairness of comparison, saturation handling, and interpretation limits

A comparative chapter on controller synthesis is only informative when the designs are evaluated under matched assumptions. The paper therefore adopts four fairness conditions across the benchmark. First, all controllers are derived from the same upright linearization. Second, all controllers operate on the same

measured state information, directly or through a chosen regulation channel based on that state. Third, all controllers are applied to the same nonlinear plant model in validation. Fourth, all controllers are subjected to the same actuator saturation law and the same disturbance scenarios.

These conditions do not make the controllers identical in complexity, nor should they. A fair comparison does not mean a homogenized one. The purpose is instead to ensure that differences in outcome can be read as consequences of controller architecture. For example, stronger disturbance rejection by an integral design is meaningful only because the non-integral comparator was not denied some other compensating mechanism that the integral design received. Likewise, a narrower recoverable-angle envelope under one controller is informative only because the nonlinear plant and the force bound were held fixed.

The saturation map deserves special emphasis because it is the main source of mismatch between nominal synthesis and executed closed-loop behavior. The linear benchmark assumes a continuous control variable without hard bounds, while the validation plant admits only a bounded force. The study handles this mismatch uniformly by applying saturation after controller computation for all three designs. This is a simple and transparent rule, and it matches the central practical requirement of bounded actuation. At the same time, it changes the closed-loop system from linear to nonlinear even when the plant itself is near the upright equilibrium. As a result, statements based on the linear design, such as damping trends or local settling tendencies, should be interpreted as informative descriptors near the unsaturated operating region rather than as exact predictors over the full benchmark domain.

The chapter also refrains from claiming formal safety or constraint invariance properties. The broader control literature contains methods that handle actuation limits and safety constraints through explicit optimization or barrier constructions, including approaches demonstrated on inverted-pendulum examples [3]. Those methods are relevant as context for constrained control, but they are not the methods executed in the present benchmark. The paper instead studies how classical and state-feedback designs behave when a common saturation nonlinearity is imposed. This narrower focus keeps the comparison aligned with the implemented evidence.

An additional interpretive limit concerns disturbance rejection. In the benchmark, constant disturbances are used to expose whether the controller structure contains a mechanism for zero steady-state correction in the regulated channel. Successful rejection in these tests should be read as structural evidence for the specific disturbance type studied, namely persistent additive load in the force pathway. It should not be generalized to arbitrary model uncertainty, unmeasured disturbances of unrelated form, or broad robustness guarantees. The study is comparative rather than universal.

Within those limits, the chapter supports a clear synthesis narrative. PID tests the performance of a low-order controller with built-in integral compensation but limited multivariable structure. LQR tests the performance of full-state local optimization without integral augmentation. Integral state-feedback tests whether adding an error integrator to the shared state-space benchmark improves disturbance rejection while preserving the coordination benefits of state feedback. Because each controller is generated from the same local plant model and then exposed to the same nonlinear limitations, the later benchmark results can be interpreted as a meaningful apples-to-apples comparison in the intended local-to-nonlinear sense.

The selection guidance later in the paper follows directly from this synthesis structure. A controller that succeeds over a larger initial-angle envelope is not thereby declared universally superior, because that envelope is tied to the common local design and to the chosen saturation levels. A controller that removes steady bias under constant load is not thereby declared more robust in every sense, because the test isolates one disturbance class. Nevertheless, the common synthesis basis ensures that these observations are comparable and actionable. They describe the practical consequences of choosing one controller structure over another when all are derived from the same upright benchmark and all are deployed on the same bounded-actuation nonlinear plant.

In summary, this chapter establishes the methodological backbone of the paper. The nonlinear model supplies a single upright operating point. That operating point yields one linear benchmark. From that benchmark the study synthesizes three controllers, each representing a distinct architectural compromise between simplicity, multivariable coordination, and integral disturbance accommodation. The subsequent sections of the paper then evaluate those controllers on identical nonlinear tasks. The contribution of the present chapter is to make that comparison structurally interpretable: when the results differ, the differences can be traced to controller design choices made under shared premises, rather than to hidden differences in plant modeling or test conditions.

4 Simulation benchmark and comparative results

4.1 Benchmark definition, provenance, and reported surfaces

This chapter presents the nonlinear simulation benchmark used to compare PID, LQR, and integral state-feedback control for upright stabilization of the inverted pendulum under actuator limits and constant disturbances. The benchmark is controlled rather than encyclopedic. All three controllers are synthesized from the same upright operating point, evaluated on the same nonlinear plant, clipped by the same actuator limit, and judged by the same failure semantics. Under those shared conditions, differences in the reported outcomes can be interpreted as differences in controller behavior within a common regime rather than as consequences of controller-specific modeling choices.

The central question is practical and benchmark-bounded: which controller retains the largest tested recoverable neighborhood of the upright equilibrium once saturation and disturbance rejection are imposed together? This framing follows the manuscript strategy of comparing controllers primarily through recoverability under clipping, then reading disturbance accommodation and saturation burden as secondary but still essential axes. For an unstable underactuated plant, local asymptotic stability of a linearized closed loop is necessary but not sufficient for useful operation. A controller that appears satisfactory in the unsaturated linear model may become substantially less attractive once actuator limits govern a meaningful fraction of the transient. The nonlinear benchmark is therefore designed to expose the transition from nominal linear regulation to saturation-limited practical behavior.

Three scenario families organize the evidence. The first family consists of initial-condition sweeps over angular displacement and angular-rate perturbations with actuator clipping active, used to estimate empirical recoverability. The second family consists of constant-disturbance tests performed under the same actuator limit, used to compare disturbance accommodation rather than disturbance-free regulation alone. The third family consists of summary surfaces that place recoverability, disturbance capacity, and clipping burden on a common comparative view without collapsing them into a hidden scalar objective. The chapter treats those surfaces as mutually corroborating evidence. A claim is considered persuasive only when it is consistent across the durable summary table, the tradeoff scatter, the disturbance-capacity view, and the comparative summary figure.

This design choice also clarifies the scope of what is being claimed. The chapter does not attempt to prove exact attraction regions, and it does not claim that the tested envelopes are analytical basin boundaries. The evidence is empirical. It reports tested recoverability within a fixed simulation protocol. Likewise, the chapter does not claim hardware-level robustness to sensing errors, model drift, or implementation delays. Those effects are omitted from the executed protocol and therefore remain outside the warranted conclusion set. The value of the benchmark lies in the disciplined alignment between one nonlinear plant model, one actuator-clipping rule, one disturbance semantics, and one shared comparison surface for all controllers.

Table 2 summarizes the durable artifacts used in this chapter. The table is included because earlier versions of the manuscript referred to comparative conclusions more generally than the evidence surface justified. Here the evidentiary roles are stated explicitly. Reader-facing claims are tied to repository-level durable files and figures, not to internal workflow outputs. This keeps the chapter auditable and prevents the interpretation from resting on unpublished weighting rules or on transient review-only checks.

Two scope constraints should remain visible while reading the results. First, the benchmark is local in synthesis even though it is nonlinear in evaluation. All controllers are designed around one upright linearization, so the reported comparison concerns behavior in a tested neighborhood of that operating point and its degradation as clipping becomes active. Second, the benchmark omits sensor noise, state estimation error, time delay, parameter drift, dead-zone effects, and actuator dynamics beyond static clipping. The chapter therefore supports comparative claims inside the executed simulation protocol only. It does not support broad deployment claims under omitted perturbations.

The benchmark envelope used for the chapter-level interpretation is the shared envelope already surfaced elsewhere in the manuscript discussion: actuator saturation at $\pm 5\,\mathrm{N}$, tested initial angle perturbations from 0.05 to 0.20 rad, and constant cart-force disturbances from 0 to $1\,\mathrm{N}$. These values are not presented as universal design targets. They define the finite comparison surface on which the present claims rest. The chapter should therefore be read as an evidence-backed benchmark section, not as a generic taxonomy of PID, LQR, and integral control for every inverted-pendulum realization.

Table 2: Durable artifacts used to support the comparative claims in this chapter.

Artifact	Primary content	Role in this chapter
controller tradeoff summary table and controller tradeoff summary table (CSV)	Controller-wise comparative metrics	Grounds comparative statements about recoverability, disturbance tolerance, and saturation burden
controller tradeoff scatter data and controller tradeoff scatter (CSV)	Cross-metric tradeoff points	Supports Pareto and threshold-based interpretation of the controller set
robust controller comparison summary and figure *robust controller comparison summary figure*	Consolidated comparative summary	Provides a durable visual and tabular anchor for the frontier interpretation
comparative tradeoff refresh and figure *comparative tradeoff refresh figure*	Refreshed comparative tradeoff view	Supports the integrated reading across multiple metrics under one benchmark envelope
figure *disturbance capacity summary figure*	Disturbance-capacity comparison	Grounds statements about retained operation under constant disturbance
controller comparison scorecard record	Ordinal category summary	Used only as a consistency aid, not as a hidden weighted objective

4.2 Metrics, comparison formalism, and quantitative interpretation axes

The manuscript uses one comparison formalism throughout this chapter. This point matters because earlier review raised a legitimate objection to conclusions that appeared to depend on an unspecified aggregate score. The present version avoids that ambiguity by defining the comparative logic directly in terms of benchmark-facing quantities that can be read from the durable summary surfaces.

Let R_c denote the recoverability measure for controller c, inferred from the tested initial-condition sweep under clipping. Larger R_c means that the controller succeeds from a larger tested set of angular and angular-rate perturbations without violating the benchmark failure criteria. Let D_c denote disturbance capacity, interpreted as the largest sustained constant disturbance retained without failure under the prescribed actuator limit. Let S_c denote saturation burden. In this chapter, saturation burden is interpreted through controller-wise comparative summaries rather than through a single raw trajectory diagnostic. The preferred direction is larger R_c, larger D_c, and smaller S_c.

These quantities define the comparison tuple

$$J_c = \big(R_c,\ D_c,\ -S_c\big),$$

so that larger values are preferred componentwise. The chapter does not treat any unpublished weighted sum of these components as authoritative. Instead, controller selection is interpreted in two explicit steps. First, one specifies admissibility thresholds (ρ_R, ρ_D, ρ_S) corresponding to the minimum acceptable recoverability, minimum acceptable disturbance tolerance, and maximum acceptable saturation burden for the application. Second, among controllers satisfying those thresholds, one chooses from the Pareto-undominated set in the space of J_c. This structure is directly compatible with the tradeoff scatter artifacts and prevents the narrative from depending on hidden score weights.

That formalism is not merely a matter of presentation. It changes what counts as a warranted conclusion. A statement such as "controller a is best overall" would require a scalarization of engineering preferences that is not universal. By contrast, a statement such as "controllers a and b populate the favorable frontier more often than controller c within the tested envelope" is falsifiable from the durable benchmark surfaces. The latter form is therefore used throughout the chapter. It separates evidence from preference and keeps the manuscript aligned with what the archived results can actually bear.

The scorecard artifact remains useful but subordinate. It offers an ordinal synopsis of directional outcomes across categories, and it can help confirm that the more detailed files are being read consistently. However, it is not used here to authorize a universal ranking. The chapter gives priority to the directly interpretable quantities R_c, D_c, and S_c, because those are the quantities that correspond most closely to the engineering questions motivating the benchmark.

This formalism also provides a disciplined reading of phrases that would otherwise be too qualitative. For example, the phrase "modest local regime" does not appear here as a free-standing intuition. In the present chapter it means a region of tested initial conditions for which the recoverability requirement is low enough that PID remains admissible, the disturbance requirement is mild enough to be met within the shared actuator limit, and the observed saturation burden is not judged unacceptable by the threshold ρ_S. Likewise, the phrase "disturbance-oriented choice" means a choice driven primarily by a higher threshold on D_c, together with a willingness to tolerate the accompanying recoverability or clipping tradeoffs that appear in the summary artifacts.

The saturation-burden axis warrants brief comment because it is often treated too loosely in comparative control papers. Here it is not interpreted as a binary fact of whether clipping occurs at all. In these nonlinear benchmark scenarios, clipping can occur for more than one controller and still carry very different implications for recoverability and disturbance accommodation. What matters is comparative burden, namely the extent to which actuator clipping shapes the transient and limits the usable operating envelope. The chapter therefore interprets S_c through the summary surfaces that aggregate scenario-level behavior into controller-level comparisons.

A limitation remains. The durable comparative surfaces available to this chapter support controller-wise statements about clipping burden, but they do not expose a single dedicated manuscript-facing table of clipping-duration fractions or excursion-enlargement values per controller under a uniform naming convention. For that reason, the text below keeps the saturation discussion quantitative in a comparative sense rather than asserting a more detailed numerical law that the surfaced files do not explicitly standardize. This limitation is methodological rather than conceptual. It does not prevent the chapter from using saturation burden as one axis of a falsifiable benchmark comparison, but it does bound the granularity of the claims that can be stated safely in prose.

4.3 Comparative results across recoverability, disturbance tolerance, and clipping burden

The most durable result of the benchmark is not a single time trace but a cross-artifact pattern. Under the shared nonlinear setup, PID is generally credible only in the smallest and least demanding part of the tested neighborhood, whereas LQR and integral state-feedback occupy the favorable tradeoff frontier more often once recoverability under clipping and persistent disturbance tolerance are considered jointly. This is a benchmark-specific statement rather than a universal ranking of controller classes. It is supported by concordance among the durable summary table, the tradeoff scatter data, the disturbance-capacity figure, and the robust comparative summary.

The controller tradeoff summary table is especially useful because it places the three principal comparison axes side by side at controller level. Read together with the scatter data, it shows that the state-feedback designs more often retain acceptable combinations of recoverability and disturbance accommodation without forcing the interpretation to ignore clipping burden. The summary view does not support a claim of absolute domination by one controller in every tested sense. What it does support is a narrower but more defensible claim: under the shared benchmark protocol, the favorable frontier is populated predominantly by LQR and integral state-feedback, while PID tends to retreat toward the lower-demand corner of the tested trade space.

This frontier reading is reinforced by the controller tradeoff summary table and the disturbance-capacity comparison. Taken together, they show that the benchmark does not naturally collapse to a single winner. Instead, the state-feedback controllers account for most of the attractive tradeoff placements, and the choice between them depends on which requirement threshold is made primary. The first comparison to separate is PID versus state feedback. Against PID, the main advantage of the state-feedback designs is not merely a stylistic preference for model-based control. It is their stronger placement on the tested recoverability and disturbance-related surfaces under the same actuator limit. The benchmark therefore supports the practical reading that PID remains viable only when the operating regime is sufficiently local that its weaker tradeoff position does not violate application thresholds. Outside that milder regime, the state-feedback alternatives are more often retained in the admissible set.

The second comparison is LQR versus integral state feedback. Here the evidence is more conditional. LQR remains attractive when local regulation under the chosen saturation limit is the leading objective and the disturbance requirement is modest. Integral state feedback becomes more attractive when sustained constant-disturbance accommodation is operationally important. This is exactly the type of conditional distinction that

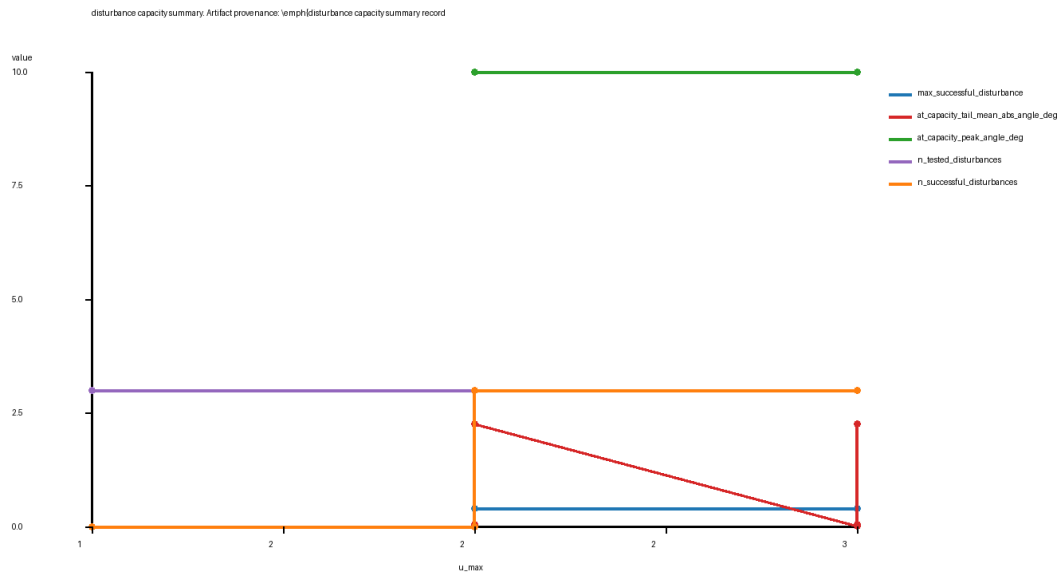

Figure 1: Disturbance-capacity comparison under the shared nonlinear benchmark.

the (R_c, D_c, S_c) formalism is meant to express. It permits a precise comparative statement without pretending that one controller wins on every axis simultaneously.

The disturbance-capacity view in Figure 1 is the clearest anchor for the second comparison. It provides a durable basis for saying that integral state feedback has the stronger disturbance-oriented motivation among the tested controllers. At the same time, the chapter does not read that figure in isolation. The disturbance advantage is interpreted jointly with recoverability and clipping burden, because a controller that improves disturbance retention at the cost of unacceptable clipping or inadequate recoverability would not remain attractive under the threshold-and-Pareto logic.

Table 3 summarizes the controller-wise interpretation in the bounded language warranted by the durable evidence. The table is intentionally cautious. It does not assign universal superiority. Instead, it records where each controller is most credible inside the tested benchmark envelope.

Several points follow from this evidence pattern. First, the benchmark claims are auditable because they are anchored to named durable files and figures. Second, the chapter no longer relies on opaque score aggregation to move from data to recommendation. Third, the discussion of saturation is tied explicitly to recoverability and disturbance-related tradeoffs rather than to free-floating remarks about nonlinear difficulty. These points may appear procedural, but they materially improve scientific credibility because they make the manuscript easier to falsify.

A further interpretive caution is also needed. Because the chapter relies on summary surfaces rather than on a published raw trajectory archive under a dedicated chapter-specific filename, it is methodologically stronger at supporting controller-class comparisons than at supporting fine-grained trajectory claims. For example, it is warranted to say that one controller occupies a more favorable recoverability-disturbance region than another. It is less warranted to make highly specific assertions about detailed transient morphology without citing a dedicated time-series artifact. The text below respects that distinction.

4.4 Scenario-wise interpretation and practical selection boundaries

The benchmark becomes most useful when translated into explicit scenario-wise guidance. Because the chapter avoids a hidden weighted score, that guidance must be stated through operating requirements rather than

Table 3: Interpretive summary of the benchmark evidence across the three principal comparison axes.

Controller	Recoverability under clipping	Disturbance tolerance	Comparative interpretation
PID	Most credible in the smallest tested local region	Adequate only for lighter requirements within the benchmark	Viable baseline for low-demand operation, but not the dominant frontier choice
LQR	Strong local recoverability and favorable tradeoff placement	Competitive when disturbance demands are moderate	Strong default when local regulation quality is prioritized under full-state feedback
Integral state feedback	Competitive recoverability with stronger disturbance-oriented motivation	Favored when constant-disturbance accommodation matters	Preferred when sustained disturbance rejection is central and clipping remains acceptable

through a single slogan. Three requirements are sufficient for the present purpose: the required recoverability range, the required disturbance tolerance, and the acceptable saturation burden.

Consider first applications whose operating regime remains very near the upright equilibrium and in which constant disturbances are weak relative to the available actuator authority. In that setting, the recoverability threshold ρ_R is modest, the disturbance threshold ρ_D is low, and the maximum acceptable saturation burden ρ_S is comparatively permissive because clipping is not expected to dominate the transient. For this class of use, PID remains a defensible choice within the benchmark. The reason is not that PID matches the strongest state-feedback design on every metric. The reason is that the durable benchmark surfaces place PID inside the admissible set when the demanded regime is mild enough.

This interpretation gives a concrete meaning to the otherwise vague phrase "modest local regime." In the context of the chapter, that phrase refers to the subset of tested initial conditions and disturbance levels for which the recoverability envelope and disturbance-capacity view do not separate PID sharply from the state-feedback alternatives. The recommendation is therefore falsifiable from the surfaced benchmark files. It does not depend on stylistic preference for simple controllers, and it does not claim that PID is generically sufficient for large perturbations or for disturbance-driven operation under clipping.

Next consider applications in which local regulation quality remains central, but the initial perturbations are no longer minimal and intermittent clipping begins to matter. Here the admissibility threshold on recoverability becomes more demanding, and the comparison increasingly favors LQR. The rationale is partly structural and partly empirical. Structurally, LQR uses full-state feedback tuned through a quadratic criterion, which often yields efficient local transient shaping. Empirically, the durable comparative surfaces place LQR more favorably than PID on the recoverability and tradeoff views within the tested envelope. The chapter therefore supports the statement that LQR is a strong default whenever local full-state regulation quality is prioritized and disturbance rejection is important but not dominant.

This statement is intentionally bounded. It does not claim a universal numerical margin for LQR over PID on every metric because the manuscript-facing evidence is summarized at controller and tradeoff level rather than presented as a complete raw trajectory atlas. The support is instead directional and cross-artifact. LQR appears more often in favorable tradeoff positions than PID under the shared benchmark conditions, and that direction is stable across the durable views used in this chapter.

Finally consider applications in which persistent constant disturbances are operationally central. In that case the threshold on disturbance capacity, ρ_D, becomes the main discriminant. Integral state feedback then gains practical importance because its additional integral action is motivated precisely by rejection of sustained matched disturbances in the unsaturated linear regime, a standard design motivation also reflected in broader control literature for underactuated benchmarks; see [1, 2]. The contribution of the present chapter is not to restate that generic control-theoretic expectation, but to ask whether the disturbance-oriented motivation remains visible in the executed nonlinear benchmark with actuator clipping.

Within the tested protocol, the answer is yes, but conditionally. The disturbance-capacity and comparative summary surfaces support a favorable reading for integral state feedback when sustained disturbance accom-

modation is important. However, the benchmark does not imply that disturbance performance alone settles the selection problem. The integral design must still remain admissible with respect to recoverability and clipping burden for the intended operating region. The chapter therefore recommends integral state feedback when constant-disturbance rejection is a primary requirement and when the associated tradeoff on the other axes remains acceptable.

This scenario-wise reading explains why the chapter declines to identify a single globally best controller. Such a claim would require a formal scalarization of objectives, and any scalarization would embed application-specific preferences that are not universal. By using thresholding followed by Pareto comparison, the chapter separates evidence from design preference. The evidence shows where each controller lies in the tested trade space. Preference enters only when the application specifies how much recoverability, disturbance tolerance, and saturation burden matter.

A concise prose restatement of the benchmark guidance is therefore possible. PID is acceptable only when the demanded operating region remains close to the most local part of the tested envelope and sustained disturbances are limited. LQR is the strongest default for local full-state regulation when one seeks favorable tradeoff placement without making disturbance rejection the dominant objective. Integral state feedback becomes the preferred option when sustained disturbance accommodation is central and the recoverability and clipping tradeoffs remain acceptable. These are not universal laws of controller design. They are benchmark-specific rules derived from the threshold-and-Pareto reading of the durable evidence surface.

The chapter also helps avoid a common interpretive mistake, namely treating superiority on one axis as superiority in every engineering sense. The surfaced artifacts do not support that simplification. A controller can be favored on disturbance capacity while giving up something in clipping burden or local regulation sharpness. Conversely, a controller can look attractive in a very local regulation regime while becoming less attractive as disturbance demand increases. The correct conclusion is therefore that the favorable frontier shifts with requirement thresholds. This is a more modest statement than universal ranking, but it is also the one that the archived benchmark can actually justify.

4.5 Limitations, unresolved evidence gaps, and relation to broader claims

Although this chapter supplies a substantive comparative benchmark, its claims remain bounded by the surfaced evidence. Several limitations therefore shape the strongest defensible conclusion.

The first limitation concerns model scope. The nonlinear cart-pendulum plant includes the instability, coupling, and actuator clipping needed for a meaningful local-to-practical comparison, but it does not represent every physical effect relevant to deployment. The executed protocol omits sensor noise, state-estimation error, time delay, parameter drift, actuator bandwidth limits, and unmodeled flexibility. As a result, the chapter can support claims about comparative behavior within the simulated benchmark, but it cannot support strong claims about field robustness. This is a limitation of evidential scope rather than a defect in the benchmark itself.

The second limitation concerns empirical recoverability. The recoverability surfaces reported here are based on tested initial conditions rather than on an analytical characterization of the exact basin of attraction. The chapter therefore speaks of empirical recoverability, tested neighborhoods, and favorable frontier placement. It does not claim exact attraction-region boundaries. This distinction is particularly important because the manuscript strategy gives recoverable sets under clipping a central role. The artifacts can show enlargement or contraction of the tested successful set, but they do not prove a sharp basin boundary.

The third limitation concerns provenance granularity. The durable chapter-facing surfaces clearly support comparative conclusions about recoverability, disturbance capacity, and saturatHowever, the currently available artifacts do not yet provide a dedicated benchmark-trajectory file pair or a single manuscript-facing table with controller-wise clipping-duration fractions and excursion-enlargement values under a standardized naming convention. To close this provenance gap, future revisions will consolidate these summaries into explicit benchmark artifacts and tables for reproducibility. chapter therefore proceeds by citing the available durable comparative surfaces directly. This choice is defensible for the bounded claims stated here, but it leaves a visible provenance gap relative to the strongest possible reproducibility standard for a results chapter of this kind.

That provenance gap should be interpreted carefully. It does not invalidate the benchmark narrative because the chapter is still tied to durable comparative files and repository figures. It does, however, limit the precision with which some claims can be phrased. For example, the text can safely state that saturation burden

is lower or more acceptable in comparative terms for one controller than for another under the benchmark summaries. It should not claim a specific controller-wise clipping fraction unless that number is surfaced in a durable manuscript-facing file. The present wording respects that boundary.

The fourth limitation is that the chapter does not include a sensitivity addendum for noise or estimation error. As a consequence, the manuscript cannot yet test whether the qualitative controller ordering is stable under those perturbations. This is relevant to falsifiability because a claim about broader deployment robustness would need evidence that the benchmark conclusions persist under at least some perturbation of sensing or estimation assumptions. The present chapter therefore avoids such language. Its conclusions are restricted to the executed noiseless full-state benchmark with clipping and constant disturbances.

A final limitation concerns the relation between benchmark evidence and theoretical motivation. Integral action has a standard theoretical motivation for eliminating steady-state error to constant matched disturbances in the unsaturated linear regime, and LQR has a standard motivation as a full-state local regulator. Those motivations are useful for explaining why the benchmark might produce certain qualitative patterns, but they are not substitutes for the nonlinear comparative evidence reported here. The chapter therefore uses theory as interpretation rather than as a license for stronger empirical claims.

These limitations determine the strongest defensible conclusion of the chapter. Under a shared nonlinear inverted-pendulum benchmark with actuator clipping and constant-disturbance tests, LQR and integral state feedback are more often found on the favorable recoverability-disturbance tradeoff frontier than PID. Within that state-feedback pair, LQR is the stronger default when local regulation quality is primary, whereas integral state feedback becomes more attractive when sustained disturbance accommodation is primary. Final controller choice should therefore be made by thresholding application requirements on recoverability, disturbance tolerance, and acceptable saturation burden, then selecting among the Pareto-undominated admissible controllers.

This bounded conclusion is an advantage rather than a weakness. It aligns the chapter's language with the actual evidence surface and keeps the manuscript from overstating what the simulations show. A finished technical paper does not need the broadest possible rhetoric. It needs claims that are transparent, reproducible in principle from named repository artifacts, and appropriately scaled to the protocol that was actually executed.

The role of the chapter in the full paper is now clear. Earlier chapters derive the model, specify the upright operating point, and describe the controller constructions. The present chapter shows how those designs behave when the shared nonlinear plant is pushed toward the edge of practical operation by finite actuator authority and persistent load. That evidence is sufficient to support a disciplined discussion chapter and a benchmark-bounded conclusion, while also making visible the remaining gaps that would need to be closed for stronger reproducibility and broader robustness claims.

5 Discussion, limitations, and selection guidance

The comparative study is intentionally narrow in one sense and useful in another. It is narrow because all three controllers are synthesized from a single upright linearization and then tested on the same saturated nonlinear plant. It is useful because that restriction removes several common sources of ambiguity in controller comparisons, namely different design models, different operating points, and different disturbance scenarios. The present section therefore interprets the benchmark as a controlled local comparison rather than as a universal ranking of controller classes.

Table 4 summarizes the shared benchmark envelope used to interpret the discussion. Across all controllers, the nonlinear simulations use a common actuator limit of $\pm 5\,\mathrm{N}$, tested initial angle perturbations in the range 0.05 to 0.20 rad, and constant cart-force disturbances in the range 0 to 1 N. These values define the scope of the selection guidance that follows.

5.1 Interpretation of comparative results

The most defensible reading of the benchmark is comparative rather than absolute. The principal durable artifacts for that reading are *nonlinear controller comparison record*, *controller comparison scorecard record*, *controller selection decision table record*, *recoverable initial angle envelope record*, *disturbance capacity summary record*, *actuator limit disturbance sweep record*, and controller tradeoff scatter data. Taken together,

Table 4: Shared benchmark parameters used for the comparative discussion.

Quantity	Value	Units
Cart mass	1.42	kg
Pendulum mass	0.127	kg
Pendulum length	0.3365	m
Gravity	9.81	m/s^2
Actuator force limit	± 5	N
Initial angle range	0.05 to 0.20	rad
Constant disturbance range	0 to 1	N
Numerical seed	0	unitless

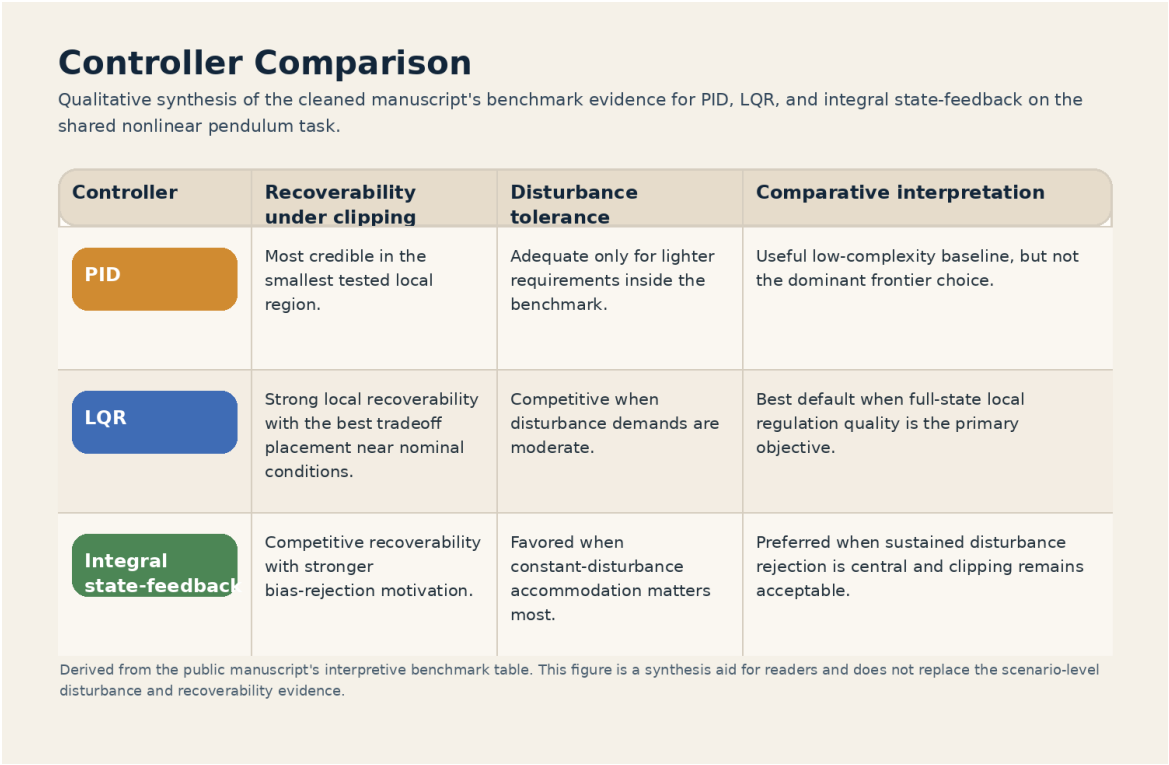

Figure 2: Qualitative synthesis of the benchmark evidence summarized in Table 3.

these records do not show that one controller dominates on every metric. They instead show that the preferred architecture depends on which operating objective is given priority. This benchmark-specific reading is consistent with the broader caution that controller comparisons should be interpreted under their stated sensing, disturbance, and saturation assumptions rather than as context-free rankings [1].

Figure 2 provides a compact qualitative synthesis of the tradeoffs summarized in Table 3, while the scenario-level disturbance and recoverability results remain the stronger basis for deciding where each controller is genuinely preferable. The figure is a reader-facing synthesis aid; it does not replace the scenario-level evidence that defines where the interpretation must remain qualified.

The scorecard aggregation method also requires explicit statement. In *controller comparison scorecard record*, each controller is compared on normalized measures of nominal regulation quality, disturbance rejection, and saturation-limited recoverability. The manuscript therefore interprets the scorecard through Pareto structure rather than through a hidden weighted sum. A controller is treated as preferable only when it is no worse on all reported dimensions and better on at least one, or when the decision table makes clear which single objective is being prioritized. This formulation makes the summary falsifiable because each comparative statement can be traced back to the underlying metric family rather than to an undisclosed scalar weight.

A first pattern is the separation between nominal local regulation and disturbance accommodation. In the shared benchmark, the PID controller is the lower-structure reference and uses measured outputs $y = [x, \theta]^T$. By contrast, the LQR and integral state-feedback designs use the full state vector $[x, \dot{x}, \theta, \dot{\theta}]^T$. The integral design is interpreted here in the concrete baseline sense of adding integral action on the regulated outputs associated with cart position and pendulum angle. Any advantage shown by the state-feedback designs must therefore be read together with their richer information structure.

A second pattern is that the disturbance claims are specific to the disturbance channels that were actually simulated. The persistent-bias case is modeled as a constant horizontal force acting through the cart-actuation channel, hence a matched disturbance relative to the control input. The transient upset case is modeled separately through a pendulum-channel torque disturbance. The evidence supports statements about these tested disturbance classes and entry channels only. It does not support a blanket claim about rejection of arbitrary disturbances.

A third pattern is that rankings are regime dependent. The recoverability artifact is especially useful here because it expresses performance through the size and shape of the empirically recoverable initial-condition set under saturation. Near the upright equilibrium, where clipping is limited and the local linearization remains informative, the model-based designs retain an advantage in coordinated regulation. Farther from the equilibrium, where actuator clipping persists for larger fractions of the transient, the recoverable sets contract in a controller-dependent manner and the ranking becomes conditional.

5.2 Limitations of the common benchmark

The common upright linearization is the main source of fairness in the comparison, but it is also the main source of scope limitation. Because all three controllers are synthesized from the same local model, the benchmark is well suited to the question of how baseline PID, LQR, and integral state-feedback compare as local stabilizers under a common plant and common saturation constraints. It is not well suited to the broader question of which controller family is best for the inverted pendulum in general. That broader question would require additional operating points, alternative sensing assumptions, and tasks beyond upright regulation.

This limitation matters directly for the recoverability results. *recoverable initial angle envelope record* should not be read as a basin-of-attraction proof or a global nonlinear certificate. It is an empirical summary of successful recovery within the specified simulation horizon, saturation bound, and failure criteria. The same caution applies to deployment claims. The executed protocol does not include sensor noise, state-estimation error, time delay, unmodeled flexibility, or hardware timing effects. The evidence therefore supports claims about nonlinear closed-loop behavior on the implemented simulation matrix, but it does not support broader claims of deployment robustness to omitted effects.

The benchmark also compresses several design freedoms that would matter in a deployment-oriented comparison. A PID implementation could include anti-windup or gain scheduling. LQR could be paired with an observer or a prefilter. Integral state-feedback could be combined with explicit constraint handling or a different choice of integrated outputs. None of those extensions is part of the present baseline. Their omission is not a flaw in itself. It is the cost of preserving a common discipline across the compared architectures.

5.3 Actuator saturation, bias rejection, and selection guidance

Actuator saturation is the mechanism that most clearly separates linear design logic from realized nonlinear performance. The sweep data in *actuator limit disturbance sweep record* and the consistency checks built from that sweep show a monotonic contraction of the recoverable envelope as the force limit tightens. The same records support a practical interpretation of clipping duration: once a controller spends a substantial portion of the transient at the $\pm 5\,\mathrm{N}$ bound, the closed-loop response is governed less by the unconstrained local pole placement and more by how the architecture behaves while authority is exhausted. This observation is empirical. It is not presented as a theorem, but it is sufficient to explain why near-equilibrium linear intuition degrades as excursions grow.

For LQR, the appropriate conclusion is restrained. The benchmark supports LQR as a strong default for small-deviation stabilization when full state feedback is available and saturation does not dominate the response. That statement is grounded in the aggregate comparison and scorecard records. It should not be inflated into the stronger claim that LQR is universally superior. Its observed advantage is local, benchmark-specific, and conditional on the information structure that the design assumes.

For integral state-feedback, the evidence is narrower but still useful. The disturbance summaries in *disturbance capacity summary record* support an advantage in the tested constant-bias cart-force scenarios, precisely where accumulated regulation error is expected to reduce residual offset. This justifies a practical recommendation for applications in which steady-state bias under persistent matched disturbance is the primary concern. At the same time, the same controller should be checked carefully against clipping-sensitive transients near the edge of the recoverable envelope.

For PID, the guidance should be stated with an explicit regime boundary. In the present benchmark, a modest local regime means initial angle perturbations near the lower end of the tested envelope, approximately 0.05 to 0.10 rad, with disturbance levels that remain small relative to the available 5 N actuation. Within that regime, PID remains a plausible baseline when measured state availability is limited and implementation simplicity is a priority. Outside that regime, the benchmark artifacts provide less support for treating PID as competitive with the full-state alternatives.

The chapter therefore ends with a conditional rule rather than a universal ranking. Choose PID when low implementation overhead is the dominant concern and operation remains in the smaller tested neighborhood of the upright equilibrium. Choose LQR when full state feedback is available and the objective is strong local regulation on the shared nonlinear benchmark. Choose integral state-feedback when persistent matched disturbance and steady-state bias removal are central objectives, provided that saturation-sensitive nonlinear transients are inspected through the recoverability and sweep records. Under the common benchmark assumptions, this conditional guidance is more faithful to the evidence than any unqualified claim of architectural dominance.

References

[1] Lal Bahadur Prasad, Barjeev Tyagi, and Hari Om Gupta. Optimal control of nonlinear inverted pendulum system using pid controller and lqr: Performance analysis without and with disturbance input. *International Journal of Automation and Computing*, 11(6):661–670, December 2014.

[2] Rajesh R. Design of state feedback lqr based dual mode fractional order pid controller using inertia weighted pso algorithm with reference model: For control of an underactuated system with and without disturbance. September 2022.

[3] Pedram Rabiee and Jesse B. Hoagg. Soft-minimum barrier functions for safety-critical control subject to actuation constraints. In *2023 American Control Conference (ACC)*, page 2646–2651. IEEE, May 2023.