

The Pupillary Light Reflex as a Dynamical Arousal Index: A Delay-Differential Framework for Non-Invasive Locus Coeruleus Estimation

Sour Ismail

github.com/Hidokun

`ismail.sour@etu.unilim.fr`

May 4, 2026

Abstract

Standard pupillometry measures resting diameter and discards everything else. The recovery trajectory following a light pulse, however, encodes the loop gain of the underlying delayed negative-feedback system, and that gain is set by tonic Locus Coeruleus output through the sympathetic dilator pathway. The recovery time constant τ_{return} is therefore a non-invasive dynamical index of LC arousal state, inaccessible to instruments that treat the pupil as a static aperture.

This paper develops the computational framework that makes that index rigorous. The pupillary light reflex is modelled as a nonlinear delay-differential equation following Longtin and Milton [5], with a gain structure that pins the linearised loop gain exactly to a single dimensionless parameter at every operating point. The exact analytical solution of the characteristic equation replaces the classical zero-delay approximation, relocating the bifurcation threshold and providing a calibration curve whose sensitivity is concentrated precisely where arousal resolution matters most. Simulation confirms that the recovery time constant tracks this curve closely under realistic noise conditions, and a spectral analysis of continuous pupil traces cross-validates the bifurcation geometry through an independent computational pathway. A three-variable non-invasive instrument architecture is specified, combining resting diameter, recovery dynamics, and chromatic post-illumination response into a self-consistent output whose internal validity is guaranteed by the physiology rather than imposed by design. All results are computational; three falsifiable predictions are identified for empirical validation.

1 Introduction

The locus coeruleus (LC) is a bilateral pontine nucleus containing the largest concentration of norepinephrine-synthesizing neurons in the human brain. Its tonic discharge rate sets the ambient neuromodulatory gain across cortical and subcortical targets, modulating the signal-to-noise ratio of neural processing in a manner formalized by the adaptive gain theory [1]. In that framework, intermediate tonic LC activity corresponds to an optimal gain regime supporting focused, exploitation-oriented behavior, while excessively low or high tonic rates shift the system toward disengagement or diffuse, exploration-oriented processing respectively. Dysregulation of tonic LC output is implicated in the pathophysiology of anxiety disorders, major depressive disorder, and attention-deficit hyperactivity disorder, and the noradrenergic system is the primary pharmacological target of atomoxetine, reboxetine, and several first-line antidepressants.

Despite this clinical centrality, no non-invasive instrument currently provides continuous real-time estimation of tonic LC state in awake humans. Neuroimaging approaches achieve spatial specificity at the cost of temporal resolution and ecological validity; peripheral biomarkers such as salivary or urinary catecholamine levels integrate over hours and cannot resolve moment-to-moment fluctuations. The gap between the clinical importance of tonic LC monitoring and the available measurement technology is the motivating problem of this project.

The pupillary system offers a partial but tractable proxy. The afferent limb of the pupillary light reflex (PLR) shares its pretectal relay with projections that ascend to noradrenergic and cholinergic nuclei, and the efferent limb receives direct sympathetic input from LC-projecting preganglionic neurons whose tone is set by tonic LC discharge. Four converging lines of evidence establish the empirical basis for pupil-based LC state estimation. Single-unit recordings in non-human primates [4] demonstrate that LC firing rate reliably predicts pupil dilation on timescales from hundreds of milliseconds to several seconds, with the strongest coupling during the tonic baseline period between task events. In humans, [7] spontaneous fluctuations in resting pupil diameter co-vary with BOLD signal in the LC region identified by high-resolution neuroimaging, confirming that the primate relationship generalizes to the human system. Baseline pupil diameter tracks the adaptive gain curve predicted by Gary Aston-Jones and Jonathan D. Cohen [1], with intermediate diameter corresponding to optimal task performance and deviations in either direction predicting performance decrement [2]. At the pharmacological level, atomoxetine, a selective norepinephrine reuptake inhibitor, modulates the phasic pupil response to task-relevant stimuli in a manner consistent with enhanced LC-NE signaling [8]. Taken together, these results establish that pupil size is a valid, if imperfect, index of LC-linked arousal state. The important qualification is that pupil diameter reflects distributed arousal network activity and not LC spiking rate in isolation [6]; the coupling is graded, state-dependent, and subject to non-LC confounds including cholinergic tone, cognitive load, and emotional valence. The instrument proposed here is designed with this limitation explicitly encoded in its output framing.

What current pupillometry does not exploit is the dynamical structure of the PLR recovery trajectory. Standard clinical and research pupillometers reduce each light reflex to two static scalar measurements: resting diameter D_0 and peak constriction amplitude. Both quantities are extracted from the first few hundred milliseconds of the response and the subsequent re-dilation trajectory is discarded. This is a significant information loss, because the re-dilation trajectory is not a passive mechanical relaxation but the transient response of a nonlinear delayed negative-feedback control system returning to its fixed point. the PLR is governed by a delay-differential equation whose stability properties depend on a dimensionless loop gain G that is itself modulated by sympathetic dilator tone, and therefore by tonic LC-NE output [5]. The main hypothesis of this project is that the recovery time constant τ_{return} , extracted from the re-dilation envelope, constitutes a continuous dynamical proxy for tonic LC state that is theoretically grounded, computationally validable and instrumentally realizable. The project delivers a theoretical derivation of the measurement principle via linearization of the Longtin-Milton model, bifurcation analysis, and the exact Lambert W eigenvalue solution (Section 2); a computational validation of the $\tau_{\text{return}}(G)$ relationship across the full physiological gain range using the full nonlinear model with gain-pinned γ scaling (Sections 3 and 4); and a complete optical instrument specification for non-invasive simultaneous measurement of D_0 , τ_{return} , and the chromatic post-illumination pupil response ΔR as a melanopsin-linked functional index (Section 5).

2 Theory: The PLR as a Dynamical System

2.1 Physiology and Feedback Topology

Viewing the pupillary light reflex as a closed negative-feedback loop whose anatomical structure directly determines the gain and delay parameters of the dynamical system. A light pulse incident on the retina propagates through a neural relay chain, from the retina through the pretectal nucleus and the Edinger-Westphal nucleus to the ciliary ganglion and finally the iris sphincter muscle, which functions collectively as the feedback controller. The sphincter contracts in response, reducing the pupil aperture, attenuating the retinal flux, and closing the loop. The total conduction delay accumulated across the afferent and efferent limbs is $\delta = 0.300$ s, and the iris mechanical time constant is $\tau_{\text{iris}} = 0.311$ s [5].

A second pathway is superimposed on this feedback loop through the sympathetic nervous system, which permits tonic LC state to modulate the system. Tonic LC discharge drives a sympathetic chain that terminates at the iris dilator muscle, where norepinephrine release onto adrenoceptors generates a sustained dilator force opposing the sphincter contraction. The net resting pupil area A^* results from the balance between these two competing drives, and because the LC sets the dilator tone, it directly sets the operating point from which the feedback loop runs. Higher tonic LC output shifts the resting pupil area A^* upward, moves the operating point into a lower-sensitivity region of the sphincter response function, and reduces the local sensitivity of the feedback loop to luminance perturbations. This is how tonic LC state sets loop gain G , and it is the causal chain the instrument measures.

2.2 The Longtin-Milton Delay-Differential Equation

The full nonlinear governing equation is the delay-differential equation of [5]:

$$\tau_{\text{iris}} \frac{dA}{dt} = -A(t) + A^* + \gamma [f(\Phi(t - \delta)) - f(A^*)] \quad (1)$$

where $\Phi(t) = A(t)(1 + s(t))$ is the effective retinal light flux, $s(t)$ is the normalized stimulus intensity, and f is the Hill-type sphincter response function:

$$f(\Phi) = \frac{c \theta^n}{\theta^n + \Phi^n} \quad (2)$$

The resting fixed point A^* is an equilibrium for any value of γ , since $\gamma [f(A^*) - f(A^*)]$ vanishes identically. The scaling factor γ is determined by the requirement that the linearized loop gain at A^* equals exactly G .

2.3 Linearization and the Loop Gain G

Writing $A(t) = A^* + a(t)$ with $|a| \ll A^*$ and expanding f to first order around A^* :

$$\tau_{\text{iris}} \dot{a}(t) = -a(t) + \gamma f'(A^*) a(t - \delta) \quad (3)$$

The Hill function is monotonically decreasing, so $f'(A^*) < 0$. Its analytic value is:

$$f'(A^*) = -\frac{cn \theta^n (A^*)^{n-1}}{(\theta^n + (A^*)^n)^2} \quad (4)$$

Defining the dimensionless loop gain $G \equiv \gamma |f'(A^*)|$, the linearized equation becomes:

$$\tau_{\text{iris}} \dot{a}(t) = -a(t) - G a(t - \delta) \quad (5)$$

with $G > 0$, and γ is fixed by:

$$\gamma = \frac{G}{|f'(A^*)|} \quad (6)$$

This construction pins the local dynamics at A^* to exactly G regardless of which Hill parameters (c, θ, n) are chosen. For large perturbations and post-bifurcation dynamics the full nonlinear Hill function remains active, allowing the model to saturate into a stable limit cycle above G_c rather than diverging. The physiological reading follows directly from Section 2.1: as tonic LC output increases, A^* increases, $|f'(A^*)|$ decreases, and G decreases. Reduced LC tone drives G upward toward the bifurcation.

2.4 The Characteristic Equation and Hopf Bifurcation

Substituting $a(t) = e^{\lambda t}$ into the linearized equation yields:

$$\tau_{\text{iris}}\lambda + 1 + Ge^{-\lambda\delta} = 0 \quad (7)$$

At a Hopf bifurcation, a complex conjugate pair of eigenvalues crosses the imaginary axis. Substituting $\lambda = i\omega_c$ and separating real and imaginary parts:

$$\begin{aligned} 1 + G_c \cos(\omega_c\delta) = 0 &\Rightarrow \cos(\omega_c\delta) = -\frac{1}{G_c} \\ \tau_{\text{iris}}\omega_c = G_c \sin(\omega_c\delta) \end{aligned} \quad (8)$$

Dividing gives $\tan(\omega_c\delta) = -\tau_{\text{iris}}\omega_c$. For $\delta = 0.300$ s and $\tau_{\text{iris}} = 0.311$ s, numerical solution yields:

$$G_c = 2.318, \quad f_c = \frac{\omega_c}{2\pi} = 1.07 \text{ Hz} \quad (9)$$

The predicted oscillation frequency $f_c = 1.07$ Hz falls within the 1–3 Hz band of spontaneous pupillary oscillations observed clinically as hippus, cross-validating the model independently of the gain sweep. Above G_c , the nonlinear Hill saturation clamps the post-bifurcation instability into a stable limit cycle at this frequency, consistent with observed hippus amplitude.

2.5 Critical Slowing Down and the Recovery Time Constant

For $G < G_c$, the dominant eigenvalue λ^* is the root of the characteristic equation with the largest (least negative) real part. The recovery time constant is defined as

$$\tau_{\text{return}}(G) = -\frac{1}{\text{Re}(\lambda^*)} \quad (10)$$

As $G \rightarrow G_c^-$, $\text{Re}(\lambda^*) \rightarrow 0^-$ and $\tau_{\text{return}} \rightarrow \infty$. This is the critical slowing down phenomenon identified by Scheffer [10]: a system approaching a bifurcation recovers anomalously slowly from perturbation, and τ_{return} encodes that proximity as a continuous, measurable quantity.

The characteristic equation (7) admits no closed-form solution for general G , but the Lambert W function yields the exact eigenvalue through the following substitution. Setting $u = 1 + \tau_{\text{iris}}\lambda$, so that $\lambda = (u - 1)/\tau_{\text{iris}}$, and substituting into Eq. (7):

$$\begin{aligned} (u - 1) + 1 + Ge^{-(u-1)\delta/\tau_{\text{iris}}} &= 0 \\ ue^{u\delta/\tau_{\text{iris}}} &= -Ge^{\delta/\tau_{\text{iris}}} \end{aligned}$$

Multiplying both sides by $\delta/\tau_{\text{iris}}$ casts the left side into the form Xe^X :

$$\frac{u\delta}{\tau_{\text{iris}}} \exp\left(\frac{u\delta}{\tau_{\text{iris}}}\right) = -\frac{G\delta}{\tau_{\text{iris}}} e^{\delta/\tau_{\text{iris}}}$$

Since $Xe^X = Y$ implies $X = W(Y)$ by definition of the Lambert W function,

$$\frac{u\delta}{\tau_{\text{iris}}} = W_0\left(-\frac{G\delta}{\tau_{\text{iris}}} e^{\delta/\tau_{\text{iris}}}\right)$$

where W_0 denotes the principal branch of the Lambert W function. Recovering λ from $u = 1 + \tau_{\text{iris}}\lambda$ yields the dominant eigenvalue:

$$\lambda^*(G) = \frac{W_0\left(-\frac{G\delta}{\tau_{\text{iris}}} e^{\delta/\tau_{\text{iris}}}\right)}{\delta} - \frac{1}{\tau_{\text{iris}}} \quad (11)$$

Substituting into Eq. (10) gives the exact recovery time constant:

$$\tau_{\text{return}}(G) = \frac{\tau_{\text{iris}}}{1 - \frac{\tau_{\text{iris}}}{\delta} \operatorname{Re}\left[W_0\left(-\frac{G\delta}{\tau_{\text{iris}}} e^{\delta/\tau_{\text{iris}}}\right)\right]} \quad (12)$$

No approximation beyond the linearization of Section 2.3 enters this result. Equation (12) is the calibration function of the instrument and is used throughout this paper.

2.6 The Gain Inversion

The measurement principle works only if the map $\tau_{\text{return}}(G)$ is injective on the operative window $G \in [0.15, 2.1]$. Two distinct gain states that produce the same recovery time constant make the inversion ill-posed and the instrument undefined as a measurement device. Injectivity here is not a modeling assumption but follows from the Lambert W structure of Eq. (12).

As G increases from 0 toward G_c , the argument of W_0 in Eq. (12), namely $-\frac{G\delta}{\tau_{\text{iris}}} e^{\delta/\tau_{\text{iris}}}$, moves monotonically in the negative direction along the real axis. On the principal branch, $\operatorname{Re}[W_0(x)]$ is a strictly increasing function of x for $x \in (-1/e, 0)$, so $\operatorname{Re}[W_0]$ increases toward zero as the argument traverses this interval. The denominator of Eq. (12) decreases monotonically toward zero and τ_{return} increases without bound. The map $\tau_{\text{return}} : [0, G_c) \rightarrow [\tau_{\text{iris}}, \infty)$ is therefore strictly monotonically increasing, and every measured τ_{return} in the operative window corresponds to exactly one G . The inversion is well-posed by construction.

Given a measured $\tau_{\text{return}}^{\text{meas}}$, the gain estimate is recovered by solving

$$-\frac{1}{\operatorname{Re}(\lambda^*(G))} = \tau_{\text{return}}^{\text{meas}} \quad (13)$$

where $\lambda^*(G)$ is given by Eq. (11). This is a one-dimensional root-finding problem on a strictly monotone function, so convergence from any bracketed starting interval within $[0.15, 2.1]$ is guaranteed. The implementation in `analysis.py` solves Eq. (13) with a bracketed root-finder initialized at the window boundaries, propagating uncertainty in G from the residuals of the exponential envelope fits across the ten-pulse protocol.

The inversion does not remain equally well-conditioned across the full operative window. As $G \rightarrow G_c$, the Lambert W curve steepens and a small error in $\tau_{\text{return}}^{\text{meas}}$ produces a large displacement in the recovered G . This is a structural consequence of critical slowing down [10]: the divergence of τ_{return} that makes

bifurcation proximity detectable and the amplification of inversion sensitivity to noise are two expressions of the same underlying mechanism, and neither can be present without the other. The practical upper bound $G \approx 2.1$ established in Section 2.7 therefore marks not only the clinical feasibility limit of the measurement protocol but the onset of meaningful inversion uncertainty. For sessions approaching this boundary, a point estimate of G reported without an accompanying uncertainty interval is not a complete output of the instrument.

2.7 Validity Conditions

Two bounds define the operative measurement window theoretically, with a third imposed by the noise floor established in Section 4.4.

Below $G \approx 0.15$ the dominant eigenvalues are purely real and the recovery from perturbation is monotone, with no oscillatory envelope. The node-focus transition at $G \approx 0.15$ is where the eigenvalues first acquire a nonzero imaginary part and the oscillatory structure appears. Below this transition τ_{return} is well-defined but carries negligible gain information. The reason is visible directly from Eq. (12): at low G the argument of W_0 is small, $\text{Re}[W_0(-z)] \approx -z$ to first order, and the denominator of Eq. (12) varies slowly with G , producing a nearly flat calibration curve across this entire regime. Reliable gain estimation begins above this transition.

Above $G \approx 2.1$, τ_{return} exceeds approximately 10 s and recovery times become clinically impractical within a finite session protocol. This sets the upper operational bound. The theoretical measurement window is therefore $G \in [0.15, 2.1]$.

This window narrows under realistic noise conditions. Section 4.4 establishes that below $G \approx 0.5$, noise-induced displacement of the delayed area term generates perturbations comparable in magnitude to the recovery signal being fitted, and gain estimation becomes unreliable at any noise level. The practical operative window under realistic instrument noise is $G \in [0.5, 2.1]$, and it is this narrowed window that defines the honest clinical claim of the instrument.

Post-bifurcation states are identified by a spectral exclusion criterion applied before the PLR sweep begins. Sustained power in the 1–3 Hz band of the continuous pupil trace exceeding a defined threshold flags the session as operating near or above G_c and excludes τ_{return} estimation entirely. Section 4.3 validates this criterion computationally.

2.8 The Chromatic Signal ΔR

A subset of retinal ganglion cells, constituting approximately 0.4–1.5% of the human retinal ganglion cell population, express the photopigment melanopsin and are intrinsically photosensitive independently of rod and cone input [9]. These intrinsically photosensitive retinal ganglion cells (ipRGCs) produce a sustained membrane depolarisation that persists for seconds to tens of seconds after light offset, unlike the rapidly adapting responses of rods and cones. When this sustained drive reaches the Edinger-Westphal nucleus via the olivary pretectal nucleus, it produces a post-illumination pupil response (PIPR): a maintained constriction outlasting the stimulus by a duration far exceeding what the cone-driven sphincter pathway alone would produce. The PIPR measured at 5–7 s post-stimulus offset is dominated by melanopsin-driven ipRGC activity and is the standard metric used in validated chromatic pupillometry protocols [9].

The differential measurement ΔR exploits the steep wavelength dependence of melanopsin absorption. Human melanopsin has a peak absorption at approximately 479 nm [9]. At 627 nm, melanopsin sensitivity is approximately 1–2% of its peak value, making a 627 nm stimulus a cone-dominated drive with negligible melanopsin activation. At 480 nm, the full photoreceptor activation profile is as follows: melanopsin at peak sensitivity, M-cones at approximately 55% of their peak ($\lambda_{\text{max}} \approx 530$ nm), rods at approximately

57% of their scotopic peak (CIE 1951 scotopic luminosity function $V'(\lambda)$, normalised at 507 nm), and S-cones at approximately 22% of their peak (Stockman-Sharpe S-cone fundamental, normalised at 420 nm). The 627 nm control wavelength is chosen because at 555 nm, which maximises photopic cone drive, melanopsin retains approximately 15% of its peak sensitivity, sufficient to contaminate the differential measurement. At 627 nm this contamination falls to 1–2%, meeting the spectral isolation requirement for PIPR measurement.

The PIPR differential is formally defined as:

$$\Delta R = R_{480} - R_{627} \quad (14)$$

where R is the normalised pupil area relative to pre-stimulus baseline:

$$R = \frac{A_{\text{PIPR}} - A_0}{A_0} \times 100\% \quad (15)$$

with A_0 the mean pupil area over a 2 s pre-stimulus window and A_{PIPR} the mean pupil area over the 5–7 s post-stimulus offset window, consistent with the validated chromatic pupillometry convention [9]. A more negative ΔR indicates a larger sustained melanopsin-driven constriction relative to the cone-dominated baseline, indexing higher functional ipRGC gain.

3 Simulation Design

3.1 Model Implementation

The simulator integrates the full nonlinear Longtin-Milton delay-differential equation directly, retaining the Hill function in every time step rather than substituting the linearized deviation form. The governing equation for each step is:

$$\tau_{\text{iris}} \frac{dA}{dt} = -A(t) + A^* + \gamma [f(\Phi(t - \delta)) - f(A^*)] \quad (16)$$

where $f(\Phi) = c\theta^n / (\theta^n + \Phi^n)$ is evaluated at the delayed retinal flux $\Phi(t - \delta) = A(t - \delta)(1 + s(t - \delta))$. Integration uses a forward Euler scheme with time step $dt = 0.001$ s. The delay history is stored in a circular buffer of length $N_{\text{delay}} = \text{round}(\delta/dt) = 300$ samples, updated at each step. A physical boundary clamp enforces $A \geq 0.1 \text{ mm}^2$ to prevent negative areas under noise.

The scaling factor γ is computed analytically at initialization from the Hill derivative at the fixed point:

$$\gamma = \frac{G}{|f'(A^*)|} \quad (17)$$

$$f'(A^*) = -\frac{cn\theta^n (A^*)^{n-1}}{(\theta^n + (A^*)^n)^2} \quad (18)$$

This construction pins the linearized loop gain at A^* to exactly G for every value in the sweep, establishing an exact correspondence between the simulation and the Lambert W eigenvalue theory of Section 2.5. For small perturbations around A^* the model is mathematically equivalent to the linearized DDE with gain coefficient G ; for large perturbations and post-bifurcation dynamics the Hill saturation takes over, clamping the instability into a stable limit cycle above G_c rather than permitting unbounded growth. The Hill parameters (c, θ, n) therefore control the nonlinear saturation behavior but do not affect the local gain structure, which is entirely determined by γ .

When noise is active, zero-mean Gaussian noise scaled to $\sigma \cdot A^*$ is added to the delayed area value before the Hill function is evaluated, representing photon noise and iris mechanical variability in the measurement channel. The delayed value is clamped to 0.1 mm^2 after noise injection to maintain physical consistency.

3.2 Stimulus Protocol and Gain Sweep

Each simulation trial consists of a sequence of brief light pulses delivered to a system initialized at rest at A^* . Individual pulses have duration 200 ms and normalized intensity $s = 0.3$, producing a clearly resolved recovery trajectory while remaining small enough relative to A^* to keep the dominant dynamics within the linearized regime. The inter-pulse interval at each gain value is set adaptively to $T_{\text{IPI}} = 5 \times \tau_{\text{pred}}$, where τ_{pred} is the predicted recovery time constant computed from the exact Lambert W formula of Eq. (11), ensuring the system returns to within 1% of A^* before the next pulse arrives. Ten pulses are delivered per trial and τ_{return} is estimated from the mean of the individual recovery fits to reduce sensitivity to transient initialization effects.

The gain sweep runs from $G = 0.10$ to $G = 0.99 \cdot G_c = 2.295$ in uniform steps of $\Delta G = 0.05$, spanning the full range from below the node-focus transition through the near-bifurcation regime. For the noise robustness analysis, a zero-noise sweep is first run to produce the ground truth $\tau_{\text{return}}(G)$ curve that serves as the reference baseline in Section 4.4. The sweep is then repeated independently at four noise levels: $\sigma = 1\%$, 5% , 10% , 20% of A^* , where $\sigma = 1\%$ represents the near-zero instrumental noise floor rather than a distinct noise condition. At each step, zero-mean Gaussian noise at the specified level is added to the delayed area value before the Hill function is evaluated, representing photon noise and iris mechanical variability in the measurement channel. The noisy estimates at each σ level are then compared against the zero-noise baseline in Section 4.4 to establish the operative noise tolerance of the instrument.

Table 1: Simulation Parameters

Parameter	Symbol	Value
Iris time constant	τ_{iris}	0.311 s
Conduction delay	δ	0.300 s
Resting fixed point	A^*	15.0 mm^2
Hill coefficient	n	4
Hill threshold	θ	15.0 mm^2
Hill saturation	c	30.0 mm^2
Critical gain	G_c	2.318
Critical frequency	f_c	1.07 Hz
Node-focus transition	G_{nf}	0.15
Integration time step	dt	0.001 s
Stimulus duration	–	0.200 s
Stimulus intensity	s	0.3

The Hill parameters $\theta = 15.0 \text{ mm}^2$ and $c = 30.0 \text{ mm}^2$ used in the simulation differ from the canonical Longtin-Milton values ($n = 3$, $\theta = 50 \text{ mm}^2$, $c = 200 \text{ mm}^2$) [5]. This difference is inconsequential for the gain sweep because γ absorbs it entirely: what the simulation implements at A^* is the gain G , not any particular set of Hill parameters. The Hill parameters affect only the nonlinear saturation regime, relevant for the hippus dynamics of Section 4.3.

4 Results

4.1 Recovery Dynamics Across G

Figure 1 shows pupil area traces following a single 200 ms light pulse at five representative gain values spanning the measurement window. The grey shaded region marks the stimulus duration and the dashed line marks the resting area $A^* = 15 \text{ mm}^2$. For the four stable traces, constriction begins approximately $\delta = 300 \text{ ms}$ after stimulus onset, reflecting the neural conduction delay in the feedback loop, and reaches its minimum at roughly $t = 0.5 \text{ s}$, where both the delay and the finite iris response time τ_{iris} contribute to the observable trough position. At $G = 0.30$, the recovery appears nearly monotone. The eigenvalues are technically complex above the node-focus transition at $G \approx 0.15$, but the oscillatory structure at this gain value is too small to be clearly resolved visually. The ringing becomes unambiguous above $G \approx 0.60$, where the oscillatory envelope is clearly visible in the traces for $G = 0.60, 0.85,$ and 0.95 . As G increases through these values, the envelope stretches progressively while the oscillation frequency remains essentially invariant, consistent with the theoretical prediction that $\text{Im}(\lambda^*)$ stays close to $f_c = 1.07 \text{ Hz}$ across the full gain range. The envelope decay rate, governed by $\text{Re}(\lambda^*)$, varies with G and constitutes the measurement target. At $G = 2.35$, beyond the Hopf boundary $G_c = 2.318$, the system does not recover. The instability is bounded by the Hill nonlinearity into a stable limit cycle, with a visible transient growth phase over roughly the first four seconds before the oscillation amplitude saturates into persistent oscillations that fill the remainder of the simulation window. This is a qualitatively different dynamical regime from a slow recovery, and it defines the hard upper wall of the measurement envelope.

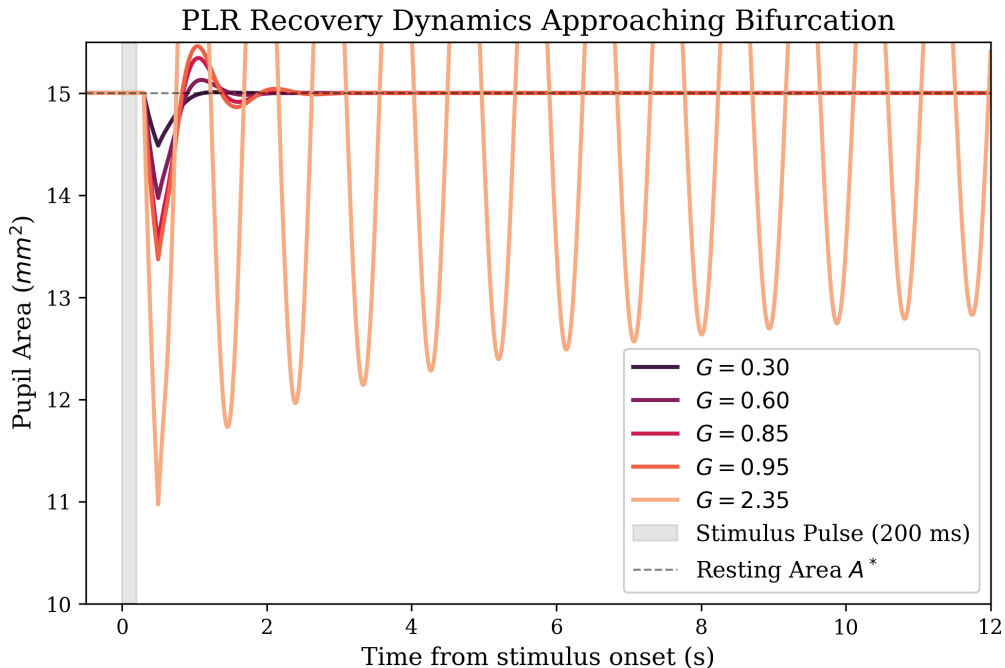


Figure 1: Pupil area recovery traces following a 200 ms light pulse at five representative gain values ($G = 0.30, 0.60, 0.85, 0.95, 2.35$). The oscillatory envelope slows as G increases, illustrating critical slowing down. At $G = 2.35 > G_c = 2.318$, the system enters a sustained limit-cycle oscillation bounded by the Hill nonlinearity rather than returning to equilibrium.

4.2 $\tau_{\text{return}}(G)$: Simulation vs Theory

Figure 2 shows the simulated τ_{return} values as orange dots against the exact Lambert W theoretical prediction of Section 2.5 as a black dashed curve, plotted on a log scale. Across the valid measurement window, the simulation closely follows the theoretical curve, confirming that the construction γ pins the linearized loop gain to G at every operating point. The dots sit systematically slightly above the curve throughout this range, which is physically expected because the Lambert W prediction is exact only for infinitesimal displacements and the finite-amplitude $s = 0.3$ stimulus drives the system into the nonlinear regime of the Hill function during recovery, extending τ_{return} slightly beyond the linear lower bound. Below the node-focus transition marked by the red dotted line, the simulated values diverge upward from the theoretical curve and the two separate visibly, reflecting the degenerate low-gain regime where the calibration curve is nearly flat and τ_{return} carries negligible gain information. As $G \rightarrow G_c = 2.318$, both simulation and theory rise steeply without bound, with the last simulated dots approaching the top of the measurement window and marking the onset of practical inaccessibility consistent with the upper operational bound of $G \approx 2.1$ stated in Section 2.7. The Lambert W curve is the calibration function that is inverted in the simulation to recover G from a measured τ_{return} , and the strict monotonicity visible across the valid window confirms that the inversion is well-posed throughout the clinical range.

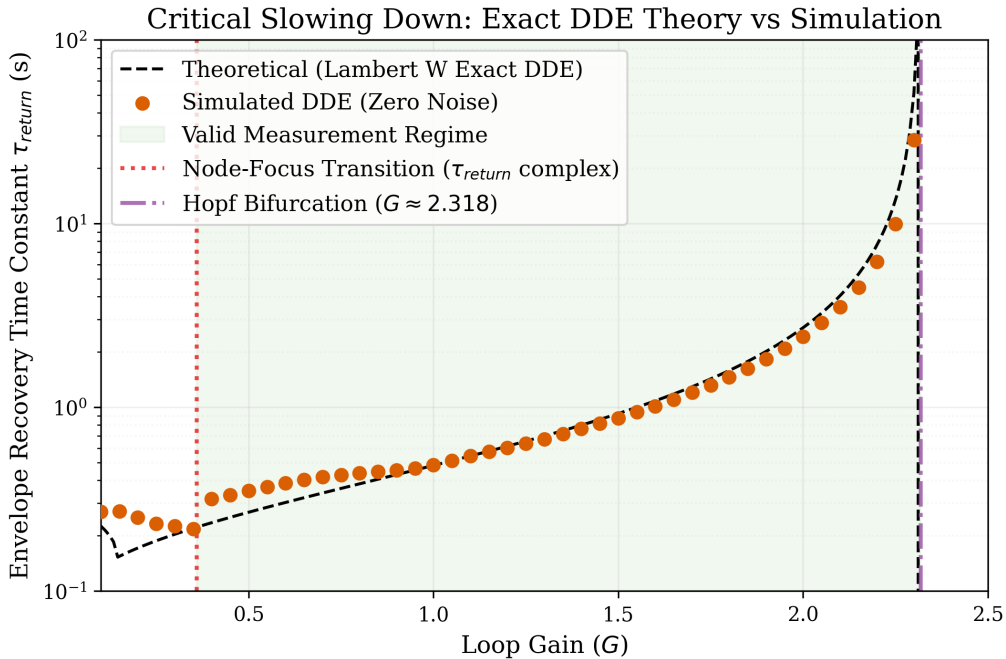


Figure 2: Simulated τ_{return} vs loop gain G (against the exact Lambert W theoretical prediction). The simulation closely tracks the theory; small deviations above it reflect finite-amplitude nonlinear effects from the $s = 0.3$ stimulus. The green shaded region marks the valid measurement window.

4.3 Hippus Emergence and Spectral Validation

Figure 3 shows the power spectral density of 60-second continuous pupil traces at two gain values. At $G = 0.50$, the spectrum shows no frequency-selective amplification, with power declining smoothly across the band and no peak structure visible anywhere in the 1–3 Hz clinical hippus range. At $G \approx 0.97 \cdot G_c = 2.249$, a broad spectral peak is centred at $f \approx 1.07$ Hz, coinciding with the theoretical Hopf frequency f_c derived in Section 2.4 and falling within the 1–3 Hz clinical hippus band. The breadth of the peak

is consistent with a lightly damped eigenmode: the system at this gain remains stable since $\text{Re}(\lambda^*)$ is still negative and no limit cycle exists, but broadband ambient noise is selectively amplified near f_c and attenuated elsewhere, producing a noise-driven resonance rather than spontaneous oscillation. The limit cycle and its sharp spectral signature emerge only above $G_c = 2.318$ where $\text{Re}(\lambda^*)$ crosses zero. The peak's location at exactly f_c , generated through a 60-second continuous trace with no stimulus forcing and entirely independently of the τ_{return} fitting procedure of Section 4.2, constitutes a cross-validation of the theoretical bifurcation geometry through a separate computational pathway. Sessions exhibiting elevated 1–3 Hz band power relative to the broadband trend are therefore operating near or above G_c , where τ_{return} estimation breaks down.

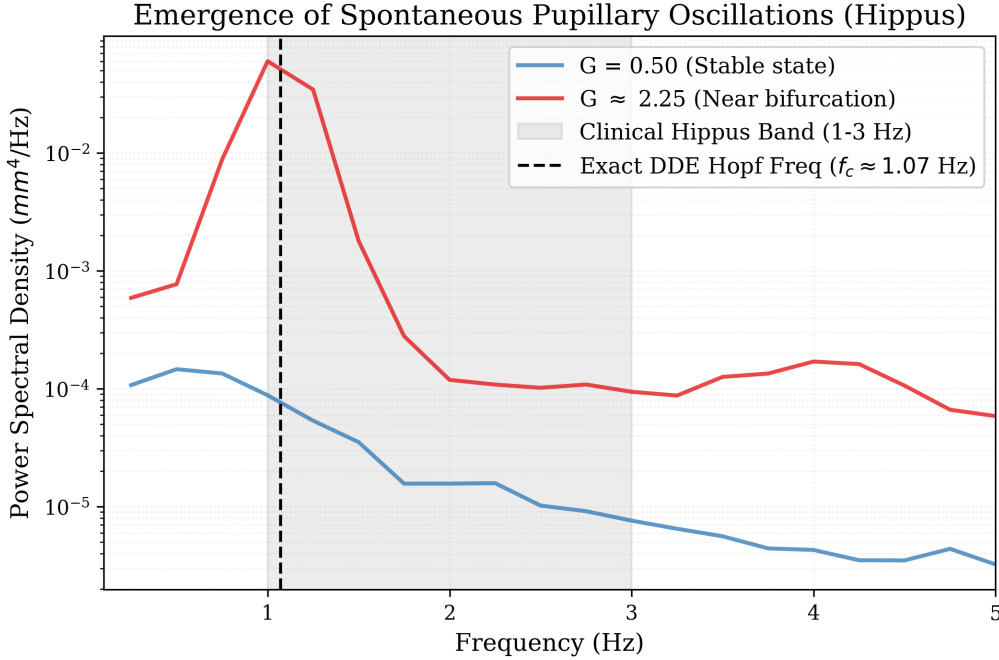


Figure 3: Power spectral density of the pupil area trace at $G = 0.50$ and at $G = 0.97G_c = 2.249$. Near bifurcation, ambient noise is selectively amplified at the lightly-damped eigenfrequency f_c , producing a spectral peak within the clinical hippus band without requiring a true limit cycle. The dashed vertical line marks f_c .

4.4 Noise Robustness

Figure 4 shows τ_{return} estimates under four noise levels plotted against the zero-noise simulation baseline on a log scale. The figure covers the operative measurement window from $G \approx 0.45$ to $G \approx 2.3$, with the low-gain failure regime below $G \approx 0.5$ falling outside the displayed range. Within the visible window, two noise regimes are distinguishable. For $\sigma \leq 5\%$, the estimates are essentially indistinguishable from the zero-noise baseline across the full operative range, with shading bands that remain narrow throughout. For $\sigma = 10\%$ and $\sigma = 20\%$, the estimates track the baseline with increasing spread, with $\sigma = 20\%$ producing clearly wider confidence intervals particularly above $G \approx 1.5$ and in the near-bifurcation region where inversion sensitivity amplifies noise into larger G uncertainty. The practical noise floor for a well-engineered NIR pupillometry system is $\sigma \leq 5\%$ of A^* , achievable with the silicon photodiode array specified in Section 5.2, and at this level the figure confirms reliable τ_{return} estimation across the full operative window $G \in [0.5, 2.1]$. The narrowing of this window relative to the theoretical $[0.15, 2.1]$ reflects the low-gain structural failure established separately, and the near-bifurcation upper

bound reflects the inversion sensitivity discussed in Section 2.6 rather than noise-induced estimation failure.

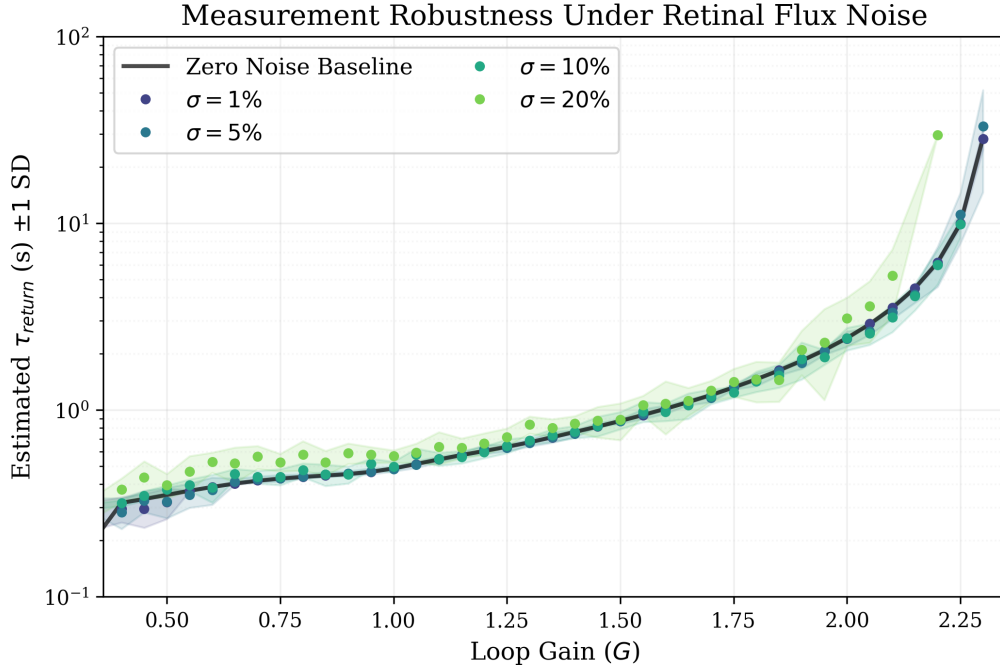


Figure 4: Noise robustness of τ_{return} estimates as mean estimates of four noise levels ($\sigma = 1\%$, 5% , 10% , 20% of A^*) with ± 1 SD shading; the black solid curve is the zero-noise simulation baseline. Reliable estimation begins above $G \approx 0.5$ for realistic noise levels ($\sigma \leq 5\%$).

5 Instrument Design

5.1 Measurement Architecture

The instrument produces three output variables that are jointly constrained by LC tonic state through the physiological model of Section 2, each operating on a distinct timescale and through a distinct physiological pathway.

Resting pupil diameter D_0 reflects the ambient sympathetic dilator tone set by tonic LC output. Within the present framework it provides a direct estimate of the resting fixed point A^* from which the loop gain analysis proceeds, connecting the geometric measurement to the operating-point argument of Section 2.3.

The recovery time constant τ_{return} is the central contribution of this project. Extracted from the PLR re-dilation envelope following each light pulse, it encodes loop gain G via the Lambert W inversion of Section 2.6 and constitutes a medium-timescale dynamical index of LC-linked arousal state that is inaccessible to instruments that discard the recovery trajectory.

The chromatic post-illumination pupil response $\Delta R = R_{480} - R_{627}$ is the difference in sustained constriction amplitude between 480 nm and 627 nm stimulation measured in the 5–7 s post-offset window [9]. It reflects melanopsin-driven ipRGC activation through a pathway functionally independent of the LC-sympathetic mechanism, providing a third measurement dimension for within-session consistency verification.

The algebraic coupling between D_0 and G follows directly from Section 2.3: both are set by A^* through the operating-point relationship, so the A^* implied by D_0 and the A^* implied by the G estimate

via the Hill function must be mutually consistent within a session. Systematic disagreement between the two quantities flags individual differences in iris mechanics or non-LC arousal confounds, providing a self-consistency check unavailable to single-variable pupillometers.

5.2 Optical Transduction

The fundamental transduction problem is extracting pupil area as a continuous time series with sufficient spatial resolution and temporal bandwidth to support exponential envelope fitting across the full operative range of τ_{return} . This problem cannot be solved with a single-element photodiode. A device of this type integrates total reflected flux over its active area and produces a scalar output that encodes luminous power, not geometry. Pupil diameter and area are inaccessible from this signal independently of scene illuminance, which rules out the simplest photodiode implementation entirely.

Among spatially resolving silicon detector architectures, position-sensitive detectors exhibit nonlinear response at pupil edges and are calibrated for centroid tracking rather than edge detection; quadrant detectors provide insufficient spatial sampling for sub-millimeter diameter resolution; two-dimensional CCD or CMOS imaging arrays introduce unnecessary bandwidth and processing constraints for what is fundamentally a one-dimensional diameter measurement along each axis. The appropriate architecture is a linear silicon photodiode array, which produces a one-dimensional intensity profile across the pupil plane from which edge positions are extracted by threshold detection, giving diameter directly in a single axis per array. Two arrays in orthogonal configuration give both pupil axes simultaneously, handling non-circular pupils and providing redundancy against blink artifacts.

This architecture is validated by the clinical pupillometer of [3], which used dual 256-element linear CCD arrays in orthogonal configuration illuminated by infrared diodes, achieving diameter accuracy better than ± 0.1 mm over the range 3.2–9.5 mm at a maximum sampling rate of 50 Hz. At 50 Hz, a 200 ms stimulus pulse is resolved across 10 samples, and re-dilation envelopes spanning 1 to 50 s are resolved across 50 to 2500 samples, providing adequate temporal resolution for exponential fitting across the full operative gain range $G \in [0.5, 2.1]$.

The tracking channel illuminates the eye off-axis at 850 nm in a dark-pupil geometry: the pupil appears as a dark region against the brighter iris in the reflected intensity profile, producing sharp edges for threshold-based diameter extraction. Silicon photodiodes have peak spectral responsivity near 850 nm, approximately 0.5–0.6 A/W, maximizing signal-to-noise ratio in the tracking channel. With adequate illumination power the shot noise floor places measurement uncertainty comfortably below $\sigma = 5\%$ of A^* , the noise level at which Section 4.4 shows reliable τ_{return} estimation beginning at $G \approx 0.5$. The complete signal chain from illumination to analysis input is: 850 nm LED array illuminating the corneal plane \rightarrow dual orthogonal linear CCD arrays reading the reflected intensity profile \rightarrow threshold-based edge detection algorithm \rightarrow diameter in mm along each axis \rightarrow pupil area $A(t)$ in mm^2 \rightarrow DDE analysis pipeline in `analysis.py`.

The 850 nm tracking wavelength is spectrally isolated from both stimulus channels by construction. The nearer stimulus channel at 627 nm is separated from the tracking channel by 223 nm, a gap large enough that detector cross-response to stimulus light is negligible without additional optical filtering.

The first stimulus channel operates at 480 nm, near the melanopsin absorption peak at 479 nm. This wavelength drives both cone pathways and ipRGC melanopsin simultaneously, producing the fast cone-mediated PLR used for τ_{return} extraction and the sustained melanopsin-driven post-illumination pupil response used for ΔR measurement.

The second stimulus channel operates at 627 nm, driving cone pathways with negligible melanopsin activation. It produces a PLR through the cone-sphincter pathway with minimal sustained response, serving as the spectrally controlled reference. The chromatic signal $\Delta R = R_{480} - R_{627}$ is extracted

from the 5–7 s post-offset window, where the fast cone-driven component has decayed and the sustained melanopsin-driven component dominates [9].

5.3 Acquisition Protocol

A session proceeds in three sequential blocks with a total duration of approximately 30 minutes.

The baseline block runs for 30 s under ambient illumination with no stimulation. Pupil area is sampled continuously at 50 Hz and D_0 is extracted as the mean over this window, converted to diameter. Spectral analysis of the same trace provides the hippus exclusion flag: the 1–3 Hz band-power ratio relative to the broadband baseline is compared against a predefined threshold, and sessions exceeding it are operating near or above G_c and excluded from τ_{return} estimation before the PLR sweep begins.

The PLR sweep block delivers ten light pulses per stimulus intensity level using an adaptive inter-pulse interval set to $5 \times \tau_{\text{pred}}$ at each gain value, ensuring the pupil returns to within 1% of A^* before the next pulse arrives. τ_{return} is extracted from the re-dilation envelope of each pulse by exponential fitting and the ten estimates are mean-aggregated per intensity level to reduce sensitivity to transient initialisation effects. Total duration across the operative window $G \in [0.5, 2.1]$ is approximately 26.6 minutes, with session time concentrated at high gain values where the adaptive interval is longest: at $G = 2.05$, $\tau_{\text{pred}} \approx 10$ s and the inter-pulse interval reaches 50 s, accounting for the bulk of the sweep.

The chromatic block interleaves 480 nm and 627 nm pulses with a minimum inter-stimulus interval of 70 s to allow full melanopsin recovery between trials [9]. PIPR amplitude is extracted in the 5–7 s post-offset window for each wavelength and ΔR is computed as the difference. This block takes approximately 2 minutes for two trials per wavelength.

5.4 Output and Inversion

A completed session produces a four-tuple $(D_0, G, \Delta R, \text{flag})$ with associated uncertainty estimates. D_0 is reported directly in mm from the baseline block mean. G is recovered from the mean τ_{return} estimate by direct numerical evaluation of the Lambert W inversion of Eq. (11): given a measured τ_{return} , the corresponding G is the unique root of the characteristic equation (13) within the operative window $G \in [0.5, 2.1]$, located by a bracketed root-finder initialised at those boundaries. Uncertainty in G is propagated from the residuals of the exponential envelope fits across the ten pulses. ΔR is reported in mm^2 of sustained constriction difference. The hippus flag is binary, set during the baseline spectral analysis.

The self-consistency check is not a redundancy check. Both D_0 and G are functions of the same resting fixed point A^* , set by tonic LC output through the sympathetic dilator pathway, so the A^* inferred from D_0 via pupil geometry and the A^* inferred from the G estimate through the Hill function operating-point relationship of Section 2.3 must agree within a session. This coupling is exploited as an internal consistency constraint: the instrument is checking whether its two measurement channels are seeing the same underlying physiological state. Sessions where the two A^* estimates disagree beyond a calibrated tolerance are flagged for review, identifying subjects whose iris mechanics or ambient arousal state fall outside the model’s validity assumptions before clinical interpretation proceeds.

6 Limitations

6.1 Validity Range of τ_{return}

The τ_{return} measurement is operative only for $G \in [0.5, 2.1]$ under realistic noise conditions, narrowed from the theoretical window $[0.15, 2.1]$ by the noise floor result of Section 4.4. This boundary is structural

rather than instrumental: the information content of τ_{return} as a gain index is concentrated near the bifurcation, where the Lambert W curve is steep and small changes in G produce large, resolvable changes in τ_{return} . In the low-gain regime the curve is flat and τ_{return} carries negligible gain information regardless of measurement quality. The clinical consequence is that hyperarousal states, where high tonic LC output suppresses G well below 0.5, are not resolved by this instrument. The experimental hypothesis this boundary generates is that pharmacological elevation of tonic LC tone to levels corresponding to $G < 0.5$ should produce τ_{return} values indistinguishable from those at $G \approx 0.5$, a prediction testable under controlled noradrenergic manipulation.

6.2 Coupling Between D_0 and G

Both D_0 and G are functions of the resting fixed point A^* , which is set by LC tone through the sympathetic dilator pathway. They are not independent measurements. Within a session this coupling is exploited as the self-consistency check of Section 5.4; across subjects it limits the interpretability of absolute G values without individual calibration. Individual differences in baseline iris viscoelastic mechanics, resting sympathetic tone from non-LC sources, and iris pigmentation all modulate A^* and therefore confound the cross-subject mapping from D_0 to G . The experimental hypothesis is that the D_0 - G self-consistency residual, defined as the discrepancy between A^* inferred from geometry and A^* inferred from the gain estimate, should correlate with known sources of inter-individual iris mechanical variability, and that this residual should be stable within a subject across sessions but variable across subjects.

6.3 LC Non-Specificity

The pupil reflects distributed arousal network activity rather than LC spiking rate in isolation. Cholinergic tone from the basal forebrain, cognitive load, emotional valence, and task engagement all modulate pupil size and PLR dynamics through pathways that partially overlap with the LC-sympathetic channel [6]. The instrument therefore indexes LC-linked arousal state, which is a clinically meaningful quantity, but it does not provide a readout specific to LC firing rate. This distinction is not resolvable within a purely pupillometric framework. The experimental hypothesis it generates is that G estimates should show selective sensitivity to noradrenergic pharmacological challenge (atomoxetine, reboxetine, clonidine) relative to cholinergic challenge (scopolamine, physostigmine), a dissociation that would establish the degree to which G tracks the noradrenergic component of the distributed arousal signal specifically.

6.4 Individual τ_{iris} Variance

The inversion formula uses a fixed $\tau_{\text{iris}} = 0.311$ s from the Longtin-Milton parameterization [5]. Individual variation in iris viscoelastic mechanics introduces systematic bias in G estimates whose magnitude depends on how far a given subject's true τ_{iris} departs from this value. The bias propagates through the Lambert W inversion nonlinearly: at moderate G the calibration curve is steep and a small τ_{iris} error produces a proportionally small G error; near G_c the curve diverges and the same error produces a large G displacement. The experimental hypothesis is that individual τ_{iris} can be estimated from a brief dark-adapted recovery measurement in the absence of a light stimulus, and that substituting subject-specific τ_{iris} values into the inversion reduces the variance of repeated G estimates within a subject below the variance obtained with the fixed population value.

6.5 Parameter Sensitivity: δ and τ_{iris}

The critical gain G_c and the Hopf frequency f_c both depend on the ratio $\delta/\tau_{\text{iris}}$, and the Lambert W calibration curve shifts with any variation in these parameters. The present analysis treats $\delta = 0.300$ s and $\tau_{\text{iris}} = 0.311$ s as fixed constants, but in a human population both will vary. Correlated variation in δ and τ_{iris} , where both increase or decrease together, produces a smaller shift in G_c than anticorrelated variation because the ratio $\delta/\tau_{\text{iris}}$ is partially stabilized. A sensitivity analysis sweeping both parameters over $\pm 20\%$ of their nominal values is fully computable within the existing simulation framework and constitutes a natural extension of the present work. The experimental hypothesis is that population-level variation in δ and τ_{iris} , measurable from latency and recovery kinetics respectively, produces predictable shifts in the individual G_c value that can be corrected by subject-specific parameterization of the calibration curve.

6.6 Hill Parameter Choice and Spectral Threshold

The simulation uses Hill parameters $n = 4$, $\theta = 15$ mm², $c = 30$ mm², which differ from the canonical Longtin-Milton values ($n = 3$, $\theta = 50$ mm², $c = 200$ mm²) [5]. As established in Section 3, the γ construction makes the local gain structure independent of this choice. The nonlinear saturation behavior above G_c , however, does depend on the Hill parameters: the limit cycle amplitude and waveform of the simulated hippus in Figure 1 are specific to the chosen (n, θ, c) combination. The spectral exclusion threshold defined in Section 2.7 is therefore calibrated to these simulation parameters and would require retuning if the Hill parameters were changed to match a different physiological dataset. The experimental hypothesis is that the hippus limit cycle amplitude predicted by the nonlinear model under the chosen parameters can be compared against published clinical hippus amplitude measurements to assess whether the chosen Hill parameters produce physiologically realistic post-bifurcation behavior.

6.7 G as a Dimensionless Index: No Firing Rate Mapping

The theoretical framework establishes G as a dimensionless loop gain set by the product of iris mechanical sensitivity and sympathetic dilator tone, the latter driven by tonic LC output. The paper provides a physiological argument for why higher LC firing rates produce lower G values through the operating-point mechanism of Section 2.3, but it does not provide a quantitative mapping from LC firing rate in spikes/s to a specific G value. Such a mapping would require knowing the gain of the sympathetic preganglionic pathway, the α_1 adrenoceptor sensitivity of the iris dilator, and the relationship between LC discharge rate and norepinephrine concentration at the neuroeffector junction, none of which are fixed by the current framework. The instrument therefore produces a dimensionless arousal index whose ordinal properties are theoretically grounded but whose absolute calibration against LC electrophysiology remains an open experimental question. The hypothesis this generates is precise: simultaneous recording of LC single-unit activity and τ_{return} -derived G estimates in a non-human primate preparation, following the methodology of [4], would produce a monotonic G -to-firing-rate mapping whose slope encodes the sympathetic pathway gain.

6.8 Computational Scope

All results in this paper are theoretical and computational. The physiological parameters are drawn from established literature, the simulation validates the Lambert W prediction numerically across the full gain range, and the instrument design is grounded in validated hardware precedents. No human subject data exist for this instrument and no empirical validation of the τ_{return} -to- G mapping has been performed. Three experimental hypotheses follow directly from the framework and would constitute its empirical validation: that τ_{return} changes monotonically and in the predicted direction under pharmacological

manipulation of tonic LC-NE tone; that the τ_{return} -derived G estimate correlates with simultaneous LC-BOLD signal measured by high-resolution neuroimaging following the protocol of [7]; and that test-retest reliability of G within a subject exceeds the between-subject variance, establishing the index as a stable individual trait measure rather than a noise floor artifact.

6.9 The Light Pulse as an Arousal Confound

The 200 ms light stimulus is treated in the model as a purely retinal perturbation to the pupil feedback loop: it shifts the effective retinal flux $\Phi(t-\delta)$, drives the Hill function away from its operating point, and the recovery trajectory encodes G as the system returns to A^* . This treatment is incomplete at the level of the full physiological circuit. The pretectal olivary nucleus, described in Section 2.1 as the first relay of the afferent PLR arc, projects not only to the Edinger-Westphal nucleus but also to arousal-modulating nuclei including the LC itself through collateral projections in the ascending arousal network. A 200 ms light pulse is therefore not only a retinal perturbation to the sphincter feedback loop but potentially a brief sensory arousal event that transiently elevates LC firing rate, shifts sympathetic dilator tone, and displaces G during the recovery trajectory being measured. The model cannot distinguish between a recovery trajectory shaped by the pre-stimulus G and one partially contaminated by a stimulus-evoked G transient. The adaptive inter-pulse interval design of Section 5.3 provides a partial structural mitigation: spacing pulses at $5 \times \tau_{\text{pred}}$ allows any arousal transient driven by a preceding pulse to decay before the next recovery measurement begins, preventing transient accumulation across the ten-pulse protocol. Within a single recovery trajectory the confound remains uncontrolled. The experimental hypothesis it generates is that τ_{return} estimates should show a systematic dependence on stimulus intensity exceeding the operating-point shift predicted by the linearized model alone, a dissociation testable by measuring τ_{return} across a range of stimulus intensities within the same subject while holding independently estimated G constant through pharmacological stabilization of LC tone.

7 Conclusion

Standard pupillometry discards the dynamical structure of the reflex. The pupil does not merely dilate and constrict; it recovers along a trajectory whose time constant is set by the loop gain of a nonlinear delayed feedback system, and that gain is physiologically controlled by tonic LC output through the sympathetic dilator pathway. Treating the pupil as a static aperture is therefore not just a simplification but also a loss of information.

The classical approximation placed the bifurcation at $G = 1$ by assuming zero delay; the exact Lambert W solution of the DDE characteristic equation places it at $G_c = 2.318$, relocates the informative measurement regime to $G \in [0.5, 2.1]$, and provides a calibration curve whose steep gradient near bifurcation makes the arousal index most sensitive precisely where clinical resolution matters most. The simulation confirms that τ_{return} tracks this curve closely across the operative window and survives realistic noise. The spectral resonance at $f_c = 1.07$ Hz near $G = 2.249$, arising through a completely independent computational pathway, cross-validates the bifurcation geometry without appeal to the fitting procedure.

Because resting diameter, recovery dynamics, and chromatic response are all constrained by the same underlying LC state, the three-variable instrument architecture is self-consistent rather than additive: the algebraic coupling between D_0 and G through the shared fixed point A^* is a consequence of the physiology, and the instrument exploits it as an internal validity check on every session rather than treating it as a nuisance. The framework thus generates three claims: τ_{return} should change monotonically under noradrenergic pharmacological challenge, the τ_{return} -derived G estimate should correlate with simultaneous LC-BOLD signal, and repeated G measurements within a subject should show lower variance than

measurements across subjects.

code availability:

All simulation code, analysis scripts, and figure generation code are openly available at: https://github.com/Hidokun/Critical_Slowing_Down_PLR_Analysis_Simulation

References

- [1] Gary Aston-Jones and Jonathan D. Cohen. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28:403–450, 2005. doi: 10.1146/annurev.neuro.28.061604.135709.
- [2] Mark S. Gilzenrat, Sander Nieuwenhuis, Marieke Jepma, and Jonathan D. Cohen. Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cognitive, Affective, & Behavioral Neuroscience*, 10(2):252–269, 2010. doi: 10.3758/CABN.10.2.252.
- [3] D. P. Jones, S. J. Harris, and S. R. Scicinski. Pupillometer for clinical applications using dual 256-element linear CCD arrays. *Medical and Biological Engineering and Computing*, 30(5):487–490, 1992. doi: 10.1007/BF02457826.
- [4] Siddhartha Joshi, Yin Li, Rishi M. Kalwani, and Joshua I. Gold. Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, 89(1):221–234, 2016. doi: 10.1016/j.neuron.2015.11.028.
- [5] André Longtin and John G. Milton. Modelling autonomous oscillations in the human pupil light reflex using nonlinear delay-differential equations. *Bulletin of Mathematical Biology*, 51(5):605–624, 1989. doi: 10.1007/BF02459969.
- [6] Maxime Megemont, Jim McBurney-Lin, and Hongdian Yang. Pupil diameter is not an accurate real-time readout of locus coeruleus activity. *eLife*, 11:e70510, 2022. doi: 10.7554/eLife.70510.
- [7] Peter R. Murphy, Redmond G. O’Connell, Michael O’Sullivan, Ian H. Robertson, and Joshua H. Balsters. Pupil diameter covaries with BOLD activity in human locus coeruleus. *Human Brain Mapping*, 35(8):4140–4154, 2014. doi: 10.1002/hbm.22466.
- [8] Matthew R. Nassar, Katherine M. Rumsey, Robert C. Wilson, Kinjan Parikh, Benjamin Heasley, and Joshua I. Gold. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, 15(7):1040–1046, 2012. doi: 10.1038/nn.3130.
- [9] Martina Romagnoli, Luigi Bonarini, Maurizio Baccega, Stefano Govoni, and Stefania Bianchi-Marzoli. Chromatic pupillometry for evaluating melanopsin retinal ganglion cell function in Alzheimer’s disease and other neurodegenerative disorders: a review. *Frontiers in Psychology*, 14:1295129, 2024. doi: 10.3389/fpsyg.2023.1295129.
- [10] Marten Scheffer, Jordi Bascompte, William A. Brock, Victor Brovkin, Stephen R. Carpenter, Vasilis Dakos, Hermann Held, Egbert H. van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461:53–59, 2009. doi: 10.1038/nature08227.