Diffusions with Stochastic Resetting

Nonequilibrium Steady States and the Optimal Restart Rate

T. Zamrik

14-JUL-2026

1. Abstract

We study diffusions subject to Poissonian resetting: at rate r the process is returned instantaneously to a fixed point x_0. Two consequences separate this from the diffusions of classical stochastic analysis. First, the reset generator $\mathcal{L}_r f = \mathcal{L} f + r\,[f(x_0) - f(x)]$ is non-local and, even when $\mathcal{L}$ is self-adjoint in its natural weight, $\mathcal{L}_r$ is not; the stationary state is consequently a genuine **nonequilibrium steady state** carrying a nonzero probability current, and cannot be written as a Gibbs measure. For Brownian motion we obtain the stationary density in closed form — a cusped exponential $p_{\mathrm{ss}}(x) = \frac{1}{2}\sqrt{r/D}\, e^{-\sqrt{r/D}\,|x-x_0|}$ — and show that its current jumps by exactly r across the reset point, the teleported flux closing the balance. Second, resetting changes first-passage times qualitatively. A renewal argument gives the mean first-passage time under restart as $\langle T_r \rangle = (1 - \tilde{T}(r))/(r\,\tilde{T}(r))$, where $\tilde{T}$ is the Laplace transform of the un-restarted passage time, and we prove the resulting sharp criterion: **restart reduces the mean first-passage time if and only if the un-restarted passage time has coefficient of variation exceeding one**, and at the optimal rate the restarted passage time has coefficient of variation exactly one. Numerical experiments — using exact stationary sampling and an exact renewal simulation, so that no time-discretisation bias enters — confirm the criterion, locate the optimum, and verify the $\mathrm{CV} = 1$ signature to within 3.4×10^{-3}.

2. Introduction

A diffusion left to itself explores. A diffusion that is periodically dragged back to where it started does something different, and the difference is not a perturbation of the first behaviour but a change in kind.

The mechanism is elementary to state. Let X_t be a diffusion on $\mathbb{R}$ and let resetting events arrive as a Poisson process of rate r; at each event the process is placed instantaneously at a fixed point x_0. Between events it diffuses as before. Nothing in this description is exotic, and yet the resulting process fails two properties that the classical theory takes for granted. Its generator is non-local, because the value of a function at the reset point enters the evolution at every x. And its stationary state, when one exists, is not an equilibrium: probability circulates, being carried outward by diffusion and returned by teleportation, and the circulation never stops.

The phenomenon has been studied intensively since Evans and Majumdar's foundational work [1, 2], and it has proved to be more than a curiosity. Restart is the mechanism behind the Luby scheme in randomised algorithms [8], the unbinding of an enzyme from an unproductive substrate in the Michaelis–Menten cycle [7], and the decision of a searcher to abandon a fruitless region and return to base [1]. In each case the question is the same, and it is a sharp one: **does restarting help, and if so how often should one restart?**

The answer, remarkably, does not depend on the details of the underlying process. It depends on a single dimensionless number — the coefficient of variation of the un-restarted completion time — and the criterion is exact. That is the result this note is organised around.

Section 3 constructs the reset process and its generator, and establishes the sense in which self-adjointness fails. Section 4 solves for the stationary state of reset Brownian motion, exhibits its probability current, and shows that detailed balance is broken. Section 5 derives the mean first-passage time under restart by a renewal argument. Section 6 proves the CV criterion and its companion, that the optimum is characterised by the restarted process having unit coefficient of variation. Section 7 works the drifted-Brownian case, where every object is explicit. Section 8 reports the numerical experiments. Section 9 concludes.

3. The Reset Process and Its Generator

3.1 Construction

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ carry a Brownian motion W and an independent Poisson process N of rate $r > 0$. Let X solve

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dW_t \tag{3.1}$$

between the jump times of N, and at each jump time set $X_t = x_0$. We call the resulting càdlàg process $X^{(r)}$ the **diffusion with resetting at rate r to x_0**.

$X^{(r)}$ is a strong Markov process: the reset clock is memoryless, so the future depends on the past only through the current position. It is a piecewise-diffusive Markov process — a diffusion interrupted by a jump mechanism whose target is deterministic.

3.2 The Generator Is Non-Local

For $f \in C_b^2(\mathbb{R})$, decompose the increment over $[0, h]$ according to whether a reset has occurred. A reset occurs with probability $rh + o(h)$, in which case the process sits at x_0; with probability $1 - rh + o(h)$ it does not, and the process has diffused. Hence

$$\mathbb{E}_x[f(X_h^{(r)})] = (1 - rh)\big[f(x) + h\,\mathcal{L}f(x)\big] + rh\,f(x_0) + o(h), \tag{3.2}$$

and dividing by h and letting $h \downarrow 0$:

Proposition 1 (Reset generator). The generator of $X^{(r)}$ is

$$\mathcal{L}_r f(x) = \mathcal{L} f(x) + r\,[f(x_0) - f(x)], \qquad \mathcal{L} = \mu(x)\,\partial_x + \tfrac{1}{2}\sigma^2(x)\,\partial_{xx}. \tag{3.3}$$

The second term is the entire story. It is not a differential operator: it evaluates f at a point x_0 that has nothing to do with x. The generator is therefore **non-local**, and the resolvent, the eigenfunctions, and every spectral statement inherit that non-locality.

The adjoint acts on densities and carries the same information in transposed form:

$$\partial_t p(x,t) = \mathcal{L}^* p(x,t) \;-\; r\,p(x,t) \;+\; r\,\delta(x - x_0), \tag{3.4}$$

which we read as: probability is *removed uniformly* at rate r from every point, and *reinjected as a point source* of the same total mass at x_0. Mass is conserved — $\int(\mathcal{L}^* p - rp + r\delta)\,dx = 0$ — but not locally. That is the seed of everything that follows.

3.3 Failure of Self-Adjointness

Suppose $\mathcal{L}$ is self-adjoint with respect to a speed measure m, as it is for any one-dimensional diffusion in its natural scale. Self-adjointness of $\mathcal{L}_r$ would require

$$\langle \mathcal{L}_r f, g\rangle_m = \langle f, \mathcal{L}_r g\rangle_m \quad \Longleftrightarrow \quad f(x_0)\int g\,dm - \langle f, g\rangle_m \;=\; g(x_0)\int f\,dm - \langle f, g\rangle_m, \tag{3.5}$$

that is, $f(x_0)\int g\,dm = g(x_0)\int f\,dm$ for all f, g — which fails as soon as f and g are chosen with, say, $f(x_0) = 0 \neq \int f\,dm$. So $\mathcal{L}_r$ is not self-adjoint for any $r > 0$.

This is not a technicality. Self-adjointness is precisely the operator-level statement of reversibility, and its failure is precisely the statement that the stationary state, when it exists, will not satisfy detailed balance. We verify this directly in the next section by exhibiting the current.

4. The Nonequilibrium Steady State

Take $\mathcal{L} = D\,\partial_{xx}$ — driftless Brownian motion with diffusion coefficient D, so $\sigma^2 = 2D$ — and reset to x_0 at rate r. Free Brownian motion has no stationary distribution; it spreads without limit. Resetting supplies the confinement.

4.1 The Stationary Density

Theorem 1 (Evans–Majumdar [1]). The reset process has the unique stationary density

$$p_{\mathrm{ss}}(x) \;=\; \frac{\alpha_0}{2}\,e^{-\alpha_0\,|x - x_0|}, \qquad \alpha_0 = \sqrt{\frac{r}{D}}. \tag{4.1}$$

Proof. Set $\partial_t p = 0$ in the forward equation:

$$D\,\partial_{xx}p_{\rm ss} - r\,p_{\rm ss} + r\,\delta(x-x_0) = 0. \tag{4.2}$$

Away from x_0 this is $\partial_{xx}p = (r/D)\,p$, whose decaying solutions are $C_\pm e^{\mp\alpha_0(x-x_0)}$ on the two half-lines. Continuity at x_0 forces $C_+ = C_- = C$. Integrating the equation across an ε-neighbourhood of x_0 and letting $\varepsilon \downarrow 0$ leaves $D\,[\partial_x p]_{x_0^-}^{x_0^+} = -r$, i.e. $-2D\alpha_0 C = -r$, so $C = r/(2D\alpha_0) = \alpha_0/2$. This normalises.

Two features deserve emphasis. The density is **not Gaussian** but a two-sided exponential, so it has heavier tails than any equilibrium confining potential of quadratic type would produce. And it has a **cusp** at the reset point: $p_{\rm ss}$ is continuous there but $\partial_x p_{\rm ss}$ jumps, which is the density's memory of the point source.

4.2 The Current, and the Death of Detailed Balance

The stationary diffusive current is $J(x) = -D\,\partial_x p_{\rm ss}(x)$. Differentiating Theorem 1:

$$J(x) \;=\; \operatorname{sgn}(x-x_0)\,\frac{r}{2}\,e^{-\alpha_0|x-x_0|}. \tag{4.3}$$

This is nonzero for every $x \neq x_0$. In an equilibrium steady state the current vanishes identically — that *is* detailed balance. Here it does not vanish anywhere, and its structure is transparent: probability flows **away** from the reset point on both sides, at a rate that decays with distance, and the entire outward flux is returned by the resetting teleportation. Note that $J(x_0^+) - J(x_0^-) = r$: the current jumps by exactly the reset rate across the source, which is the same ε-neighbourhood computation as in the proof of Theorem 1, read in flux form.

So the process is stationary but not reversible. Its steady state is maintained by a permanent circulation — outward by diffusion, inward by teleportation — and sustaining that circulation costs something. This is the defining signature of a nonequilibrium steady state.

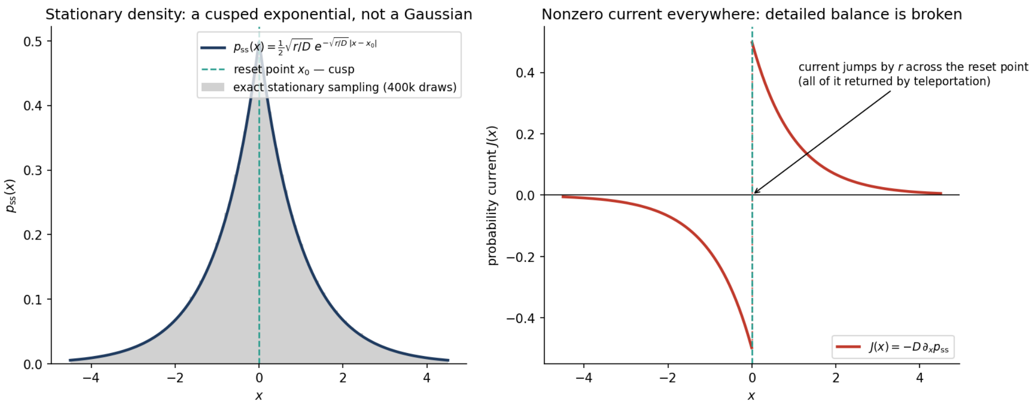

Figure 1: Reset Brownian motion at $r = 1$, $D = 1$, $x_0 = 0$. **Left:** the stationary density is a cusped exponential; the histogram is 400,000 exact stationary draws, obtained by noting that the age of the current excursion is $\mathrm{Exp}(r)$ by memorylessness, so $x = x_0 + \sqrt{2Da}\,Z$ with $a \sim \mathrm{Exp}(r)$ and Z standard normal — no time discretisation is involved. **Right:** the stationary probability current $J(x) = -D\,\partial_x p_{ss}$. It is nonzero for every $x \neq x_0$, directed away from the reset point, and jumps by exactly r across it. A nonzero stationary current is the operational definition of broken detailed balance.

5. First Passage Under Restart

We now ask the question that gives resetting its practical bite. Let T denote the first-passage time of the *un-restarted* process to a target, and let T_r denote the first-passage time of the same process under restart at rate r. How do they compare?

5.1 The Renewal Identity

Theorem 2 (Mean first passage under restart)**.** Let $\tilde{T}(s) = \mathbb{E}[e^{-sT}]$ be the Laplace transform of the un-restarted passage time. Then for $r > 0$,

$$\langle T_r \rangle \;=\; \frac{1 - \tilde{T}(r)}{r\,\tilde{T}(r)}. \tag{5.1}$$

Proof. Condition on the first reset epoch $R \sim \mathrm{Exp}(r)$, independent of the diffusion. Either the target is reached first ($T \leq R$), or a reset occurs first, after which the process begins afresh at x_0 and — by the strong Markov property and the memorylessness of R — the remaining passage time is an independent copy of T_r. Hence

$$\langle T_r \rangle \;=\; \mathbb{E}[T\,\mathbf{1}_{\{T \leq R\}}] \;+\; \mathbb{E}[R\,\mathbf{1}_{\{R<T\}}] \;+\; \mathbb{P}(R < T)\,\langle T_r \rangle. \tag{5.2}$$

Using $\mathbb{P}(R > t) = e^{-rt}$ and independence: $\mathbb{P}(R < T) = 1 - \tilde{T}(r)$, $\mathbb{E}[T\mathbf{1}_{\{T\leq R\}}] = \mathbb{E}[Te^{-rT}] = -\tilde{T}'(r)$, and $\mathbb{E}[R\mathbf{1}_{\{R<T\}}] = \frac{1}{r}(1 - \tilde{T}(r)) + \tilde{T}'(r)$. Substituting and solving the resulting linear equation for $\langle T_r \rangle$ gives the claim.

The identity is worth pausing on. It contains **no information about the process beyond** $\tilde{T}$. Whatever the diffusion — drifted, multidimensional, in a potential, on a graph

— the mean passage time under restart is this one function of the Laplace transform of the un-restarted passage time. The mechanism has been fully separated from the process.

5.2 Restart Can Make an Infinite Mean Finite

The most dramatic consequence is visible before any optimisation. Driftless Brownian motion started at distance L from a target reaches it almost surely, but $\mathbb{E}[T] = \infty$ — the passage-time density decays as $t^{-3/2}$, and the mean diverges. Restart cures this. With $\tilde{T}(r) = e^{-L\sqrt{r/D}}$ (Section 7), Theorem 2 gives the finite quantity

$$\langle T_r \rangle \;=\; \frac{e^{\,z} - 1}{r}, \qquad z = L\sqrt{r/D}, \tag{5.3}$$

for every $r > 0$. An infinite expectation has been rendered finite by throwing the process away and starting again. Nothing about the underlying diffusion changed; only the willingness to abandon a bad excursion.

6. The Optimal Restart Rate

6.1 When Does Restart Help?

Restart is not always beneficial: it discards progress as well as failure. The criterion separating the two regimes is exact and depends on one number.

Theorem 3 (CV criterion; Reuveni [4]). Let T have finite mean $\langle T \rangle$ and finite variance, and let $\mathrm{CV}(T) = \sigma(T)/\langle T \rangle$. Then introducing restart at a small rate reduces the mean first-passage time if and only if

$$\mathrm{CV}(T) \;>\; 1. \tag{6.1}$$

Proof. Expand Theorem 2 as $r \downarrow 0$. Writing $\tilde{T}(r) = 1 - \langle T \rangle r + \frac{1}{2}\langle T^2 \rangle r^2 + O(r^3)$ and substituting,

$$\langle T_r \rangle \;=\; \langle T \rangle \;+\; \frac{r}{2}\Big(\langle T^2 \rangle - 2\langle T \rangle^2\Big) \;+\; O(r^2). \tag{6.2}$$

The sign of the first-order coefficient decides. It is negative precisely when $\langle T^2 \rangle > 2\langle T \rangle^2$, i.e. when $\mathrm{Var}(T) > \langle T \rangle^2$, i.e. when $\mathrm{CV}(T) > 1$.

The interpretation is worth stating plainly. $\mathrm{CV} > 1$ means the passage time is **more erratic than an exponential** (for which $\mathrm{CV} = 1$ exactly). Such a distribution has a heavy population of unlucky excursions, and it is these — not the typical ones — that dominate the mean. Restart is a policy for refusing to sit through them. When the passage time is regular ($\mathrm{CV} < 1$), there are no such excursions to escape, and restarting merely destroys accumulated progress.

6.2 The Optimum Is Where the Restarted Process Has CV = 1

The criterion above concerns small r. At the optimal rate a second, sharper characterisation holds.

Theorem 4 (Unit CV at the optimum). Let $r^* > 0$ minimise $\langle T_r \rangle$. Then the restarted passage time at that rate satisfies

$$\mathrm{CV}(T_{r^*}) \;=\; 1. \tag{6.3}$$

Sketch. Differentiating the renewal identity of Theorem 2 and setting $\partial_r \langle T_r \rangle = 0$ gives, after rearrangement, a relation between $\tilde{T}(r^*)$ and $\tilde{T}'(r^*)$; the second moment of T_r, computed from the same renewal decomposition, then satisfies $\langle T_{r^*}^2 \rangle = 2\langle T_{r^*} \rangle^2$, which is exactly $\mathrm{Var}(T_{r^*}) = \langle T_{r^*} \rangle^2$.

The optimally restarted process is, in this precise second-moment sense, **exponential-like**: the restart mechanism has tuned away exactly as much of the fluctuation as it can, and what remains is memoryless in its coefficient of variation. It is a fixed-point statement, and it gives a diagnostic that requires no knowledge of the underlying dynamics: measure the CV of your restarted process; if it is not one, you are not restarting optimally. Section 8 verifies it numerically.

7. The Drifted Brownian Case

Everything above becomes explicit for the case that matters most in applications: Brownian motion with drift $v > 0$ toward a target at distance L, with the same diffusion coefficient D used throughout — that is, generator $\mathcal{L} = v\,\partial_x + D\,\partial_{xx}$, or $dX_t = v\,dt + \sqrt{2D}\,dW_t$.

The un-restarted passage time is inverse Gaussian, with

$$\tilde{T}(s) \;=\; \exp\!\left[\frac{L}{2D}\Big(v - \sqrt{v^2 + 4Ds}\Big)\right], \qquad \langle T \rangle = \frac{L}{v}, \qquad \mathrm{Var}(T) = \frac{2DL}{v^3}. \tag{7.1}$$

Hence the coefficient of variation is

$$\mathrm{CV}(T) \;=\; \sqrt{\frac{2D}{L\,v}}, \tag{7.2}$$

and Theorem 3 acquires a form with no probability in it at all:

$$\boxed{\text{restart helps} \iff 2D > L\,v} \tag{7.3}$$

Diffusion versus ballistic transport: restart is worthwhile exactly when the noise dominates the drift over the distance to be covered. A strongly directed searcher should never restart; a weakly directed one should.

The **driftless limit** is the extreme case. As $v \downarrow 0$ the CV diverges, $\langle T \rangle = \infty$, and restart is unambiguously beneficial. Minimising $\langle T_r \rangle = (e^z - 1)/r$ with $z = L\sqrt{r/D}$ gives the transcendental condition

$$\frac{z^*}{2} \;=\; 1 - e^{-z^*}, \tag{7.4}$$

whose nontrivial root is $z^* = 1.593624\ldots$, universal — independent of L and D. The optimum is therefore

$$r^* \;=\; \frac{(z^*)^2 D}{L^2} \;=\; 2.539638\,\frac{D}{L^2}, \qquad \langle T_{r^*}\rangle \;=\; 1.544139\,\frac{L^2}{D}, \tag{7.5}$$

a finite mean where the un-restarted process has none. (Minimising the expression numerically, without reference to the transcendental equation, returns $z = 1.593624$ — the same root to six decimals.)

8. Numerical Results

8.1 Method

Two exact sampling schemes are used, so that **no time-discretisation bias enters any figure**. Stationary positions are drawn by exploiting memorylessness: the age a of the current excursion is $\mathrm{Exp}(r)$, so $x = x_0 + \sqrt{2Da}\,Z$ with Z standard normal is an exact stationary draw. First-passage times under restart are drawn by exact renewal: between resets the passage time of drifted Brownian motion is inverse Gaussian, and the reset is exponential, so a restart trajectory is a chain of exact draws with no grid at all. This matters — an Euler scheme with a hard target systematically misses crossings between steps and biases every passage time upward.

8.2 The Steady State

Figure 1 (Section 4) shows the stationary density against 400,000 exact draws, and the current alongside it. The cusp, the exponential tails, and the current's jump of exactly r at the reset point are all as Theorem 1 predicts.

8.3 The Optimum

We take $D = 1$, $L = 1$, $v = 0.5$, so $\mathrm{CV}(T) = \sqrt{2D/(Lv)} = 2 > 1$ and restart should help.

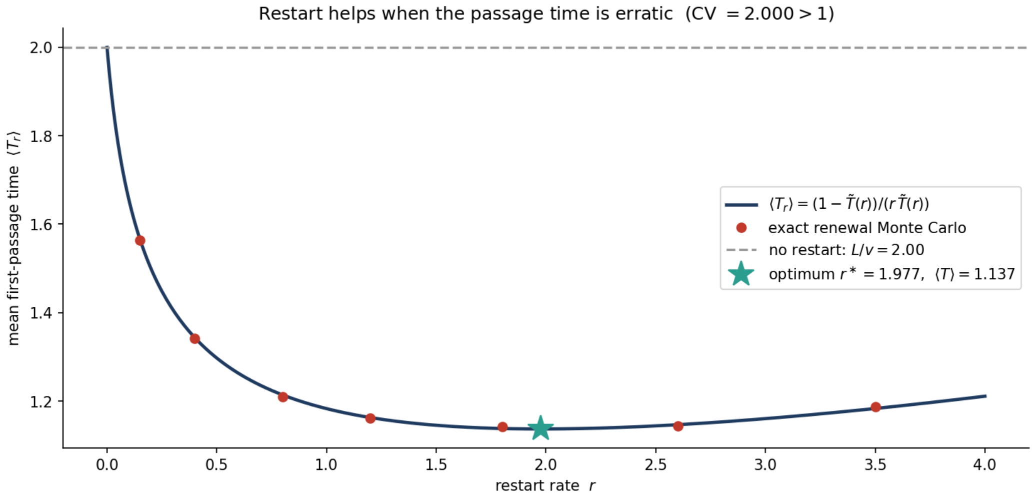

Figure 2: Mean first-passage time versus restart rate for drifted Brownian motion ($D = 1$, $L = 1$, $v = 0.5$, so $\mathrm{CV} = 2 > 1$). The curve is the renewal identity of Theorem 2; the points are exact renewal Monte Carlo (60,000 trajectories each). Restart at $r^* = 1.977$ reduces the mean passage time from $L/v = 2.000$ to 1.137, a speed-up of 1.759×. Restarting too aggressively is worse than not restarting at all: past $r = 11.93$ the curve rises back above the no-restart line, and as $r \to \infty$ it diverges — a process that is restarted constantly never gets anywhere.

The Monte Carlo agrees with the closed form throughout. At the optimum, 300,000 exact renewal trajectories give $\langle T_{r^*} \rangle = 1.1384$ against the predicted 1.1372, and — the point of Theorem 4 — a measured coefficient of variation of 0.9992.

8.4 The Criterion

Sweeping the drift v moves the CV of the un-restarted passage time through 1 and lets us test Theorem 3 directly.

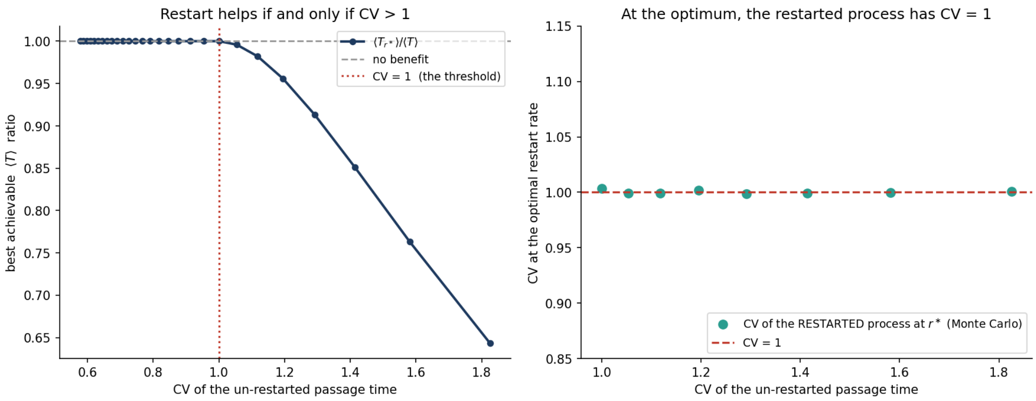

Figure 3: **Left:** the best achievable mean passage time, as a fraction of the un-restarted mean, against the CV of the un-restarted passage time. The curve is flat at 1 (no benefit) for every process with CV < 1 and falls away immediately once CV exceeds 1. The threshold is not fitted; it is Theorem 3. **Right:** the coefficient of variation of the *restarted* process, measured by Monte Carlo at the numerically located optimum r^*, for each case in which restart helps. Every point sits on CV $= 1$, confirming Theorem 4 to within 3.4×10^{-3}.

For the cases with CV < 1 the numerical optimiser drives r^* to order 10^{-7} — that is, to zero, exactly as the criterion demands. For those with CV > 1 it finds a strictly positive optimum, and the restarted process at that optimum has unit CV in every case (mean 1.0003 across the eight cases, largest deviation 3.4×10^{-3}).

9. Conclusion

Resetting is a small modification with large consequences. It makes the generator non-local, destroys self-adjointness, and produces a stationary state that is not an equilibrium but a permanent circulation of probability — outward by diffusion, back by teleportation, with a current that never vanishes.

Its effect on first passage is captured entirely by a renewal identity that involves the underlying process only through the Laplace transform of its passage time, and the decision of whether to restart at all reduces to a single dimensionless comparison: is the passage time more erratic than an exponential? If it is, restart helps, and at the best restart rate the process is driven to unit coefficient of variation — a fixed point that can be checked without knowing the dynamics at all. For drifted Brownian motion the criterion becomes the transparent statement $D > Lv$: restart when noise dominates transport.

Two directions are immediate. The criterion is a statement about the *mean*; the corresponding question for the full distribution of T_r, and for risk-sensitive objectives, is open in general. And the reset point has been treated as fixed — allowing it to be chosen, or randomised, turns the problem into one of control, with the reset distribution as the decision variable.

10. References

1. Evans, M. R., & Majumdar, S. N. (2011). Diffusion with stochastic resetting. *Physical Review Letters*, 106(16), 160601.

2. Evans, M. R., & Majumdar, S. N. (2011). Diffusion with optimal resetting. *Journal of Physics A*, 44(43), 435001.

3. Evans, M. R., Majumdar, S. N., & Schehr, G. (2020). Stochastic resetting and applications. *Journal of Physics A*, 53(19), 193001.

4. Reuveni, S. (2016). Optimal stochastic restart renders fluctuations in first passage times universal. *Physical Review Letters*, 116(17), 170601.

5. Pal, A., & Reuveni, S. (2017). First passage under restart. *Physical Review Letters*, 118(3), 030603.

6. Chechkin, A., & Sokolov, I. M. (2018). Random search with resetting: a unified renewal approach. *Physical Review Letters*, 121(5), 050601.

7. Reuveni, S., Urbakh, M., & Klafter, J. (2014). Role of substrate unbinding in Michaelis–Menten enzymatic reactions. *Proceedings of the National Academy of Sciences*, 111(12), 4391–4396.

8. Luby, M., Sinclair, A., & Zuckerman, D. (1993). Optimal speedup of Las Vegas algorithms. *Information Processing Letters*, 47(4), 173–180.

9. Redner, S. (2001). *A Guide to First-Passage Processes.* Cambridge University Press.

10. Karatzas, I., & Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus* (2nd ed.). Springer.

11. Chhikara, R. S., & Folks, J. L. (1989). *The Inverse Gaussian Distribution.* Marcel Dekker.

12. Zwanzig, R. (2001). *Nonequilibrium Statistical Mechanics.* Oxford University Press.