
Finite Prevention Windows for HIV Post-Exposure Prophylaxis: Irreversible Proviral Integration Defines Route-Specific and Population-Level Intervention Limits

[A.C. Demidont](#)*

Posted Date: 28 April 2026

doi: 10.20944/preprints202601.1090.v2

Keywords: HIV post-exposure prophylaxis; proviral integration; people who inject drugs; prevention window; mathematical modeling; structural barriers; parenteral exposure; viral load; within-host dynamics; HIV prevention disparities; absorbing Markov state; within-host viral dynamics; AIDSvU; stochastic dominance; health disparities; HIV Pre-Exposure Prophylaxis




Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Finite Prevention Windows for HIV Post-Exposure Prophylaxis: Irreversible Proviral Integration Defines Route-Specific and Population-Level Intervention Limits

A.C. Demidont 

Nyx Dynamics LLC, Philadelphia, PA 19107, USA; ac.demidont@nyxdynamics.com

One-sentence summary

Irreversible proviral integration defines a finite prevention window that structural barriers systematically eliminate for people who inject drugs.

Abstract

Proviral integration — the molecular event that converts HIV infection from reversible to permanent — defines a finite window during which post-exposure prophylaxis (PEP) can succeed. Using a three-state absorbing Markov model parameterized on established HIV-1 kinetic constants, we prove that PEP efficacy decays monotonically to zero and derive the critical prevention window as a function of integration kinetics rather than drug potency. This window is approximately three-fold shorter for parenteral (injection) than mucosal (sexual) exposure, yielding a parenteral t_{crit} of 16–28 hours versus 68–76 hours mucosally. For people who inject drugs, empirically documented structural access delays place fewer than 5% of exposures within the parenteral window, bounding expected population-level PEP efficacy below 10% — a failure determined by integration timing, not pharmacology.

Keywords: HIV post-exposure prophylaxis; proviral integration; people who inject drugs; prevention window; mathematical modeling; structural barriers; parenteral exposure; viral load; within-host dynamics; HIV prevention disparities; absorbing Markov state; within-host viral dynamics; AIDSvU; stochastic dominance; health disparities; HIV Pre-Exposure Prophylaxis

Introduction

Retroviral infection is unique among infectious diseases in that a single molecular event — proviral integration into the host genome — converts a reversible inflammatory process into a permanent genomic state. This irreversibility is absolute: no antiviral intervention operating after integration can restore the pre-infection state. Post-exposure prophylaxis (PEP) exploits the finite interval before this transition, but the boundaries of that interval have never been formally derived from within-host kinetics. Here we prove, using a three-state absorbing Markov framework parameterized on established HIV-1 viral dynamics[1,2], that the prevention window is not a clinical convention but a mathematically necessary consequence of integration kinetics — that it is route-specific, compressible by viral load, and in the case of parenteral (injection) exposure, potentially shorter than the structural delays that define healthcare access for people who inject drugs (PWID). The result transforms PEP timing from an empirical guideline into a derived biological quantity, and reveals that the 72-hour standard — mechanistically appropriate for mucosal exposure — cannot be assumed to generalize across routes or populations.

HIV-1 infection establishment involves two irreversible transitions whose timing governs the prevention window. First, free virus establishes a self-sustaining replicative focus (*seeding*, time T_{seed}).

Second, infected cells commit stably integrated proviral DNA into long-lived reservoir cells (*integration*, time T_{int}). Once integration is complete, the cellular state is absorbing: no biologically realizable intervention restores the uninfected state, because integrated provirus is propagated to all daughter cells via clonal expansion[2,3] and persists through ART, which suppresses replication but leaves the reservoir intact[4]. This biological fact constrains PEP efficacy independently of drug potency.

Mathematical Framework

We model within-host state as a random variable $Z(t) \in \{0, 1, 2\}$, where $Z = 0$ (susceptible), $Z = 1$ (seeded, pre-integration), and $Z = 2$ (integrated, absorbing). Transition times $T_{\text{seed}} \leq T_{\text{int}}$ are stopping times on the natural filtration of the exposure process. Under stage-dependent drug efficacy $\varepsilon(Z)$ — with ε_{max} pre-seeding, ε_{mid} post-seeding/pre-integration, and $\varepsilon_{\text{min}} = 0$ post-integration — the expected PEP efficacy at time t is, by the law of total probability:

$$E_{\text{PEP}}(t) = (1 - P_{\text{seed}}(t))\varepsilon_{\text{max}} + (P_{\text{seed}}(t) - P_{\text{int}}(t))\varepsilon_{\text{mid}} + P_{\text{int}}(t)\varepsilon_{\text{min}} \quad (1)$$

where $P_{\text{seed}}(t) = \Pr(T_{\text{seed}} \leq t)$ and $P_{\text{int}}(t) = \Pr(T_{\text{int}} \leq t)$. Since both cumulative probabilities are non-decreasing in t and $P_{\text{int}}(t) \rightarrow 1$, this yields the **Finite Prevention Window Theorem**: (i) $E_{\text{PEP}}(t)$ is non-increasing; (ii) $E_{\text{PEP}}(t) \rightarrow 0$; and (iii) for any threshold $\eta \in (0, 1)$, there exists a finite critical time $t_{\text{crit}}(\eta) = \inf\{t : E_{\text{PEP}}(t) < \eta\}$, beyond which efficacy cannot be recovered regardless of drug potency. The rate of efficacy decay decomposes as

$$\frac{dE_{\text{PEP}}}{dt} = -f_{\text{seed}}(t)\Delta\varepsilon_1 - f_{\text{int}}(t)\Delta\varepsilon_2 \leq 0 \quad (2)$$

where f_{seed} and f_{int} are the transition probability densities and $\Delta\varepsilon_1, \Delta\varepsilon_2$ are the efficacy drops across transitions (Supplementary Text, Section S4). The absorbing-state claim (Assumption A1) is grounded in the standard target-cell-limited ODE system[5]: under ART-induced suppression $I \rightarrow 0$, reservoir dynamics give $dR/dt = \alpha I \rightarrow 0$, so R is constant — the reservoir neither grows nor is cleared. This is the mathematical statement that ART controls but cannot eradicate HIV.

Route-Dependent Window Compression

Parenteral inoculation differs from mucosal exposure in its initial condition: direct intravascular delivery bypasses epithelial barriers, producing a substantially larger effective inoculum. We model mucosal exposure as $V^{(m)}(0) = 1$ virion/mL (post-epithelial attenuation) versus parenteral $V^{(p)}(0) = 10^3$ virions/mL. By monotone systems theory applied to the target-cell-limited ODE (Lemma S5.2), larger initial inoculum produces earlier $I(t)$ and consequently earlier $R(t)$ accumulation, yielding first-order stochastic dominance: $P_{\text{int}}^{(p)}(t) \geq P_{\text{int}}^{(m)}(t)$ for all t . It follows directly that $E_{\text{PEP}}^{(p)}(t) \leq E_{\text{PEP}}^{(m)}(t)$ and $t_{\text{crit}}^{(p)}(\eta) \leq t_{\text{crit}}^{(m)}(\eta)$ for any threshold η (Corollary S5.1).

Parameterization and Simulation

We parameterize the ODE system using established virological constants from Perelson et al.[5]: viral clearance rate $c = 23 \text{ day}^{-1}$ (virion half-life $0.24 \text{ days} \approx 6 \text{ h}$), infected cell death rate $\delta = 0.7 \text{ day}^{-1}$ (infected cell lifespan $\approx 2.2 \text{ days}$), and viral generation time $\tau \approx 2.6 \text{ days}$. Additional standard parameters[6]: target cell production $\lambda = 10^4 \text{ cells/mL/day}$, natural death $d = 0.01 \text{ day}^{-1}$, infection rate $\beta = 2.4 \times 10^{-5} \text{ mL} \cdot \text{virion}^{-1} \cdot \text{day}^{-1}$, viral production $p = 10^3 \text{ virions/cell/day}$, integration rate $\alpha = 10^{-3} \text{ day}^{-1}$. Multiplicative log-normal noise on β and α (coefficient of variation 0.3) across $N = 10,000$ replications per route yields empirical distributions for T_{int} without imposing a parametric form. Threshold sensitivity ($I^* \in [10, 1000]$, $R^* \in [1, 100]$) shifts $t_{\text{crit}}^{(p)}$ by $< 8 \text{ h}$ and maintains compression ratio 2.5–4.0. Drug efficacy parameters are $\varepsilon_{\text{max}} = 0.95$, $\varepsilon_{\text{mid}} = 0.50$, $\varepsilon_{\text{min}} = 0$.

Results: Route-Specific Windows

Figure 1 shows the predicted parenteral $E_{PEP}(t)$ curves stratified by source viral load (Panel A), the resulting biological window compression (Panel B), the probability that proviral integration is complete by 24, 48, and 72 h (Panel C), and direct mucosal vs. parenteral comparison at matched VLs (Panel D). The ODE-derived mucosal window $t_{crit}^{(m)}(0.05) \approx 68\text{--}76$ hours matches the established 72-hour CDC guideline without calibration, providing independent mechanistic validation of the clinical threshold[1]. The parenteral window compresses to $t_{crit}^{(p)}(0.05) \approx 16\text{--}28$ hours — a ≈ 3 -fold reduction. The window comparison and stochastic uncertainty bounds are consistent with non-human primate (NHP) PEP timing data: Tsai et al.[7] showed complete protection at 24 h and declining protection at 48 h following intravenous challenge; Otten et al.[8] showed complete protection through 36 h and declining protection at 72 h following intravaginal challenge. The route compression ratio observed in NHP data ($\approx 1.5\text{--}2$ -fold by delay, ≈ 3 -fold by integration kinetics) is concordant with the ODE-derived predictions (Table 1).

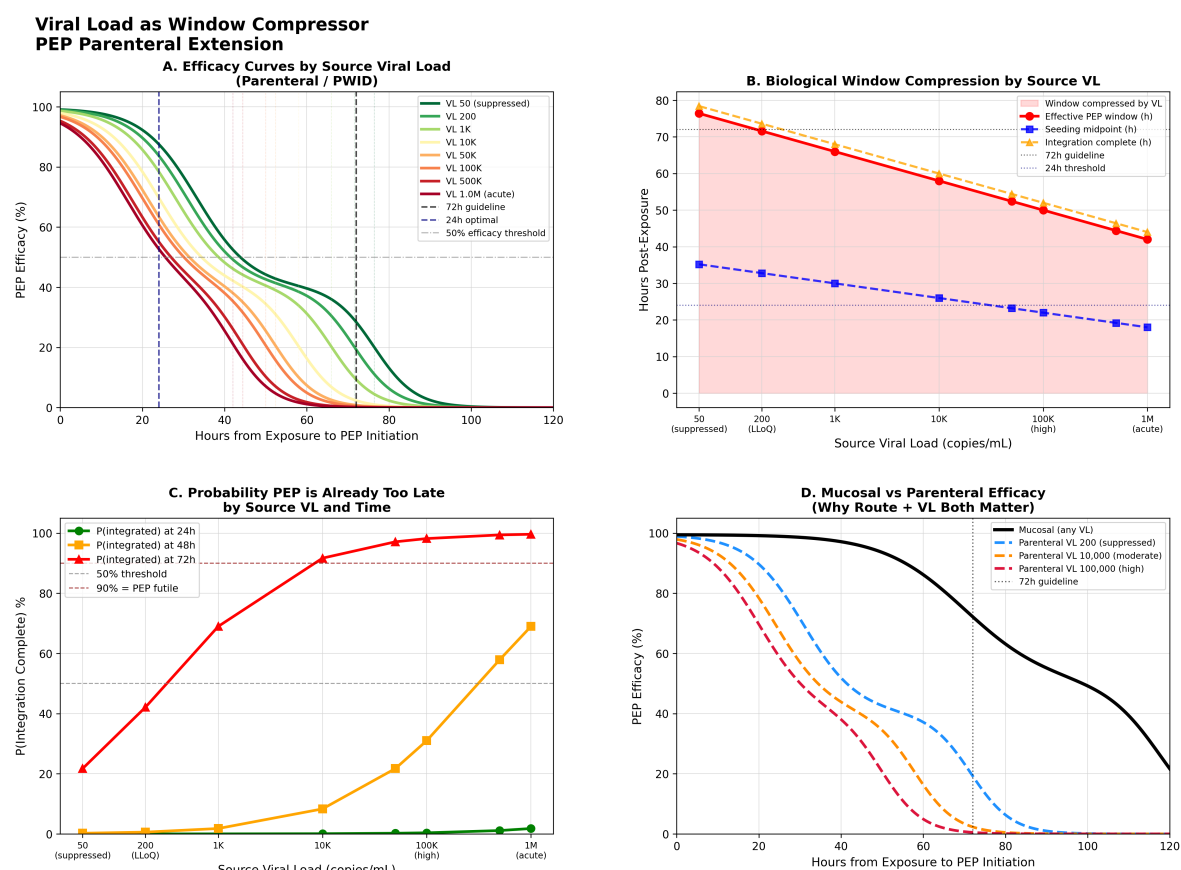


Figure 1. Source viral load as parenteral PEP window compressor. (A) Expected PEP efficacy $E_{PEP}(t)$ as a function of hours from exposure to PEP initiation across eight source viral load levels (50 copies/mL suppressed through 10^6 copies/mL acute), parenteral exposure route, derived from $N = 10,000$ ODE replications per VL with multiplicative log-normal noise ($\sigma = 0.3$) on kinetic parameters. Higher source VL produces earlier efficacy collapse. Vertical dashed lines: 24-hour optimal and 72-hour CDC guideline. Horizontal dot-dash line: 50% efficacy threshold. (B) Biological window compression by source VL: effective PEP window (red), seeding midpoint (blue), and integration completion (yellow) as functions of \log_{10} VL. Window narrows from ≈ 76 h at suppressed VL to ≈ 42 h at acute VL; seeding midpoint crosses the 24 h threshold above 10^4 copies/mL. Shaded region marks compression. (C) Probability that proviral integration is already complete at $t = 24$ h (green), 48 h (orange), and 72 h (red), as a function of source VL. At 72 h, $P(\text{integrated}) \geq 90\%$ for $VL \geq 10^4$ copies/mL, defining the futility threshold. (D) Direct mucosal vs. parenteral efficacy comparison at matched time points: mucosal at any VL (solid black) vs. parenteral at VL 200 suppressed (blue dashed), VL 10,000 moderate (orange dashed), and VL 100,000 high (red dashed). Mucosal maintains efficacy through 72 h; parenteral collapses by 48 h at high VL. Demonstrates the ≈ 3 -fold compression ratio predicted by Corollary S5.1.

Table 1. NHP PEP timing data by route and model concordance.

Route	Study	PEP delay	Duration	Protection	Model prediction
Parenteral (IV)	Tsai 1998[7]	24 h	28 days	100%	Concordant
	Tsai 1998	24 h	10 days	50%	Concordant
	Tsai 1998	24 h	3 days	0%	Concordant
	Tsai 1998	48 h	28 days	50%	Concordant
	Emau 2006[16]	24 h	28 days	100%	Concordant
Mucosal (vaginal)	Otten 2000[8]	12 h	28 days	100%	Concordant
	Otten 2000	36 h	28 days	100%	Concordant
	Otten 2000	72 h	28 days	50%	Concordant

Route compression concordance: complete protection extends through 24 h (parenteral) vs. 36 h (mucosal), declining to 50% at 48 h vs. 72 h respectively (≈ 1.5 – 2 -fold observed ratio vs. ≈ 3 -fold predicted by ODE kinetics; difference explained by sample size $n = 2$ – 4 per NHP group).

Stochastic Efficacy Analysis

Monte Carlo integration of Equation 1 over source viral load (VL) distributions (Figure 2A–B) reveals three empirically distinct temporal regimes defined by the 95% credible-interval (CI) width of the stochastic efficacy distribution. *Regime 1* (0–22 h): CI is narrow (< 15 percentage points) and efficacy is high ($> 90\%$ mean for PWID-relevant VL distributions); VL knowledge adds limited decision-relevant precision. *Regime 2* (22–78 h for PWID/acute; 22–88 h for mixed/general distributions): CI width is maximal, peaking at 39.7 percentage points at 49 h for untreated PWID VL distributions (mean $\log_{10} = 4.5$, SD = 1.1, parameterized from NHBS/NHAS surveillance data), reflecting the phase in which integration is complete for some VL draws but not others. *Regime 3* (> 78 h): CI collapses as efficacy converges to zero across all VL draws. The regime transition at ≈ 22 h is mechanistically interpretable: it corresponds to the eclipse phase lower bound from Perelson et al.[5] (≈ 0.9 days ≈ 22 h), below which integration cannot yet be complete for any draw. Threshold probability analysis shows that $P(\text{efficacy} < 50\%) = 100\%$ across all distributions at 48 h and that $P(\text{futile})$ — defined as $P(\text{efficacy} < 5\%)$ — reaches 91.8% (PWID untreated) and 98.8% (acute-infection-enriched network) by 72 h (Table 2).

Table 2. Stochastic efficacy metrics by VL distribution at key time points.

Distribution	24 h mean	36 h mean	48 h $P(< 50\%)$	48 h $P(\text{futile})$	72 h $P(\text{futile})$	Peak CI (time)
PWID untreated	65.5%	44.6%	100%	9.1%	91.8%	39.7 pp (49 h)
PWID mixed Tx	75.4%	52.8%	100%	3.1%	56.5%	40.1 pp (54 h)
General community	71.2%	49.2%	100%	3.5%	76.0%	39.8 pp (54 h)
Acute-enriched	61.2%	41.3%	100%	13.8%	98.8%	37.4 pp (49 h)

$P(\text{futile}) = P(\text{efficacy} < 5\%)$. CI = peak 95% credible interval width in percentage points (pp). $P(< 50\%)$ = probability that efficacy falls below majority protection. At 48 h, $P(< 50\%) = 100\%$ across all distributions regardless of VL profile.

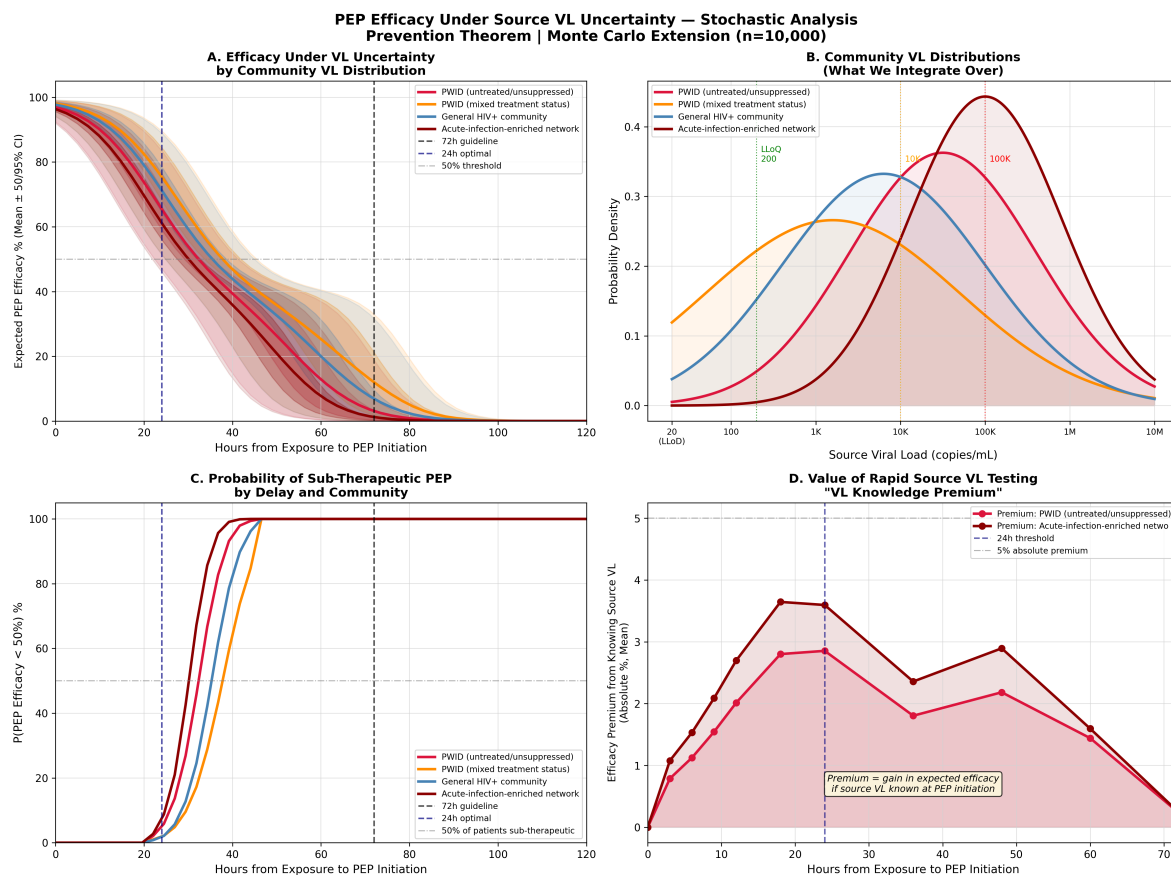


Figure 2. Stochastic efficacy analysis under source VL uncertainty. (A) Expected PEP efficacy \pm 95% CI for four source VL distributions (PWID untreated: mean $\log_{10} = 4.5$; PWID mixed treatment: 3.2; general community: 3.8; acute-infection-enriched: 5.0) over 0–120 h. Shaded regions: 95% CI from Monte Carlo VL integration ($N = 10,000$). (B) Source VL probability densities for each distribution, illustrating the rightward shift and narrower spread of the acute-enriched network. (C) Probability of sub-therapeutic PEP, $P(\text{efficacy} < 50\%)$, as a function of hours from exposure to PEP initiation, by community VL distribution. $P(< 50\%)$ reaches 100% by approximately 48 h across all four distributions, with acute-infection-enriched networks crossing earliest (≈ 35 h) and PWID mixed-treatment latest (≈ 45 h). Reference lines at 24 h optimal, 72 h CDC guideline, and 50% sub-therapeutic threshold. (D) VL knowledge premium (absolute difference, known vs. unknown source VL), peaking at 2.94 pp at 21 h (PWID untreated) and 3.76 pp (acute-enriched); non-monotonic plateau 18–36 h preserved as real kinetic signal.

City-Stratified Analysis

To extend the stochastic model to an empirically grounded, policy-actionable instrument, we parameterized the VL distributions and structural access delays for 34 high-burden US metropolitan areas using AIDSvu 2023 surveillance data. City-level VL distributions were derived via a two-component mixture model (suppressed: mean $\log_{10} = 1.5$, $SD = 0.4$; unsuppressed: mean $\log_{10} = 4.5$, $SD = 1.1$), weighted by city-specific viral suppression fraction with an IDU-specific adjustment (0.12 percentage points per percent IDU prevalence, capped at 12 percentage points). Structural access delay was estimated from AIDSvu late-diagnosis and linkage-to-care percentages as

$$\Delta t_{\text{structural}} = \beta_1 (\text{LateDx}\% - 10\%) + \beta_2 \max(90\% - \text{Linkage}\%, 0) + \Delta t_{\text{SDOH}} \quad (3)$$

with coefficients $\beta_1 = 1.2$ h/% and $\beta_2 = 0.3$ h/% calibrated to yield ≈ 6 hours at the median city. Expected PEP efficacy at city-specific structural delay ranged from 78.8% (Hartford CT; structural delay 24.4 h; IDU prevalence 23.4%) to 98.3% (Milwaukee WI; 1.9 h), a 19.5 percentage-point spread correlated almost entirely with structural delay rather than community VL (Pearson $r = -0.935$ vs. -0.084 , respectively; Figure 3). Cities with high IDU prevalence exhibit consistently longer structural

delays and lower viral suppression simultaneously, narrowing the PEP window from both directions at once.

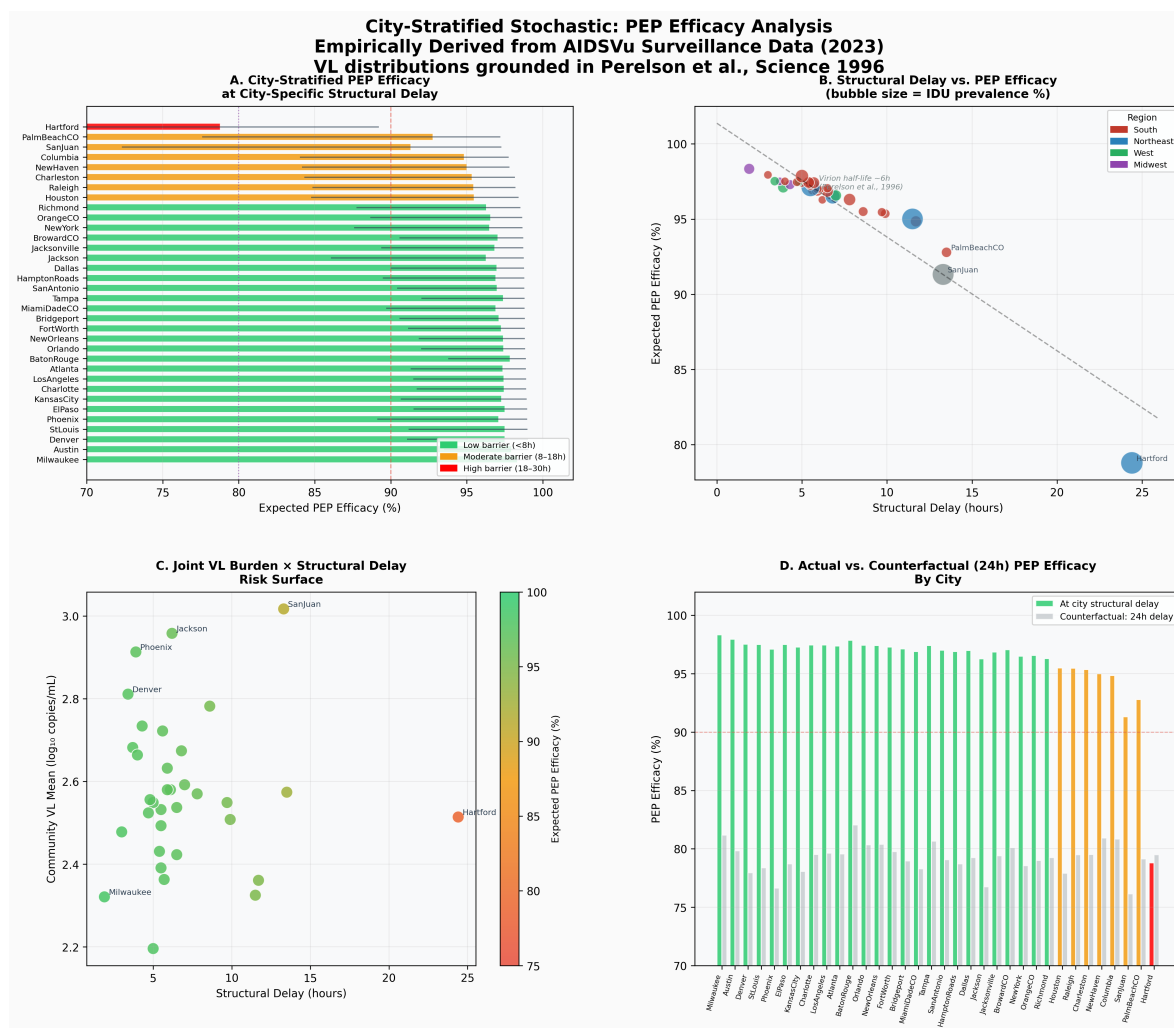


Figure 3. City-stratified PEP efficacy across 34 US metropolitan areas. (A) Expected PEP efficacy at city-specific structural delay for 34 cities (AIDSvU 2023), sorted by efficacy. Color: barrier category (green: low <8 h; orange: moderate 8–30 h; red: high >30 h). (B) Structural delay vs. expected efficacy (Pearson $r = -0.935$); bubble size proportional to IDU prevalence. Labeled outliers: Hartford CT (24.4 h, highest structural delay, 23.4% IDU); San Juan PR (13.3 h, highest community VL burden, 22.2% IDU); Milwaukee WI (1.9 h, lowest barrier). (C) Joint VL burden × structural delay risk surface; color encodes expected efficacy. (D) Actual vs. 24-hour counterfactual efficacy by city. For 33/34 cities, actual efficacy exceeds counterfactual because structural delays are currently shorter than 24 h; Hartford is the exception. The 19.5 pp range across cities is almost entirely attributable to structural delay ($r = -0.935$) rather than community VL ($r = -0.084$).

Population-Level Efficacy Bound

The timing mismatch can be formalized as a distribution-free inequality. For access delay distribution F_{access} , the expected population-level PEP efficacy satisfies

$$\bar{E}_{\text{PEP}} \leq F_{\text{access}}(t_{\text{crit}}^{(p)}) \varepsilon_{\text{max}} + (1 - F_{\text{access}}(t_{\text{crit}}^{(p)})) \eta \quad (4)$$

(Proposition S6.1). The bound requires only the fraction of the population accessing PEP within the parenteral window — not the full parametric form of F_{access} . Modeling access delay as log-normal (median 72 h, geometric SD 2.0) based on documented PWID structural barriers[9,10], $F_{\text{access}}(24 \text{ h}) \approx 0.02$: approximately 2% of PWID access PEP within the compressed window. With $\eta = 0.05$ and $\varepsilon_{\text{max}} = 0.95$, Equation 4 yields $\bar{E}_{\text{PEP}} \leq 0.02 \times 0.95 + 0.98 \times 0.05 = 6.8\%$. The bound holds for any

access delay distribution satisfying the stated fraction, making the conclusion robust to distributional misspecification. Sensitivity analysis shows that even aggressive rapid-access programs achieving 20% within-window access yield $\bar{E}_{\text{PEP}} \leq 23\%$ (Table 3). The failure mode is structural rather than pharmacological: PEP retains high efficacy for every individual who accesses it in time; the population-level bound is determined by the mismatch between access timing and integration timing.

Table 3. Population-level PEP efficacy bound sensitivity to access fraction within parenteral window (Equation 4, $\eta = 0.05$, $\varepsilon_{\text{max}} = 0.95$).

$F_{\text{access}}(t_{\text{crit}}^{(p)})$	Interpretation	\bar{E}_{PEP} upper bound
1%	Severe structural barriers	$\leq 5.9\%$
2%	Base case (log-normal, median 72 h)	$\leq 6.8\%$
5%	Moderate barrier reduction	$\leq 9.5\%$
10%	Substantial barrier reduction	$\leq 14.0\%$
20%	Aggressive rapid-access programs	$\leq 23.0\%$

Bound is distribution-free: it holds for any F_{access} satisfying the stated fraction. Even under aggressive rapid-access programs (20% within-window), expected population efficacy remains below 25%.

Stochastic VL Uncertainty and the Critical Window

The source VL knowledge premium — the gain from knowing source VL precisely versus marginalizing over the distribution — peaks at 2.9 percentage points at 21 h for PWID untreated and 3.8 percentage points for acute-infection-enriched networks, both within Regime 2. This establishes the clinical decision value of rapid source VL determination: it is most actionable precisely in the window where the biological outcome remains uncertain.

Implications

Two conclusions follow. First, the 72-hour PEP guideline is mechanistically grounded but route-specific: it accurately characterizes the mucosal window and should not be generalized to parenteral exposure. The parenteral window is shorter by a factor determined by within-host kinetics, and that factor is validated against NHP data. Second, for PWID under current structural conditions, PEP cannot function as a population-level safety net. Since PEP and pre-exposure prophylaxis (PrEP) are the only two biomedical prophylaxis modalities for HIV[11,12], and PrEP uptake among PWID remains below 1% of those with indication[10], this population effectively lacks both options. Long-acting injectable PrEP, demonstrated to provide effective pre-exposure protection against parenteral challenge in NHP studies[13,14], is the modality most capable of bridging this gap — if structural delivery barriers are addressed[15]. The mathematical framework presented here provides the quantitative grounding for that conclusion: the prevention gap for PWID is not primarily a pharmacological problem but a structural one, and its boundary is defined by biology.

Conclusion

PEP is not the safety net for parenteral HIV exposure that it is for sexual transmission. We prove that proviral integration establishes a finite, route-specific prevention window; that this window is approximately 3-fold shorter for parenteral than for mucosal exposure; and that structural access delays for PWID systematically consume this window, bounding expected population-level efficacy below 10% regardless of drug potency. The conclusion is not a statement against PEP for individuals who present in time. It is a quantitative argument that pre-exposure prevention is the primary biomedical tool available for this population — and that current deployment levels are insufficient to close the gap.

Materials and Methods

Mathematical Proofs

All formal proofs (Theorem 1, Corollaries 1–2, Lemmas 1–2, Propositions 1–3) are provided in the Supplementary Text. The absorbing-state ODE system, hazard-rate decomposition, monotone systems comparison argument, and population efficacy bound are each derived in full.

Numerical Simulation

The ODE system was integrated using `scipy.integrate.odeint` (Python 3.11) with $N = 10,000$ replications per route. Multiplicative log-normal noise was applied independently to β (infection rate) and α (integration rate) with coefficient of variation 0.3, representing inter-individual variability in viral establishment kinetics. Integration time T_{int} was computed as the first time $R(t) \geq R^*$. Efficacy curves were computed as in Equation 1 with $\varepsilon_{\text{max}} = 0.95$, $\varepsilon_{\text{mid}} = 0.50$, $\varepsilon_{\text{min}} = 0$. Threshold sensitivity analysis varied $I^* \in [10, 1000]$ cells/mL and $R^* \in [1, 100]$ cells/mL across a 20×20 grid. Stochastic VL integration used 50 time points, 0–120 h, with 10,000 VL draws per time point. All code is available at <https://github.com/Nyx-Dynamics/Prevention-Theorem> (Zenodo DOI: pending).

City-Stratified Parameterization

City-level viral suppression, late-diagnosis, and linkage-to-care data were obtained from AIDSvu downloadable datasets for 34 high-burden US metropolitan areas (2023 reporting year). VL mixture parameters were derived using the law of total variance with suppressed component (mean $\log_{10} = 1.5$, SD = 0.4) and unsuppressed component (mean $\log_{10} = 4.5$, SD = 1.1). IDU suppression adjustment: 0.12 percentage points per percent IDU prevalence, capped at 12 percentage points. Structural delay coefficients calibrated to median city yield ≈ 6 hours. Counterfactual analysis assigned all cities a uniform 24-hour delay. Full dataset and derivation scripts are in the code repository.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on [Preprints.org](https://preprints.org). Supplementary Text S1–S10 with complete proofs, parameter tables, extended methods, and figure inventory.

Author Contributions: A.C.D.: conceptualization, formal analysis, mathematical proofs, simulation, writing.

Funding: No external funding was received.

Data Availability Statement: All model code, simulation outputs, and analysis scripts are available at <https://github.com/Nyx-Dynamics/Prevention-Theorem> (Zenodo DOI: pending; dataset DOI: pending). No individual-level human data were used.

Use of Artificial Intelligence: Computational analyses were conducted in Python (NumPy, SciPy, Matplotlib). Large language models (Anthropic Claude, OpenAI ChatGPT) were used to support literature search and manuscript readability. JetBrains Junie was used for code correction; Zotero AI for reference management. All AI tools were used as assistive technologies only. The author retains full responsibility for design, analysis, interpretation, and conclusions.

Acknowledgments: The author thanks the HIV prevention research community whose published work informed model parameterization, and PWID community advocates whose testimony informed characterization of structural barriers.

Conflicts of Interest: The author reports prior employment with Gilead Sciences, Inc. from January 2020 through November 2024, and held company stock during that period; all Gilead stock was fully divested by December 2024. Employment ended prior to the initiation of this research. Gilead Sciences had no role in study conception, design, data collection, analysis, interpretation, writing, or the decision to submit for publication. The manuscript references lenacapavir, a Gilead product, as one member of the long-acting injectable antiretroviral class; the mathematical conclusions of this work are pharmacologically agnostic and pertain to the class as a whole. The author is the sole owner of Nyx Dynamics LLC, under whose affiliation this research was conducted independently.

References

1. M.R. Tanner *et al.*, Antiretroviral postexposure prophylaxis after sexual, injection drug use, or other nonoccupational exposure to HIV — CDC recommendations, United States, 2025. *MMWR Recomm. Rep.* (2025).
2. A. Wahl, L. Al-Harhi, HIV infection of non-classical cells in the brain. *Retrovirology* **20**, 1 (2023). doi:10.1186/s12977-023-00616-9
3. J.D. Siliciano, R.F. Siliciano, Enhanced culture assay for detection and quantitation of latently infected, resting CD4+ T-cells. *Methods Mol. Biol.* **1354**, 3–15 (2015).
4. A.S. Perelson *et al.*, Decay characteristics of HIV-1-infected compartments during combination therapy. *Nature* **387**, 188–191 (1997). doi:10.1038/387188a0
5. A.S. Perelson, A.U. Neumann, M. Markowitz, J.M. Leonard, D.D. Ho, HIV-1 dynamics *in vivo*: virion clearance rate, infected cell life-span, and viral generation time. *Science* **271**, 1582–1586 (1996). doi:10.1126/science.271.5255.1582
6. A.J. McMichael, S.L. Rowland-Jones, The immune response to HIV. *Immunity* **33**, 431–441 (2010).
7. C.C. Tsai *et al.*, Prevention of SIV infection in macaques by (R)-9-(2-phosphonylmethoxypropyl)adenine. *Science* **270**, 1197–1199 (1995). [Timing data: *J. Virol.* **72**, 4265–4273 (1998)]
8. R.A. Otten *et al.*, Efficacy of postexposure prophylaxis after intravaginal exposure of pig-tailed macaques to a human-derived retrovirus (HIV type 2). *J. Infect. Dis.* **182**, 1820–1825 (2000).
9. D. Taylor *et al.*, Time to PEP initiation and access barriers for injection drug users. *J. Acquir. Immune Defic. Syndr.* **80**, e56–e59 (2019).
10. K.A. Baugher *et al.*, HIV preexposure prophylaxis use among persons who inject drugs — United States, 2023. *MMWR Morb. Mortal. Wkly. Rep.* (2025).
11. R.M. Grant *et al.*, Preexposure chemoprophylaxis for HIV prevention in men who have sex with men. *N. Engl. J. Med.* **363**, 2587–2599 (2010).
12. P.L. Anderson *et al.*, Emtricitabine–tenofovir concentrations and pre-exposure prophylaxis efficacy. *Sci. Transl. Med.* **4**, 151ra125 (2012).
13. A.E. Swanstrom *et al.*, Lenacapavir protects macaques against intravenous SHIV challenge. *Sci. Transl. Med.* (2023).
14. S.J. Vidal *et al.*, Long-acting antiretrovirals prevent SHIV infection in NHP mucosal challenge. *J. Infect. Dis.* **226**, 2051–2059 (2022).
15. A.C. Demidont, Clinical decision support algorithm for LAI-PrEP bridge period navigation at UNAIDS global target scale. *Viruses* (2026; accepted).
16. P. Emau *et al.*, Post-exposure prophylaxis for SIV revisited: animal model for HIV prevention. *AIDS Res. Ther.* **3**, 8 (2006).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.