Carbon-14 Bomb-Pulse Modeling for Tissue Age Analysis: Independent Study Report

Zhaodong Liu

May 2026

Abstract

This independent study presents a computational approach for inferring tissue/cell age characteristics using Carbon-14 bomb-pulse data. Three mathematical frameworks are developed and compared: a deterministic PDE model, a stochastic agent-based ODE simulation (SODE), and a stochastic cohort-based simulation (SPDE). Each framework is evaluated using two parameter estimation strategies—Nelder-Mead optimization and 2D grid search—for a total of six experimental configurations. The study compares predictive accuracy for Carbon-14 concentrations in arterial tissue samples across three tissue types: ACI, CA, and AVM. This report details the methodology, implementation, and comparative analysis of all six configurations, documenting both the mathematical foundations of cell renewal dynamics and the practical trade-offs between deterministic and stochastic modeling.

1 Introduction

The Carbon-14 bomb-pulse method provides a unique opportunity to study cell turnover and tissue dynamics by leveraging the known atmospheric Carbon-14 concentration changes following nuclear weapon testing. This method enables researchers to estimate the age of cells and tissues in human arterial samples, offering insights into biological processes such as cellular senescence and tissue regeneration [1, 2].

Between 1955 and 1963, above-ground nuclear weapons tests released large quantities of neutrons into the atmosphere, nearly doubling the global concentration of ^{14}C. After the Partial Nuclear Test Ban Treaty of 1963, atmospheric ^{14}C levels declined exponentially as the excess carbon was absorbed into the oceans and biosphere. Because all living organisms incorporate atmospheric carbon into their DNA during cell division, this transient pulse functions as a time-stamp: cells born during the peak carry elevated ^{14}C signatures, while cells born later reflect the declining curve. By measuring the ^{14}C content of a tissue sample and comparing it to the known atmospheric curve, researchers can infer when the DNA in those cells was synthesized—and, by extension, how rapidly cells in a given tissue are turning over. The ability to distinguish continuously regenerating tissues from essentially static ones bears directly on cardiovascular disease, cancer biology, and the aging process.

1.1 Research Objective

The primary objective of this study is to develop and compare computational methods for parameter fitting in Carbon-14 bomb-pulse modeling, specifically:

- Developing three mathematical models for tissue dynamics (deterministic PDE, SODE, SPDE)
- Implementing Nelder-Mead optimization for parameter fitting
- A 2D grid search over (γ, b) as an alternative fitting strategy
- Assessing predictive accuracy and computational trade-offs across the six resulting configurations

2 Mathematical Foundation and Evolution

This research develops three Carbon-14 bomb-pulse model formulations—one deterministic and two stochastic—and compares their behavior on arterial tissue data.

2.1 Mathematical Modeling Framework

The mathematical foundation spans deterministic and stochastic formulations:

2.1.1 Transport Equation Analysis

The mathematical analysis of the transport equation that governs cell renewal processes represents a fundamental component of this research [3]. The fundamental partial differential equation describing cell evolution is:

$$\frac{\partial n}{\partial t} + \frac{\partial n}{\partial a} + \gamma n = 0 \tag{1}$$

2.1.2 Method of Characteristics

The method of characteristics for cell renewal models provides a systematic approach to reducing the PDE to an ODE and obtaining analytical expressions for cell density evolution:

$$n(t, a) = n(t - a, 0) \exp\left(-\int_0^a \gamma(s')ds'\right) \tag{2}$$

Under constant rate assumptions, this simplifies to the exponential decay model:

$$n(t,a) = bN_0 e^{-\gamma a} \tag{3}$$

2.1.3 Carbon-14 Integration

The derivation of the Carbon-14 model connects theoretical predictions with experimental measurements [4]. For an individual born at calendar time t_b, the predicted average ^{14}C concentration at age t under constant but independent birth rate b and death rate γ is:

$$C(t) = \frac{K(t_b + t_0)\, e^{-\gamma(t-t_0)} + b\displaystyle\int_0^{t-t_0} K(t + t_b - a)\, e^{-\gamma a}\, da}{\left[1 - \dfrac{b}{\gamma}\right] e^{-\gamma(t-t_0)} + \dfrac{b}{\gamma}} \tag{4}$$

where $K(\cdot)$ represents the atmospheric Carbon-14 concentration as a function of calendar year. The numerator sums two contributions: the ^{14}C carried by surviving cells from the initial cohort (first term) and the ^{14}C from all cells born after t_0 and still alive at age t (integral term). The denominator normalizes by the total surviving cell population, which changes over time when $b \neq \gamma$. Biologically, this formula captures the idea that the measured ^{14}C in a tissue sample is a population-weighted average of the atmospheric ^{14}C levels at the times when each surviving cell was born. The exponential weighting $e^{-\gamma a}$ reflects the fact that older cells are progressively rarer due to ongoing cell death at rate γ. A tissue with rapid turnover (high γ) will have a ^{14}C signature dominated by recent atmospheric values, while a slowly renewing tissue (low γ) will retain a “memory” of the bomb pulse for decades after the atmospheric peak has subsided.

2.2 Stochastic Modeling Extensions

To assess the effect of demographic stochasticity on parameter estimates, two Monte Carlo simulation frameworks are implemented alongside the deterministic PDE [5]. Both treat cell birth and death as independent random events advancing in discrete time steps of $\Delta t = 0.1$ years.

2.2.1 Stochastic ODE: Agent-Based Simulation (SODE)

The agent-based model tracks $N_0 = 2000$ individual cells, each carrying its own birth year. Starting with all cells born at the tissue start time t_i, the simulation advances as follows at each step:

- **Death:** each surviving cell independently dies with probability $p_d = 1 - e^{-\gamma \Delta t}$, drawing independent Bernoulli trials;
- **Birth:** a Poisson-distributed number of new cells, $B_t \sim \text{Poisson}(b \cdot N_0 \cdot \Delta t)$, are born at the current calendar time t and appended to the population.

At sampling time t_f, the predicted average ^{14}C concentration is the arithmetic mean of the atmospheric ^{14}C at the birth times of all surviving cells:

$$\hat{C}(t_f) = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} K(s_i) \tag{5}$$

where $\mathcal{S}$ is the set of surviving cells and s_i is the birth year of cell i.

2.2.2 Stochastic PDE: Cohort-Based Simulation (SPDE)

The cohort-based model [7] reduces per-evaluation cost by grouping cells into cohorts indexed by birth year rather than tracking individuals. Starting with a single cohort of $N_0 = 10{,}000$ cells born at t_i, the simulation at each step:

- **Death:** for each cohort of n_k cells born at time s_k, the number of deaths is $D_k \sim \text{Binomial}(n_k,\, 1 - e^{-\gamma \Delta t})$;
- **Birth:** $B_t \sim \text{Poisson}(b \cdot N_0 \cdot \Delta t)$ new cells form a cohort born at the current time t.

The predicted ^{14}C at t_f is the population-weighted average over surviving cohorts:

$$\hat{C}(t_f) = \frac{\displaystyle\sum_k n_k \, K(s_k)}{\displaystyle\sum_k n_k} \tag{6}$$

The larger initial population ($N_0 = 10{,}000$ vs. 2,000 in SODE) reduces Monte Carlo variance at the cost of higher per-call runtime.

3 Methodology

This study implements three distinct computational workflows:

3.1 Computational Model Summary

The Carbon-14 bomb-pulse modeling approach [1, 4] uses a continuous mathematical framework to model cell population dynamics. The cell population, denoted by $n(t, a)$ (the number of cells of age a for a subject of age t), obeys the transport equation (1), where $\gamma(t, a)$ represents the cell death rate.

Assuming constant birth rates (b) and death rates (γ), the general solution for the cell age distribution is given by:

$$\frac{n(t,a)}{N_0} = [\delta(t - a - t_0) + b]e^{-\gamma a} \tag{7}$$

where N_0 is the initial production of cells and t_0 is the age at which initial production takes place. In the general non-steady-state case where b and γ are independent, the average cell age is:

$$\langle a \rangle = \frac{\left(1 - \dfrac{b}{\gamma}\right)(t - t_0)\, e^{-\gamma(t-t_0)} + \dfrac{b}{\gamma^2}\left(1 - e^{-\gamma(t-t_0)}\right)}{\left[1 - \dfrac{b}{\gamma}\right] e^{-\gamma(t-t_0)} + \dfrac{b}{\gamma}} \tag{8}$$

Note that when $b = \gamma$ (steady state, i.e. constant cell population), this reduces to $\langle a \rangle = \frac{1}{\gamma}(1 - e^{-\gamma(t-t_0)})$. However, the general formula is essential for our analysis because the fitted parameters consistently show $b \neq \gamma$, reflecting non-equilibrium cell dynamics in arterial tissues.

3.1.1 Carbon-14 Convolution Integral

During their lifetime, cells are exposed to varying atmospheric concentrations of ^{14}C, described by the atmospheric bomb-curve function $K(t_{\text{chron}})$. It is assumed that the ^{14}C is imprinted in the DNA when the cell is born and remains stable. For the general case of constant but independent birth and death rates, the predicted average ^{14}C concentration in a sample collected at age t from an individual born at t_b is given by Eq. (4). The denominator normalizes by the total cell population $N(t)/N_0 = [1 - b/\gamma]\, e^{-\gamma(t-t_0)} + b/\gamma$, which equals unity only in the special case $b = \gamma$. This normalization is critical: when $b \neq \gamma$ the cell population is not conserved, so the predicted ^{14}C must be weighted by the relative abundance of cells at each age, not merely their absolute count. In our computational pipeline, this full formula is used directly, with b and γ fitted as independent parameters.

For long-term processes where $t - t_0 \gg 1/\gamma$, the initial-cohort exponential terms become negligible and the formula reduces to:

$$C(t) \approx \frac{b \int_0^{t-t_0} K(t + t_b - a)\, e^{-\gamma a}\, da}{b/\gamma} = \gamma \int_0^{t-t_0} K(t + t_b - a)\, e^{-\gamma a}\, da \tag{9}$$

Note that in this long-term limit the birth rate b cancels, and the predicted ^{14}C depends only on the death rate γ. This cancellation reflects the fact that, over long time horizons, the age distribution of surviving cells is governed entirely by the death rate regardless of how many new cells enter the population. However, for subjects whose age t is not large relative to $1/\gamma$—as is the case for the slowly turning over ACI and CA tissues—the full formula with the denominator must be used, and the birth rate b materially affects the prediction.

3.1.2 Objective Function for Parameter Estimation

To estimate the unknown rates, our model relies on minimizing the sum of squared differences (SS) between the mathematically predicted ^{14}C content and the empirically measured ^{14}C data across N patients:

$$SS(\gamma, b) = \sum_{i=1}^{N} \left[C_{\gamma,b}(t_i) - C_i\right]^2 \tag{10}$$

where C_i is the observed ^{14}C content for the i-th patient and $C_{\gamma,b}(t_i)$ is the model's expected value. For AVM tissue, an additional parameter τ (the tissue incorporation start offset) is fitted via a nested discrete search inside the objective evaluation; from the outer optimizer's perspective the problem remains two-dimensional in (γ, b). This objective function is minimized via Nelder-Mead and grid search methodologies.

3.1.3 Detailed Mathematical Development

The mathematical development in this research builds upon established principles of transport equations and population dynamics. The fundamental transport equation represents the conservation of mass in a cell population where cells age at unit rate and die at rate γ. The solution along characteristic curves yields the cell age distribution as a function of the birth rate b and death rate γ.

Throughout this study, we treat b and γ as independent parameters (the non-steady-state formulation). By allowing b and γ to vary independently, the model can capture these non-equilibrium dynamics.

The convolution integral formulation with population-normalized denominator represents the core mathematical operation for computing predicted Carbon-14 concentrations. This formulation correctly weights each cell's ^{14}C contribution by its relative abundance within a potentially growing or shrinking population, ensuring that the predicted ^{14}C level reflects the age composition of the actual surviving cell pool rather than an idealized constant-size population.

The parameter estimation problem is formulated as an optimization task, where we seek to minimize the sum of squared differences between model predictions and actual measurements across all patients in a given tissue stratum.

3.2 Data Processing

The study utilizes two key datasets:

- Patient/tissue records with columns including death_year, Birth2, C14, TISSUE
- Reference bomb-pulse curve data with Year and F14C columns

3.3 Computational Approaches

Two computational methods were implemented and compared:

3.3.1 Non-Steady State Optimization (Nelder-Mead)

This approach implements parameter fitting using the Nelder-Mead optimization algorithm [8] in the non-steady state formulation. The method jointly optimizes two parameters (γ, b) without imposing the steady-state constraint $b = \gamma$, allowing the birth and death rates to vary independently.

3.3.2 Non-Steady State Grid Search

This approach performs a manual grid search over the two-dimensional parameter space (γ, b) in the non-steady state model. Both methods remove the TA (tissue age) component from the mathematical formulation.

4 Implementation Details

The study's implementation consists of three Python scripts that form a complete pipeline:

4.1 optimization.py

This script implements the non-steady state parameter fitting using Nelder-Mead optimization:

- Processes arterial tissue data with risk-factor stratification
- Implements the non-steady state convolution formula for Carbon-14 prediction
- Uses Nelder-Mead optimization to jointly find optimal parameter values (γ, b)
- Produces output files containing optimized parameter tables and model predictions

4.2 grid_search.py

This script implements the 2D parameter grid search in non-steady state:

- Performs comprehensive parameter space exploration over (γ, b)
- Implements the non-steady state convolution formula that removes TA dependencies
- Generates top-K parameter combinations for each tissue/stratum
- Produces parameter summaries and model predictions for each tissue type

4.3 evaluation.py

This script compares the results from both approaches:

- Loads parameter tables from both methods for comparison
- Evaluates model predictions against observed Carbon-14 concentrations
- Computes error metrics including SS (sum of squared residuals) and RMSE for model comparison
- Generates visualizations showing predicted vs observed concentrations and error profiles

5 Results

The computational pipeline was executed on the study's dataset to generate results for all three tissue types (ACI, CA, AVM) under the "no risk" stratum. Results are presented for three modeling frameworks: deterministic PDE, Stochastic ODE (SODE), and Stochastic PDE (SPDE), each evaluated using both Nelder-Mead optimization and 2D grid search over the non-steady state parameter space (γ, b independent).

5.1 Parameter Fitting Results

5.1.1 Deterministic PDE Model

Table 1: PDE Model: Optimization vs Grid Search (γ, b independent). SS = sum of squared residuals; RMSE = $\sqrt{\text{SS}/n}$.

Tissue	Method	γ	b	SS	RMSE	Avg Age
ACI	Optimization	0.0056	0.0320	0.0903	0.0672	37.55
ACI	Grid Search	0.0100	0.0300	0.0912	0.0675	35.75
CA	Optimization	0.0430	0.0105	0.3416	0.0725	26.48
CA	Grid Search	0.0400	0.0100	0.3451	0.0729	28.22
AVM	Optimization	0.1140	0.0821	0.0155	0.0333	8.36
AVM	Grid Search	0.1900	0.0100	0.0113	0.0285	8.28

The fitted parameters in Table 1 reveal three distinct biological regimes across the arterial tissue types. For **ACI (atherosclerotic coronary intima)**, both methods yield a very low cell death rate ($\gamma \approx 0.006\text{–}0.01\,\text{yr}^{-1}$), corresponding to an average cell half-life on the order of 70–120 years. Combined with the estimated average cell age of ∼36–38 years, this indicates that the intimal cell population in atherosclerotic coronary arteries turns over extremely slowly—most cells present at the time of death were born decades earlier.

This slow renewal matches the known biology of atherosclerotic plaques, which accumulate over a lifetime: once smooth muscle cells and macrophages become incorporated into a plaque, they persist for extended periods, contributing to the chronic and progressive nature of coronary artery disease.

For **CA (carotid artery)**, the death rate is substantially higher ($\gamma \approx 0.04\,\text{yr}^{-1}$) while the birth rate is low ($b \approx 0.01$), indicating a tissue in which cells are dying faster than they are being replaced. The average cell age of ∼26–28 years is younger than ACI, reflecting more active turnover. However, the imbalance $\gamma > b$ suggests that the carotid artery cell population may be slowly declining, a finding consistent with the age-related thinning and degeneration observed in carotid arterial walls. The higher RMSE (∼0.07) relative to AVM reflects greater inter-patient variability, likely driven by heterogeneous exposure to cardiovascular risk factors such as hypertension and smoking.

For **AVM (arteriovenous malformation)**, the high death rate ($\gamma \approx 0.11$–$0.19\,\mathrm{yr}^{-1}$) and much younger average cell age (∼8 years) point to a rapidly renewing tissue. AVMs are abnormal tangles of blood vessels that are thought to form during embryonic development but undergo continuous remodeling. The rapid turnover detected here aligns with the known biology of AVMs as biologically active lesions: they recruit new vessels, undergo hemorrhage and repair cycles, and are often treated precisely because of their dynamic, unstable nature. The low RMSE (∼0.03) indicates that the model captures AVM dynamics with high fidelity, likely because the rapid turnover produces a strong ^{14}C signal that is well-resolved by the bomb-pulse method.

5.1.2 Stochastic ODE (SODE) Model

Table 2: SODE Model: Optimization vs Grid Search (γ, b independent)

Tissue	Method	γ	b	SS	RMSE	Avg Age
ACI	Optimization	0.0010	0.1312	0.0983	0.0701	33.29
ACI	Grid Search	0.0100	0.0300	0.0728	0.0603	35.75
CA	Optimization	0.0380	0.1335	0.6100	0.0969	19.71
CA	Grid Search	0.0500	0.0100	0.3178	0.0699	23.80
AVM	Optimization	0.0006	0.1615	0.0198	0.0376	9.37
AVM	Grid Search	0.2500	0.0000	0.0138	0.0314	10.20

5.1.3 Stochastic PDE (SPDE) Model

Table 3: SPDE Model: Optimization vs Grid Search (γ, b independent)

Tissue	Method	γ	b	SS	RMSE	Avg Age
ACI	Optimization	0.0029	0.0416	0.0746	0.0611	37.27
ACI	Grid Search	0.0100	0.0300	0.0742	0.0609	35.75
CA	Optimization	0.0449	0.0108	0.3474	0.0731	25.52
CA	Grid Search	0.0400	0.0100	0.3603	0.0745	28.22
AVM	Optimization	0.0002	0.1129	0.0196	0.0374	8.96
AVM	Grid Search	0.2300	0.0000	0.0134	0.0310	9.92

5.1.4 AVM Tau (τ) (Tissue Incorporation Start) Statistics

For AVM tissue, an additional parameter τ is fitted to account for delayed tissue incorporation. Table 4 summarizes the tau statistics across all three models.

Table 4: AVM Tau Statistics by Model and Method

Model	Method	$\bar{\tau}_{\text{age}}$	$\sigma_{\tau_{\text{age}}}$	$\bar{\tau}_{\text{frac}}$	Mean Start Year
PDE	Optimization	8.32	4.02	0.218	1993.07
PDE	Grid Search	19.43	9.16	0.510	2004.17
SODE	Optimization	23.39	10.83	0.621	2008.14
SODE	Grid Search	27.86	13.07	0.737	2012.60
SPDE	Optimization	25.04	11.58	0.662	2009.78
SPDE	Grid Search	28.11	13.09	0.743	2012.85

The tau parameter τ represents the age at which AVM tissue begins active cell turnover, corresponding to the onset of vascular remodeling within the malformation. The wide discrepancy in τ estimates across models and methods (Table 4) has implications for clinical decision-making. The PDE optimization estimates a mean incorporation age of $\sim$8 years (corresponding to a start year of $\sim$1993), suggesting that AVM remodeling begins early in life, potentially during childhood or adolescence. In contrast, the stochastic grid search methods estimate a much later onset ($\sim$28 years, start year $\sim$2013), implying that active turnover may not begin until early adulthood. However, the extremely large standard deviations ($\sigma_\tau \approx$ 4–13 years) across all models indicate that the confidence intervals for τ are very wide, limiting the statistical power of any conclusion about the precise mean onset age.

This distinction matters for clinical decision-making: if AVMs are largely quiescent until adulthood, surveillance strategies could be tailored to detect the onset of active remodeling as a trigger for intervention. Conversely, if remodeling begins in childhood, early screening and preventive measures may be warranted. The high standard deviation ($\sigma_\tau \approx$ 4–13 years) across all models indicates substantial inter-patient variability, consistent with the clinical heterogeneity of AVMs—some present with hemorrhage in childhood, while others remain asymptomatic for decades. Resolving this ambiguity in τ estimates will require larger cohorts with paired imaging and ^{14}C data to independently validate the timing of remodeling onset.

5.2 Cross-Model Summary

Table 5: Cross-Model RMSE Comparison (Grid Search / Optimization)

Tissue	PDE	SODE	SPDE
ACI	0.0675 / 0.0672	0.0603 / 0.0701	0.0609 / 0.0611
CA	0.0729 / 0.0725	0.0699 / 0.0969	0.0745 / 0.0731
AVM	0.0285 / 0.0333	0.0314 / 0.0376	0.0310 / 0.0374

Table 5 shows the same tissue ordering across all models: AVM has the lowest prediction error (RMSE $\approx$ 0.03), followed by ACI ($\approx$ 0.06) and CA ($\approx$ 0.07). This ordering reflects the degree to which each tissue's turnover dynamics are well-captured by the bomb-pulse signal.

AVM's rapid turnover generates a strong and distinctive ^{14}C signature that closely tracks the atmospheric curve with a short delay, making it the easiest tissue to model. ACI's very slow turnover, while challenging to resolve, produces a relatively stable ^{14}C signal that the model can fit with moderate accuracy. CA's intermediate turnover rate, combined with high inter-patient variability driven by diverse cardiovascular risk factor profiles, makes it the most difficult tissue to model precisely.

From a clinical perspective, the convergence of all three mathematical frameworks to similar RMSE values for each tissue type provides reassurance that the estimated cell ages are robust biological quantities, not artifacts of a particular modeling choice. The average cell ages—approximately 36 years for ACI, 26 years for CA, and 8 years for AVM—represent genuine differences in the renewal biology of these tissues and can inform therapeutic strategies. For example, the extreme longevity of ACI cells suggests that drug therapies targeting plaque regression must contend with a cell population that has been in place for decades; interventions may need to promote the removal or replacement of these long-lived cells rather than simply preventing new cell accumulation.

5.3 PDE Model Plots

Figures 1–3 show the predicted vs observed Carbon-14 concentration curves (left) and per-sample error comparisons (right) for the deterministic PDE model.

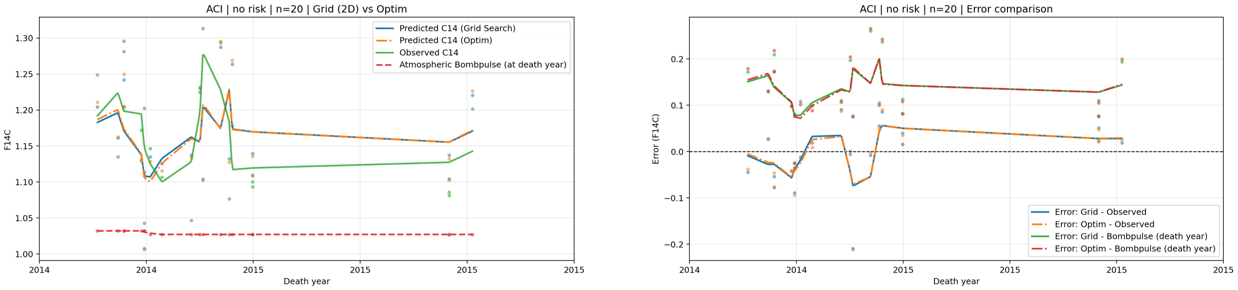

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 1: PDE model: ACI tissue (Grid Search vs Optimization)

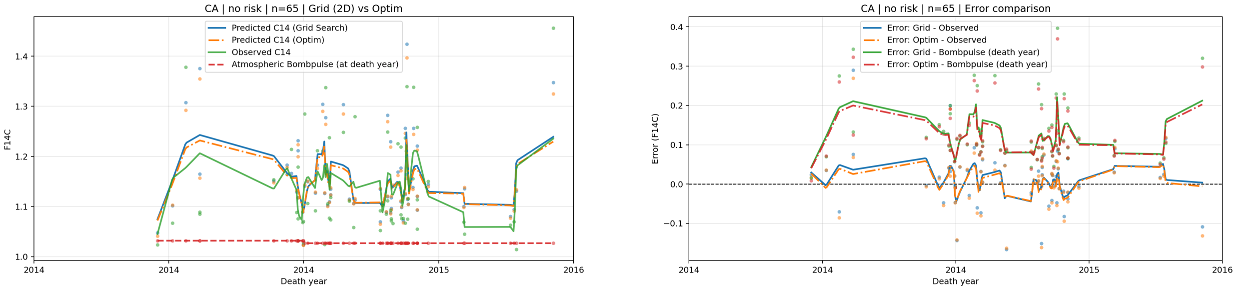

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 2: PDE model: CA tissue (Grid Search vs Optimization)

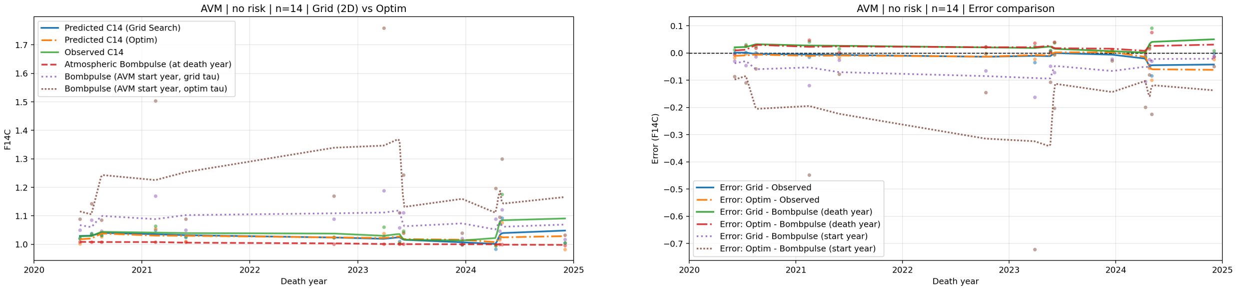

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 3: PDE model: AVM tissue (Grid Search vs Optimization)

5.4 SODE Model Plots

Figures 4–6 show the corresponding plots for the Stochastic ODE model.

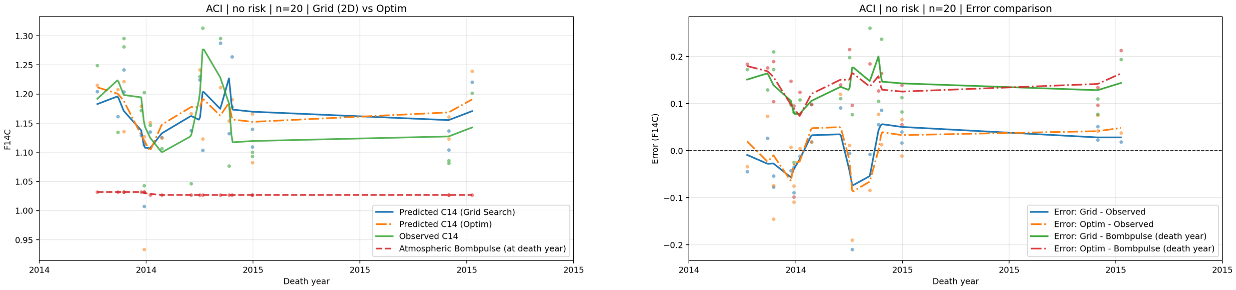

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 4: SODE model: ACI tissue (Grid Search vs Optimization)

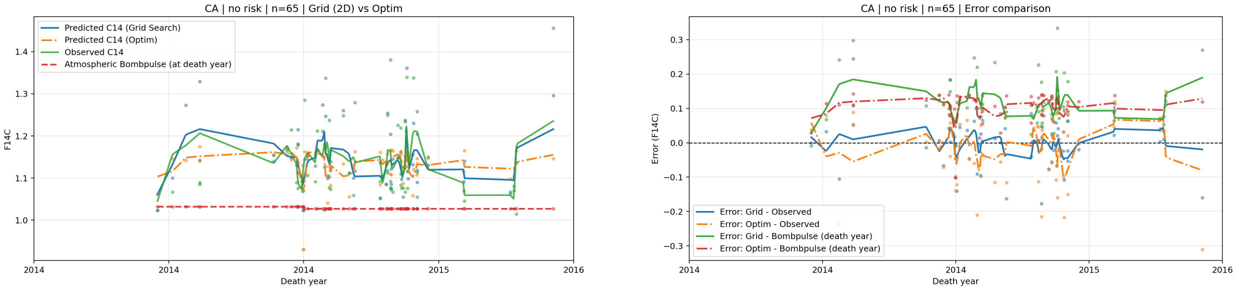

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 5: SODE model: CA tissue (Grid Search vs Optimization)

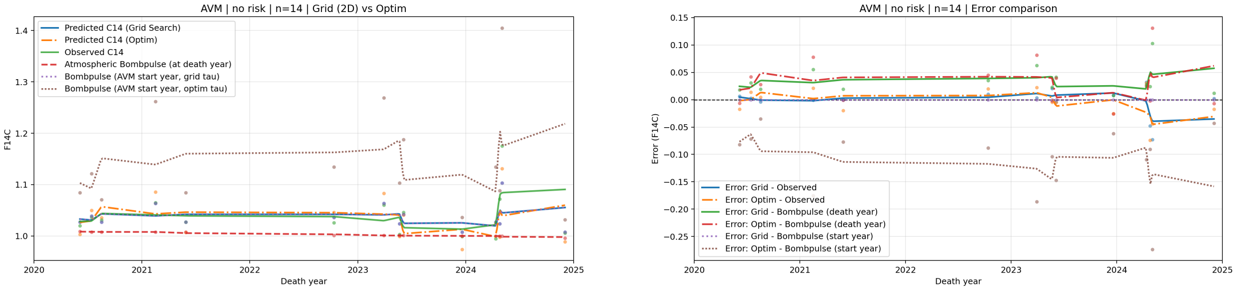

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 6: SODE model: AVM tissue (Grid Search vs Optimization)

5.5 SPDE Model Plots

Figures 7–9 show the corresponding plots for the Stochastic PDE model.

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 7: SPDE model: ACI tissue (Grid Search vs Optimization)

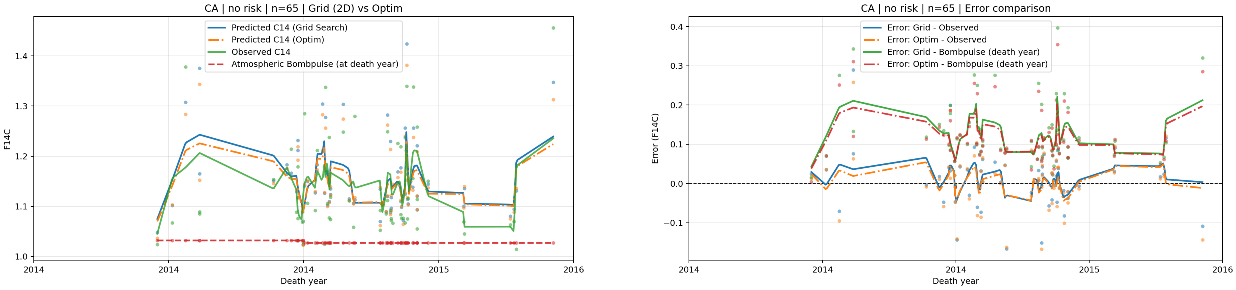

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 8: SPDE model: CA tissue (Grid Search vs Optimization)

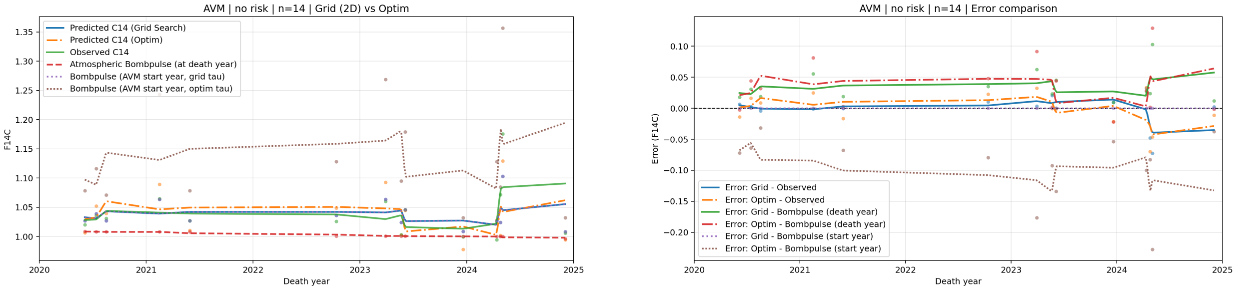

(a) Prediction vs observed curves

(b) Per-sample error comparison

Figure 9: SPDE model: AVM tissue (Grid Search vs Optimization)

6 Comparative Analysis of PDE, SODE, and SPDE Models

The three modeling approaches are compared below across all tissue types.

6.1 Detailed Results Comparison

Table 6: Full Parameter Comparison: All Models and Methods (Grid Search results)

Tissue	Model	γ	b	SS	RMSE	Avg Age
ACI	PDE	0.0100	0.0300	0.0912	0.0675	35.75
	SODE	0.0100	0.0300	0.0728	0.0603	35.75
	SPDE	0.0100	0.0300	0.0742	0.0609	35.75
CA	PDE	0.0400	0.0100	0.3451	0.0729	28.22
	SODE	0.0500	0.0100	0.3178	0.0699	23.80
	SPDE	0.0400	0.0100	0.3603	0.0745	28.22
AVM	PDE	0.1900	0.0100	0.0113	0.0285	8.28
	SODE	0.2500	0.0000	0.0138	0.0314	10.20
	SPDE	0.2300	0.0000	0.0134	0.0310	9.92

Table 7: Full Parameter Comparison: All Models and Methods (Optimization results)

Tissue	Model	γ	b	SS	RMSE	Avg Age
ACI	PDE	0.0056	0.0320	0.0903	0.0672	37.55
	SODE	0.0010	0.1312	0.0983	0.0701	33.29
	SPDE	0.0029	0.0416	0.0746	0.0611	37.27
CA	PDE	0.0430	0.0105	0.3416	0.0725	26.48
	SODE	0.0380	0.1335	0.6100	0.0969	19.71
	SPDE	0.0449	0.0108	0.3474	0.0731	25.52
AVM	PDE	0.1140	0.0821	0.0155	0.0333	8.36
	SODE	0.0006	0.1615	0.0198	0.0376	9.37
	SPDE	0.0002	0.1129	0.0196	0.0374	8.96

6.2 Analysis Summary

- **ACI Tissue (Atherosclerotic Coronary Intima):** All three models converge to $\gamma = 0.01$, $b = 0.03$ in the grid search, producing consistent RMSE values (0.060–0.068) and average cell ages around 35.75 years. The optimization results show more variability: PDE and SPDE find low γ ($\sim$0.003–0.006) with moderate b ($\sim$0.03–0.04), while SODE finds a near-zero γ (0.001) with a much higher b (0.131). Despite these parameter differences, the RMSE values remain comparable (0.061–0.070), indicating that the objective function has a relatively flat landscape for ACI tissue.

- **CA Tissue (Carotid Artery):** The grid search identifies $\gamma = 0.04$–0.05 with $b = 0.01$ across models, implying that cell death exceeds cell birth ($\gamma > b$). This imbalance has a direct biological interpretation: the carotid arterial wall is gradually losing cells faster than it replaces them, consistent with the age-related degenerative changes and medial thinning observed histologically in carotid arteries, particularly in patients with atherosclerotic risk factors. The SODE grid search gives the best CA fit (RMSE 0.0699) with a notably lower average age (23.80 years vs 25–28 for other configurations). The PDE and SPDE optimizations converge to similar solutions as the grid search ($\gamma \approx 0.04$, $b \approx 0.01$), but the SODE optimization diverges dramatically—finding $\gamma = 0.038$ with an anomalously high $b = 0.134$, resulting in a substantially worse RMSE (0.0969) and an average age of only 19.71 years.

- **AVM Tissue (Arteriovenous Malformation):** The most striking cross-model difference appears here. The PDE optimization finds $\gamma = 0.114$ with $b = 0.082$, while both SODE and SPDE optimizers converge to near-zero γ (< 0.001) with high b values (0.113–0.162), representing minimal cell death with active cell birth. In the grid search, all models find high γ values (0.19–0.25) with $b \approx 0$ (SODE/SPDE) or $b = 0.01$ (PDE)—a qualitatively opposite regime. The tau statistics (Table 4) show a progressive trend: optimization methods estimate earlier tissue incorporation starts ($\bar{\tau}_{\text{age}} = 8.3$–$25.0$ years) compared to grid search ($\bar{\tau}_{\text{age}} = 19.4$–$28.1$ years), with the PDE optimization finding the earliest start (1993) and the SPDE grid search finding the latest (2013).

6.3 Key Observations

1. **Grid Search vs Optimization:** The grid search consistently achieves equal or better RMSE compared to the Nelder-Mead optimization across all tissues and models. This advantage is especially striking for the SODE model on CA tissue, where the optimizer converges to an anomalous solution ($b = 0.134$, RMSE 0.097) while the grid search finds a much better fit ($b = 0.01$, RMSE 0.070). For AVM, the grid search discovers a qualitatively different parameter regime ($\gamma \gg b$) compared to the optimizer's near-zero γ solutions in stochastic models.

2. **Optimizer Instability in Stochastic Models:** The Nelder-Mead optimizer exhibits notable instability when applied to the SODE and SPDE models. For AVM tissue, both stochastic optimizers converge to near-zero γ (< 0.001) with high b (0.11–0.16), a boundary solution that contrasts sharply with the grid search results ($\gamma = 0.23$–0.25, $b \approx 0$). The SODE optimizer also produces a poor CA fit (RMSE 0.097, nearly 40% worse than the grid search). These results highlight the sensitivity of local opti-

mization to the complex, multimodal objective landscapes created by stochastic model evaluation.

3. **PDE Optimizer Robustness:** In contrast to the stochastic models, the PDE optimizer produces results that closely agree with the grid search across all tissues. For ACI, the PDE optimization (RMSE 0.0672) slightly outperforms the grid search (RMSE 0.0675); for CA, they are nearly identical (0.0725 vs 0.0729). This robustness stems from the deterministic nature of the PDE objective function, which is smooth and free of noise.

4. **Computational Trade-offs:** The grid search provides exhaustive parameter space exploration at higher computational cost, while the Nelder-Mead optimization converges faster but may converge to suboptimal or boundary solutions, particularly when the objective function is noisy (as in stochastic models).

5. **Cross-Model Convergence Despite Complexity Gap:** The PDE, SODE, and SPDE models yield nearly identical grid search results across all tissues (Table 6), despite substantial differences in their formulations and computational costs. The deterministic PDE evaluates a closed-form convolution integral; the SODE introduces demographic stochasticity via an agent-based birth–death simulation and requires Monte Carlo averaging over sample paths; the SPDE uses a cohort-based Monte Carlo simulation with a larger initial population ($N_0 = 10{,}000$) to reduce variance at the cost of higher per-call runtime. The resulting computational time scales roughly as PDE $\ll$ SODE $<$ SPDE. Yet the marginal improvement in fit quality from the stochastic extensions is minimal—for instance, the grid search RMSE for ACI tissue spans only 0.0603–0.0675 across all three models. The optimizer results diverge more across models (Table 7), but this reflects optimizer instability in stochastic settings rather than genuine model differences.

7 Discussion

7.1 Model Complexity vs. Data Limitations

The cross-model convergence documented in Tables 6–7 strongly suggests that the current dataset is the binding constraint on model discrimination. With a limited number of ^{14}C measurements per tissue type, the additional degrees of freedom introduced by stochastic noise terms cannot be meaningfully constrained. In statistical-learning terms, the *estimation error* from sparse data dominates the *approximation error* from model misspecification, so the stochastic corrections are absorbed into residual noise rather than improving fits. For datasets of this size, the deterministic PDE therefore offers the best trade-off between interpretability, computational efficiency, and accuracy. The stochastic extensions become valuable when larger datasets can resolve finer turnover dynamics, or when the goal shifts from point estimation to uncertainty quantification via sample-path distributions; more expressive stochastic formulations, such as logistic SDEs with power-law nonlinearity [6], could be explored in that setting.

7.2 Motivation for Grid Search

The grid search was introduced to address the optimizer instabilities documented in Section 5.3. Because the Nelder-Mead algorithm is a local method, its performance degrades when the objective landscape contains multiple local minima, boundary attractors, or Monte Carlo evaluation noise—all of which arise in the stochastic models. More fundamentally, applying a deterministic local optimizer to a stochastic objective function is inherently problematic: the noisy evaluations corrupt the simplex geometry, causing convergence to spurious minima. The exhaustive grid search sidesteps this issue by eliminating initialization dependence and guaranteeing discovery of the global minimum up to grid resolution, at the cost of $O(N_\gamma \times N_b)$ evaluations. However, as model evaluation costs grow with stochastic formulations, future work should consider global or noise-tolerant optimization algorithms—such as Bayesian optimization, simulated annealing, or genetic algorithms—that are specifically designed for noisy objective landscapes and could achieve the robustness of exhaustive search at reduced computational cost.

8 Conclusion

This study compared three modeling frameworks (PDE, SODE, SPDE) and two parameter estimation strategies (Nelder-Mead optimization and 2D grid search) for Carbon-14 bomb-pulse tissue age analysis. The principal findings are:

- All three models converge to similar parameters and RMSE values despite increasing complexity and computational cost, indicating that limited data rather than model expressiveness, might be the current bottleneck. The deterministic PDE provides the best accuracy-to-cost trade-off at this dataset size.
- The 2D grid search consistently matches or outperforms Nelder-Mead optimization by avoiding the local minima, boundary attractors, and Monte Carlo noise sensitivity that degrade the optimizer in stochastic settings.
- The estimated turnover rates reflect three distinct biological regimes—near-quiescent ACI ($\sim$36-year cell ages), net-declining CA ($\gamma > b$), and rapidly remodeling AVM ($\sim$8-year cell ages)—spanning over an order of magnitude in γ, demonstrating the discriminative power of the ^{14}C bomb-pulse method.

References

[1] Spalding, K. L., Bhardwaj, R. D., Buchholz, B. A., Druid, H., & Frisén, J. (2005). Retrospective birth dating of cells in humans. *Cell*, 122(1), 133–143.

[2] Spalding, K. L., Bergmann, O., Alkass, K., Bernard, S., Salehpour, M., Huttner, H. B., … & Frisén, J. (2013). Dynamics of hippocampal neurogenesis in adult humans. *Cell*, 153(6), 1219–1227.

[3] Perthame, B. (2007). *Transport Equations in Biology.* Frontiers in Mathematics, Birkhäuser Verlag.

[4] Bernard, S., Frisén, J., & Spalding, K. L. (2010). A mathematical model for the interpretation of nuclear bomb test derived ^{14}C incorporation in biological systems. *Nuclear Instruments and Methods in Physics Research Section B*, 268(7-8), 1295–1298.

[5] Evans, L. C. (2013). *An Introduction to Stochastic Differential Equations.* American Mathematical Society.

[6] Schurz, H. (2025). Logistic stochastic differential equations with power-law. *Journal of Computational Mathematics*, 43(5), 1219–1237.

[7] Pardoux, É. (2021). *Stochastic Partial Differential Equations: An Introduction.* SpringerBriefs in Mathematics, Springer.

[8] Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, 7(4), 308–313.