RESEARCH ARTICLE

Two chemical pulses distinguish persistent from fluctuating extracellular acidity in silico

New York General Group, Inc.

7 November 2026

Abstract

A region beside a tumour cell can remain acidic or repeatedly acidify and recover. These histories can produce the same spatial distribution of extracellular pH, yet expose a molecule anchored at one site to different conditions over time. Here we ask whether two pulses of an irreversible acylating reagent can distinguish such histories without tracking individual sites. We show that the fraction of reporter escaping both pulses, after subtraction of an independently measurable baseline, equals the covariance of the two local escape probabilities. In numerical experiments with identical pH distributions, fixed acidity produced a constant positive signal, random switching produced an exponentially decaying signal, and periodic switching produced alternating positive and negative signals. For a reporter controlled by one protonation equilibrium, a pH range of 0.20 units limits the absolute signal at any one delay to 0.705 percentage points, irrespective of pulse duration. Fluctuations in reagent concentration can, however, reproduce the entire pH-dependent response exactly. The assay therefore measures the persistence of **local chemical reactivity**; attributing that persistence to tumour acidity requires independent measurements of pH and reagent exposure.

Main

An extracellular pH map shows where a tumour is acidic, but a single map cannot show whether the same sites remain acidic. This distinction matters for a molecule tethered to the extracellular matrix. If half of the sites are acidic at any instant, a tethered molecule might spend its entire lifetime in an acidic region, or every molecule might spend half its lifetime there. The instantaneous pH distribution is identical in the two cases. What differs is how well the environment at one site predicts its environment later.

Previous experiments establish several ways to observe tumour acidity. Fluorescent nanoprobes have localized sharply acidic extracellular regions near cancer cells. Irreversibly activated probes have integrated tumour-acidosis signals over time, and ratiometric sensors have tracked extracellular pH in three-dimensional pancreatic cancer cultures and patient-derived tumourspheres. Paired measurements separated by a variable delay also have a precedent in cancer research: diffusion MRI has used this principle to study water exchange in breast tissue. The question addressed here is narrower: **can the surviving fraction of an immobilized chemical reporter encode temporal persistence, and how large can that signal be?** [1]

A second pulse tests the sites that escaped the first

Consider a matrix-tethered peptide with one accessible N-terminal primary amine. A pulse of an activated acetate ester irreversibly acetylates a fraction of the reporter. After a reagent-free interval, a second pulse of equal concentration and duration acts on the remaining amines. The endpoint is the fraction of the original reporter that remains unacetylated. Activated-ester aminolysis provides a chemical basis for this design, but competing hydrolysis and reaction with other nucleophiles make the concentration of **intact local ester**, rather than the nominal applied dose, the relevant quantity. [2]

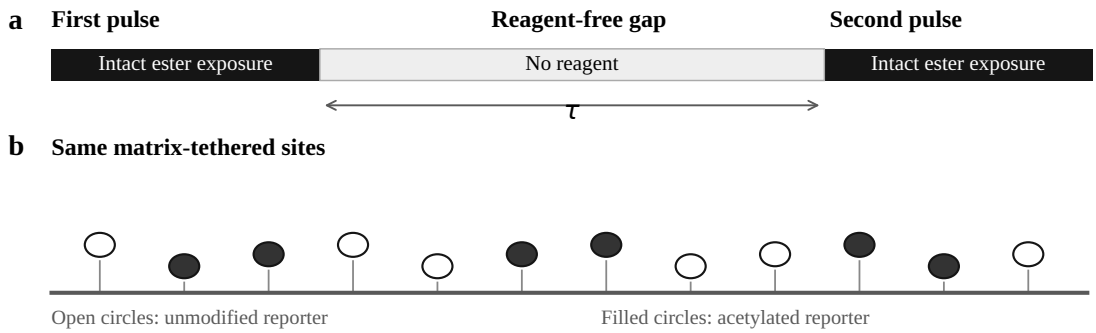


Figure 1. Paired pulses interrogate the same matrix-tethered reporter sites. The second pulse follows a reagent-free gap τ, measured from the end of the first exposure to the start of the second.

Let δ be the pulse duration and τ the interval **from the end of the first pulse to the start of the second**. For a reporter at a given site, define

$$A_j = \int_{\text{pulse } j} k(t)\,\mathrm{d}t, \qquad U_j = \exp(-A_j),$$

where $k(t)$ is the local pseudo-first-order acylation rate in s^{-1}. Thus U_j is the probability that an initially unmodified reporter would escape pulse j, conditional on its chemical history. Assume that the reporter remains at its site, that no reaction occurs during the gap and that the first pulse does not change the subsequent chemical environment. Under stationary conditions with matched pulses, the one-pulse unmodified fraction is $S_1 = \mathbb{E}[U_1] = \mathbb{E}[U_2]$. The two-pulse fraction is $S_2(\tau) = \mathbb{E}[U_1 U_2]$. It follows that

$$\boxed{D(\tau) = S_2(\tau) - S_1^2 = \mathrm{Cov}(U_1, U_2).} \tag{1}$$

The subtraction gives the signal a physical interpretation. If the sites spared by the first pulse are also preferentially spared by the second, $D > 0$. If the second pulse preferentially modifies the sites the first pulse spared, $D < 0$. Independent exposures give $D = 0$, although the converse does not follow: a single zero measurement does not establish independence. Importantly, equation (1) measures a covariance of **reaction probabilities integrated over finite pulses**, not the instantaneous autocovariance of pH.

Three specimens with the same pH distribution give different delay curves

We compared three idealized specimens. In each, a randomly selected reporter site had probability $1/2$ of being at pH 6.6 and probability $1/2$ of being at pH 6.8 at every observation time. In the **fixed** specimen, each site retained its initial pH. In the **randomly switching** specimen, each site changed state at a constant rate of $0.025\ \mathrm{s}^{-1}$ in either direction. In the **periodically switching** specimen, each site alternated between 20 s at each pH, with starting times distributed across the 40-s repeat. Thus no oscillation in the specimen-wide pH distribution was needed to produce a periodic local signal.

For the numerical experiment, both pulses lasted 10 s. The assumed apparent $\mathrm{p}K_a$ was 8.00, and the pseudo-first-order rate for a fully unprotonated reporter was $1.00\ \mathrm{s}^{-1}$. The resulting rates were $0.03829\ \mathrm{s}^{-1}$ at pH 6.6 and $0.05935\ \mathrm{s}^{-1}$ at pH 6.8. These values and switching times are simulation inputs, not measured properties of a peptide or tumour.

The one-pulse endpoints differed little: **61.715%, 61.664% and 61.602%** of reporter remained unmodified in the fixed, random and periodic specimens, respectively. The paired pulses distinguished the three cases by the *shape and sign* of the delay response (Table 1).

Table 1. Paired-pulse signal at matched instantaneous pH distributions

Reagent-free gap τ (s)	Fixed	Randomly switching	Periodically switching
0	+0.4194	+0.2597	−0.0001
10	+0.4194	+0.1575	−0.2807
20	+0.4194	+0.0955	−0.0001
30	+0.4194	+0.0579	+0.2811
50	+0.4194	+0.0213	−0.2807
70	+0.4194	+0.0078	+0.2811

Values are $100D(\tau)$, in **percentage points of the original reporter population**, not relative percentage changes. They are deterministic population averages without added measurement noise. The displayed precision is numerical, not experimental.

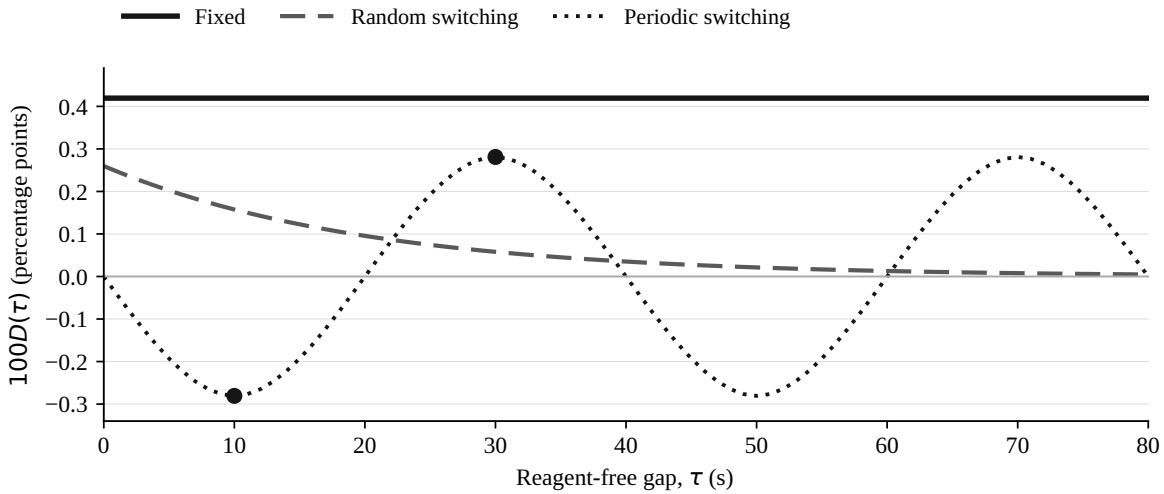


Figure 2. Paired-pulse delay responses for fixed, randomly switching and periodically switching acidity. Curves are determined by the simulation inputs in Methods; no measurement noise is included.

The periodic specimen explains why the definition of delay matters. With a **10-s gap**, the pulses start 20 s apart, exactly half a repeat. Every site therefore receives complementary exposures: a more reactive instant in one pulse is paired with a less reactive instant in the other. The combined exposure is identical at all sites, and $S_2 = 37.6674\%$. With a **30-s gap**, the pulses start 40 s apart. Each site repeats its first exposure and $S_2 = 38.2292\%$. Changing only the gap therefore raises the raw unmodified fraction by **0.5618 percentage points**. By contrast, the same change in gap produces no change in the fixed specimen and a decrease of **0.0996 percentage points** in the randomly switching specimen. This comparison requires no estimate of the one-pulse baseline.

Random switching yields a correlation time, not an acidic residence time

The exponential response in Table 1 follows exactly from the assumed two-state dynamics; it was not imposed by curve fitting. Let a site switch from state 1 to state 2 at rate a, and back at rate b. At stationarity, the occupancies are $p_1 = b/(a + b)$ and $p_2 = a/(a + b)$. If u_i is the probability of escaping a pulse that starts in state i, **including any switching during that pulse**, then

$$D(\tau) = p_1 p_2 (u_1 - u_2)^2 e^{-(a+b)\tau}. \tag{2}$$

To obtain equation (2), condition on the states at the two ends of the gap. Stationary two-state switching has the same statistics forwards and backwards in time, so the probability of escaping the *preceding* pulse conditional on its final

state i is also u_i. During the gap, the difference between the two conditional state probabilities decays as $e^{-(a+b)\tau}$. Averaging the two escape probabilities over the endpoint states gives equation (2).

For the chosen rates, the correlation time is $1/(a+b) = 20\,\mathrm{s}$, whereas the mean stay at either pH is $1/a = 1/b = 40\,\mathrm{s}$. Reporting the fitted 20 s as the duration of an acidic episode would therefore be incorrect. Equation (2) also provides a check on the proposed chemistry: changing pulse duration can change the amplitude through $u_1 - u_2$, but should not change the decay rate if the same two-state process remains unperturbed.

One protonation equilibrium limits the chemical contrast

Suppose that only the unprotonated amine reacts, protonation equilibrates rapidly relative to acylation, and the concentration $[E]$ of intact ester is spatially uniform and constant during each pulse. The assumed rate law is

$$k(\mathrm{pH}) = k_{\mathrm{chem}}[E]\,\frac{1}{1 + 10^{\mathrm{p}K_a - \mathrm{pH}}}, \tag{3}$$

where k_{chem} is the second-order rate constant of the reactive amine. Equation (3) is a hypothesis to be tested on the **immobilized** reporter; the pH dependence of aminolysis does not by itself establish this rate law in a matrix. [2]

If the entire pH excursion is confined to an interval of width $\Delta\mathrm{pH}$, the fastest and slowest rates satisfy $r = k_{\mathrm{fast}}/k_{\mathrm{slow}} \le 10^{\Delta\mathrm{pH}}$. Whatever happens within a pulse, its escape probability lies between $e^{-k_{\mathrm{fast}}\delta}$ and $e^{-k_{\mathrm{slow}}\delta}$. The variance of a quantity in an interval of width w is at most $w^2/4$. The absolute covariance of two such quantities is therefore also at most $w^2/4$, giving

$$|D(\tau)| \le \frac{1}{4}\left(e^{-k_{\mathrm{slow}}\delta} - e^{-k_{\mathrm{fast}}\delta}\right)^2. \tag{4}$$

For $r > 1$, the right-hand side is largest at $\delta_* = \ln(r)/(k_{\mathrm{fast}} - k_{\mathrm{slow}})$. Consequently,

$$\boxed{|D(\tau)| \le \frac{1}{4}\left[(1 - r^{-1})r^{-1/(r-1)}\right]^2.} \tag{5}$$

A full pH range of **0.10, 0.20 or 0.50 units** therefore gives respective ceilings of **0.179, 0.705 or 4.030 percentage points** on the absolute signal at any one delay. These are upper bounds, not expected effect sizes. For the specific rates and 10-s pulses used in Table 1, the tighter bound from equation (4) is **0.4194 percentage points**. The fixed, equally populated specimen attains it.

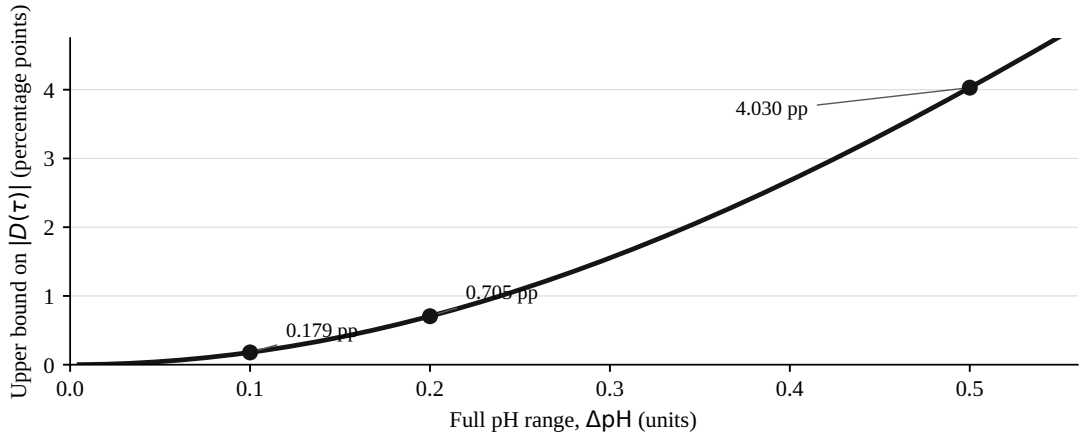


Figure 3. The single-delay covariance ceiling under a single protonation equilibrium. The labelled points correspond to full pH ranges of 0.10, 0.20 and 0.50 units; bounds are not expected effect sizes.

This distinction reconciles two apparently conflicting results: a **0.5618-point** difference between two delays is compatible with a **0.4194-point** ceiling at each delay because the periodic signal crosses zero. A peak-to-trough difference can be as large as twice the single-delay bound. Conversely, if an independent measurement confines the

full pH excursion to 0.20 units, a **single-delay** $|D|$ of two percentage points cannot be explained by equation (3) with uniform reagent delivery. A signal that large would be evidence against the assumed chemistry, not stronger evidence of pH persistence.

Discussion

The principal limitation is that an acylation rate is not a pH measurement. At a constant pH of 6.7, reducing the local intact ester concentration to **0.8022** times its nominal value reproduces the rate assigned to pH 6.6; increasing it to **1.2436** times nominal reproduces the rate at pH 6.8. Before rounding, these substitutions give exactly the same $k(t)$. If reagent delivery follows the same time course as the proposed pH changes, every one- and two-pulse endpoint is identical. Greater analytical precision cannot resolve this ambiguity.

A decisive first experiment would therefore use a **cell-free matrix**, not a tumour specimen. The tethered reporter should first be calibrated at fixed pH, with modified and unmodified peptide quantified separately by isotope-dilution liquid chromatography–mass spectrometry. Intact ester exposure and washout should be measured independently. A blinded comparison could then impose the three histories in Table 1 and prespecify $S_2(30\,\mathrm{s}) - S_2(10\,\mathrm{s})$ as the endpoint. Its predicted values are zero, negative and positive for fixed, random and periodic conditions, respectively. Precision across independently prepared specimens—not repeated injections of the same extract—would determine whether these differences are resolvable. Hydrolysis and off-target aminolysis are particular reasons to verify pulse delivery rather than infer it from the inlet concentration. [2]

Only after that test would a pancreatic tumoursphere experiment provide a meaningful biological comparison. A co-located ratiometric pH sensor, independently shown to resolve the pulse timescale, could supply local pH histories during sham and active exposures. Those histories, the measured reagent exposure and the cell-free calibration would predict $S_2(\tau)$ through equation (1), without fitting a switching time to the chemical endpoint. Optical measurements in pancreatic cancer cultures and patient-derived tumourspheres provide a precedent for this comparison; they do not establish that tumours undergo the 40-s oscillation imposed here. [3]

The value of paired pulses is thus not that they reconstruct a unique pH history. They do not. Rather, they turn an endpoint measurement into a test of whether the first exposure predicts the second at the same site. A flat response over a finite observation window cannot prove that acidity is static, and a negative response cannot by itself prove periodicity. But the sign, timescale and chemically permitted magnitude of the response together make specific explanations testable—and, in some cases, rule them out before tumour material is used.

Methods

Deterministic population calculation

For random switching, let Q contain the transition rates and K the chemical reaction rates, using row-vector probabilities:

$$Q = \begin{pmatrix} -a & a \\ b & -b \end{pmatrix}, \qquad K = \mathrm{diag}(k_1, k_2), \qquad P_\delta = \exp[(Q - K)\delta].$$

With initial probabilities $p = (p_1, p_2)$ and a column of ones 1, the unmodified fractions are

$$S_1 = pP_\delta 1, \qquad S_2(\tau) = pP_\delta e^{Q\tau} P_\delta 1.$$

The fixed calculation uses $Q = 0$ and equal initial occupancies. For the periodic calculation, 40 equally weighted sites start at successive 1-s offsets. All pulse and switching boundaries fall on this grid, so summing the exposure is exact for the specified 40-site specimen. No random trajectories were sampled, no parameters were fitted and no biological error bars are implied.

Data and code availability

All numerical inputs are given above. The following Python code reproduces Table 1 and the raw two-delay contrast; `scipy.linalg.expm` evaluates the matrix exponential. No patient data or external dataset is required. [4]

```
import numpy as np
from scipy.linalg import expm

pulse = 10  # seconds
gaps = (0, 10, 20, 30, 50, 70)
k = 1.0 / (1.0 + 10.0**(8.0 - np.array([6.6, 6.8])))  # s^-1
p = np.array([0.5, 0.5])

def switching(gap, rate):
    q = np.array([[-rate, rate], [rate, -rate]])
    exposure = expm((q - np.diag(k)) * pulse)
    after_one = p @ exposure
    after_two = after_one @ expm(q * gap) @ exposure
    return after_one.sum(), after_two.sum()

# Rate in each successive 1-s interval of a 40-s repeat.
rate_by_second = np.repeat(k, 20)
phase = np.arange(40)
u = np.exp(-sum(rate_by_second[(phase + t) % 40]
                for t in range(pulse)))

def periodic(gap):
    second = u[(phase + pulse + gap) % 40]
    return u.mean(), np.mean(u * second)

specimens = {
    "fixed": lambda gap: switching(gap, 0.0),
    "random": lambda gap: switching(gap, 0.025),
    "periodic": periodic,
}

print("gap_s     fixed    random  periodic  [percentage points]")
for gap in gaps:
    signals = []
    for specimen in specimens.values():
        s1, s2 = specimen(gap)
        signals.append(100 * (s2 - s1**2))
    print(f"{gap:5d}" + "".join(f"{x:10.4f}" for x in signals))

print("\nS2(30 s) - S2(10 s) [percentage points]")
for name, specimen in specimens.items():
    contrast = 100 * (specimen(30)[1] - specimen(10)[1])
    print(f"{name:8s} {contrast:+.4f}")
```

References

1. https://www.nature.com/articles/s41551-024-01178-7.

2. https://pubmed.ncbi.nlm.nih.gov/3667072/.

3. https://pubmed.ncbi.nlm.nih.gov/35617754/.

4. https://docs.scipy.org/doc/scipy/reference/generated/scipy.linalg.expm.html.