Preprint | Stochastic Analysis

The Contact Field: Ray-Knight Laws and Reaction in a Distributed Reactive Medium

From the Squared Bessel Description of Local Time to Survival in Partially Reactive Media

T. Zamrik 23-SEP-2026

Abstract. A particle diffusing through a reactive medium reacts not on first contact but when the contact it has accumulated crosses a threshold, so the quantity that decides its fate is a random field: the contact accumulated at every level it can react at. For one-dimensional Brownian motion that field has an exact law — the Ray-Knight theorems identify it, read in the space variable, as a squared Bessel process — and this paper applies that law to the functional physical reaction theory measures. The result is that the survival probability in an arbitrary reactivity profile $\hat{v}$ is the reciprocal of $W(a)$, where $W'' = 2\hat{v}W$ with $W(0) = 1$ and $W'(0) = 0$: one linear equation with two initial conditions, the reactivity entering as a potential, with no eigenvalue problem and no geometry remaining.

Two cases fix the result. A medium of finite width w and reduced reactivity $\hat{v}_0$ has the closed form $S = [\cosh kw + k \sinh kw\,(a - w)]^{-1}$ with $k = \sqrt{2\hat{v}_0}$; a reactive point of the same total strength has $S = (1+2\hat{\kappa}a)^{-1}$, which recovers Collins-Kimball kinetics and the Robin boundary condition. Between them, a medium of finite width behaves as a point whose reactivity is reduced by the factor $1-w/2a$ to leading order, a deficit set by the geometry of the search and independent of the strength of the chemistry. At the parameters used throughout, simulation gives 0.67569 ± 0.00113 against an exact 0.676195 for the medium and 0.66681 ± 0.00117 against an exact $2/3$ for the point, both scored on one batch of paths so that the gap between them is resolvable.

Because the same equation returns the transform at every argument, the whole law of the reacted contact is available, and with it the kinetics of a surface whose reaction threshold is not exponential — which no local boundary condition can express.

1. Introduction

A molecule diffusing towards a surface does not react the moment it arrives. It arrives, fails to react, leaves, returns, and reacts on some later visit or not at all. Half a century of chemical kinetics has treated that fact by softening the boundary condition: instead of absorbing every arrival, the surface absorbs a fraction of them, and the Smoluchowski problem becomes the partially reactive one that Collins and Kimball wrote down in 1949 [7]. The device works, and the reactivity constant it introduces is measurable, but it hides the quantity that actually decides the outcome.

That quantity is *contact*: not whether the particle has touched the surface, but how much touching it has accumulated. The encounter-based description makes this explicit [8]. A

particle carries a threshold drawn from some law, accumulates contact as it diffuses, and reacts at the instant its accumulated contact crosses the threshold. With an exponential threshold the description reproduces the partially reactive boundary condition exactly; with any other threshold law it produces kinetics no boundary condition can express. The central object is therefore not the trajectory and not the first arrival, but a random field: the contact the particle has accumulated at each place it could react.

For a surface, contact means boundary local time, and its law at a single site is classical. But a great many reactive environments are not surfaces. A reaction zone has a thickness. A trap layer coats a wall over some depth. Reactive sites sit distributed along a channel or a strand rather than concentrated at one place. In all of these the particle reacts wherever the medium is reactive, weighted by how reactive it is there, and the observable is not the contact at one point but a weighted integral of the contact over everywhere the medium reaches.

The step from a point to a medium changes what must be known. A single site needs one random variable; a medium needs the joint law of the contact at every level at once, as a process in the space variable. That object has a name and an exact description in probability, and the description is not an approximation, an asymptotic regime, or a simulation: the contact field of one-dimensional Brownian motion, read as a function of the level rather than of time, is a squared Bessel process. Confinement and geometry have long been understood to control diffusion-limited kinetics [14]; what the field law adds is that in one dimension the control is exactly computable, for any reactivity profile, from a single ordinary differential equation.

Two literatures own half of this problem each, and they do not meet. On one side, Ray [1] and Knight [2] described the law of Brownian local time in the space variable in 1963, and the theorems that carry their names are standard equipment in probability: they underlie the analysis of random walks in random environments, of branching processes, and of the Wiener sausage asymptotics of Donsker and Varadhan [12]. On the other, the encounter-based approach to diffusion-mediated surface phenomena has, since 2019, made the boundary local time the central object of physical reaction kinetics, with its distribution computed domain by domain [9].

The first literature has the law of the whole field and rarely asks what it is for. The second needs exactly that law and computes what it needs by simulation or by spectral expansion, one geometry at a time. Neither cites the other. This paper is the join: it takes the Ray-Knight description as given, applies it to the functional that physical reaction theory actually measures, and shows that for a one-dimensional medium of arbitrary reactivity profile the answer collapses to a Sturm-Liouville problem with two initial conditions and no eigenvalue to find.

Three things are established here. First, that the survival probability of a diffusing particle in a reactive medium of arbitrary profile, run until it first reaches a distant level, is the reciprocal of the solution at that level of a linear second-order equation driven by the reactivity, solved forwards from the particle's own starting point with a flat initial slope. The reactivity enters as a potential; no eigenvalue problem appears, and no geometry

beyond the profile itself is needed.

Second, that the formula is exact rather than asymptotic, and that it specialises correctly at both ends of its range: a medium of finite width has a closed form in hyperbolic functions, and as the width shrinks at fixed total reactivity that closed form converges to the classical point-reactivity survival, with a leading correction that depends on the width and the distance to the target but not on how reactive the medium is.

Third, that the field description delivers what the boundary condition cannot. Because the whole law of the contact is available and not merely its exponential transform, the same machinery answers questions a Robin condition has no way to pose: how much contact the survivors accumulated, how the contact is distributed across the medium, and what the kinetics become when the reaction threshold is not exponential.

Section 2 fixes the particle, the contact field, and the two stopping rules under which the field has an exact law, and records the scaling that lets a physical diffusion coefficient be carried in a single reduced reactivity. Section 3 states the two Ray-Knight theorems in the form used here and collects the squared Bessel facts the computation needs. Section 4 is the hinge: reaction is written as a functional of the field, and the functional is reduced to a Sturm-Liouville problem, with an algorithm that solves it for any profile.

Sections 5 and 6 are the two instantiations. The first is a medium of finite width and finite reactivity, for which the Sturm-Liouville problem is solved in closed form and the effective point reactivity of the medium is identified. The second is the point-reactivity limit, where the machinery must reproduce Collins-Kimball kinetics and the Robin boundary condition, and does. Section 7 scores both on one simulated field and exhibits the crossover between them. Section 8 is the numerical method, including an exactly computable statement of the bias of the occupation-density estimator, which is what makes the verification quantitative rather than visual. Section 9 discusses what fails in higher dimensions and what an experimentalist measures.

2. The Diffusing Particle and Its Contact

Definition 2.1 (The particle and the medium)**.** A point particle diffuses on the line with diffusion coefficient D, so that its position X_t is the Brownian motion with generator $D\,\partial_x^2$ and $\mathbb{E}[(X_t - X_0)^2] = 2Dt$. It starts at the origin, $X_0 = 0$. The medium it moves through is reactive: a non-negative, bounded, compactly supported function v on the line gives the reaction rate per unit time the particle suffers while at x.

Throughout, B denotes standard Brownian motion, the case $D = \frac{1}{2}$, whose generator is $\frac{1}{2}\,\partial_x^2$; the physical particle is $X_t = \sqrt{2D}\,B_t$. Every theorem below is stated for B, where the classical results are stated, and carried to X by the scaling recorded in the notation below. Nothing in the analysis requires v to be continuous or even piecewise continuous; boundedness and compact support are used only to keep every integral finite and every solution of the associated differential equation classical. The one-dimensional geometry is not a simplification of a three-dimensional problem but the geometry of the problems this description is sharpest for: channels, pores, strands, and layered media, where first-passage

kinetics are genuinely one-dimensional [13].

Definition 2.2 (The contact field)**.** The contact field of the particle at time t is the occupation density L_t^x, the density with respect to Lebesgue measure of the time the particle has spent at each level:

$$\int_0^t f(X_s)\,\mathrm{d}s \;=\; \int_{\mathbb{R}} f(x)\,L_t^x\,\mathrm{d}x \qquad \text{for every bounded measurable } f.$$

Equivalently $L_t^x = \lim_{\varepsilon\downarrow 0}(2\varepsilon)^{-1}\big|\{s \le t : |X_s - x| < \varepsilon\}\big|$, the limit existing almost surely and simultaneously in x.

The field is the paper's subject and the word *contact* is used for it throughout, because the quantity that governs reaction is not a duration but a density: time spent per unit of level. It has the dimensions of length divided by velocity, and it is finite and jointly continuous in (x,t) for Brownian motion, which is Trotter's theorem [3]. Read at a fixed time as a function of x it is a random function on the line; read at a fixed level as a function of t it is a continuous increasing process that grows only when the particle is at that level. Both readings are used below, and the second is the one Tanaka's formula makes computable.

Definition 2.3 (The two stopping rules)**.** For a level $a > 0$, the first hitting time is

$$T_a \;=\; \inf\{t \ge 0 : X_t = a\},$$

which is almost surely finite. For a contact level $\ell > 0$, the inverse contact time at the origin is

$$\tau_\ell \;=\; \inf\{t \ge 0 : L_t^0 > \ell\},$$

the first time the particle has accumulated contact ℓ at its starting level; it too is almost surely finite.

These two rules are not conveniences. The contact field has an exact and simple law at each of them and at no other stopping time of comparable generality: that is the content of the two Ray-Knight theorems. At a deterministic time the field is described by neither, and the difference matters for what this paper can and cannot compute. A reaction experiment run to a fixed clock time is outside the scope of the exact results below; one run until the particle escapes to a known distance, or until it has spent a prescribed amount of contact at a reference level, is inside it. The first rule is the one used for both instantiations; the second appears in the Ray-Knight section and is verified in the numerical method.

Notation 2.4 (Reduced reactivity)**.** Because $X_t = \sqrt{2D}\,B_t$, a reaction rate v suffered in physical time becomes, for the standard Brownian motion that carries the classical theorems, the reduced reactivity

$$\hat{v}(x) \;=\; \frac{v(x)}{2D},$$

and a point reactivity κ becomes $\hat{\kappa} = \kappa/(2D)$. The reduced quantities carry units of inverse length squared and inverse length respectively.

Every formula below is written in reduced variables, so that the diffusion coefficient never appears twice. To restore physical units, replace $\hat{v}$ by $v/(2D)$ and $\hat{\kappa}$ by $\kappa/(2D)$; with $D = \frac{1}{2}$ the reduced and physical reactivities coincide, which is why the probabilistic literature can suppress D altogether. The house parameters used throughout are $D = 1$, target level $a = 1$, medium width $w = 0.1$ and medium reactivity $v_0 = 5$, so that the total reactivity is $\kappa = v_0 w = 0.5$, and in reduced form $\hat{v}_0 = 2.5$ and $\hat{\kappa} = 0.25$. They are chosen so that survival is near one half, where the numerical comparisons below have the most resolution.

Assumption 2.5 (Standing hypotheses)**.** Throughout: the particle is the one-dimensional Brownian motion defined above, started at the origin; the reactivity $\hat{v}$ is non-negative, bounded and supported in $[0, a]$; the reaction threshold is a unit-mean exponential random variable independent of the path, except where the remark on non-exponential thresholds says otherwise; and the observation is stopped at T_a, except in the Ray-Knight section and again in the numerical method where τ_ℓ is used instead.

The support condition is the only one that does real work, and it is worth saying what it buys. Because $\hat{v}$ vanishes below the origin, an excursion of the particle into the negative half-line accumulates no reacted contact. Such excursions still occur, still last, and still contribute to the contact field at negative levels, but they are invisible to every observable in this paper. That is what makes the reflecting barrier of the numerical method a legitimate computational device rather than a change of model, and it is also why the results below involve the distance a to the target but no lower cut-off: there is none to involve.

The contact field has two constructions and they agree; the agreement is what makes it computable. Start from Tanaka's formula, which for standard Brownian motion and any level x states

$$|B_t - x| \;=\; |x| \;+\; \int_0^t \operatorname{sgn}(B_s - x)\,\mathrm{d}B_s \;+\; L_t^x, \tag{2.1}$$

so that the field is the compensator that turns the absolute value of a martingale into a semimartingale [3]. Rearranged, this identifies L_t^x as an increasing adapted process that grows only on $\{s : B_s = x\}$, which is a set of Lebesgue measure zero — the field measures a set the clock cannot see.

The occupation-density construction gives a second object. That the two coincide is the occupation-times formula: applying the Itô-Tanaka formula to a general convex function f and comparing second derivatives yields

$$\int_0^t g(B_s)\,\mathrm{d}s \;=\; \int_{\mathbb{R}} g(x)\,L_t^x\,\mathrm{d}x \tag{2.2}$$

for bounded measurable g, which is the occupation-density definition exactly. The consequence used everywhere below is the identity

$$\int_0^t \hat{v}(B_s)\,\mathrm{d}s \;=\; \int_{\mathbb{R}} \hat{v}(x)\,L_t^x\,\mathrm{d}x, \tag{2.3}$$

whose left side is a Feynman-Kac exponent and whose right side is a linear functional of

the field. The same number, read two ways: this paper is the consequences of reading it the second way.

Remark 2.6 (Why a point is not enough). If the reactive region is a single level, the reacted-contact functional collapses to $\hat{\kappa}L^{x_0}$ and only one random variable is needed. Its law for Brownian motion is classical and elementary, and everything the partially reactive boundary condition says follows from it.

The moment the reactive region has extent, that economy is gone. Two levels inside a medium do not carry independent contacts, and they do not carry perfectly dependent ones either: the field is a genuinely random function of the level, and the observable is an integral against it. Knowing each marginal separately is not enough to compute the expectation of an exponential of the integral, because that expectation depends on the joint law across all levels at once. This is the precise sense in which a reactive medium is harder than a reactive surface, and the precise reason the Ray-Knight theorems are the right tool: they give the joint law, as a Markov process in the level, and a Markov process in the level is exactly what makes an integral against the field computable by an ordinary differential equation.

3. The Ray-Knight Description

Theorem 3.1 (First Ray-Knight theorem)**.** *Let B be standard Brownian motion started at the origin, let $a > 0$ and let T_a be its first hitting time of a. Under the standing assumption, the contact field read downwards from the target,*

$$Z_x \;=\; L^{a-x}_{T_a}, \qquad 0 \le x \le a,$$

is a squared Bessel process of dimension two started at $Z_0 = 0$: it solves

$$\mathrm{d}Z_x \;=\; 2\,\mathrm{d}x \;+\; 2\sqrt{Z_x}\,\mathrm{d}\beta_x$$

for a Brownian motion β in the level variable. Continued below the origin, $x > a$, the same process is a squared Bessel process of dimension zero started at $Z_a = L^0_{T_a}$ and absorbed at 0.

Proof. The theorem is Ray [1] and Knight [2]; the argument in the form used here is Revuz and Yor, Chapter XI [3]. Its structure is worth recording because the two dimensions come from two different mechanisms. Above the starting level the particle must cross every intermediate level on its way to a, so the field is strictly positive on $(0, a)$ and the drift $2\,\mathrm{d}x$ is the deterministic rate at which contact is manufactured by the crossings that the target forces. Below the starting level no crossing is forced: the field is a martingale, the drift vanishes, and the dimension drops to zero, which is why the field dies at a finite depth. The diffusion coefficient $2\sqrt{Z}$ in both regimes is the Itô-Tanaka computation applied in the level variable rather than in time, the excursions straddling a level supplying the quadratic variation. Markovianity in x — the fact that the field below a level depends on the field above it only through its value at that level — is the excursion-theoretic content, and it

is what the next sections consume. □

Theorem 3.2 (Second Ray-Knight theorem)**.** *Let B be standard Brownian motion started at the origin and let τ_ℓ be the inverse contact time defined above. Under the standing assumption, the contact field read upwards from the origin,*

$$Z_x \;=\; L^x_{\tau_\ell}, \qquad x \geq 0,$$

is a squared Bessel process of dimension zero started at $Z_0 = \ell$ and absorbed at 0:

$$\mathrm{d}Z_x \;=\; 2\sqrt{Z_x}\,\mathrm{d}\beta_x.$$

In particular $\mathbb{E}[L^x_{\tau_\ell}] = \ell$ for every $x \geq 0$, and the field is extinguished above a finite level almost surely.

Proof. Knight [2]; see Revuz and Yor, Chapter XI [3]. Dimension zero here has the same origin as it does below the starting level in the first Ray-Knight theorem: stopping at τ_ℓ imposes no requirement that the particle reach any level above the origin, so no crossings are forced, no contact is manufactured, and the field is a non-negative local martingale in the level variable. A non-negative local martingale is a supermartingale, and the absorbing boundary at zero is what makes it a true martingale here with constant mean ℓ; the field is not merely decreasing on average, it is exactly conserved on average while being extinguished path by path. The two statements coexist because extinction is compensated by the growing right tail of the survivors, and the numerical method measures both. □

Lemma 3.3 (Squared Bessel laws)**.** *Let Z be a squared Bessel process of dimension $\delta \geq 0$ started at $z \geq 0$. Then for $\lambda \geq 0$ and $x > 0$,*

$$\mathbb{E}_z\big[e^{-\lambda Z_x}\big] \;=\; (1+2\lambda x)^{-\delta/2}\exp\Big(-\frac{\lambda z}{1+2\lambda x}\Big),$$

from which $\mathbb{E}_z[Z_x] = z + \delta x$. Two specialisations are used below: for $\delta = 2$ and $z = 0$, Z_x is exponentially distributed with mean $2x$; for $\delta = 0$, $\mathbb{P}(Z_x = 0) = e^{-z/(2x)}$. Squared Bessel processes are additive: the sum of independent copies of dimensions δ, δ' started at z, z' is one of dimension $\delta + \delta'$ started at $z + z'$.

Proof. The Laplace transform is the classical computation of Pitman and Yor [4]; see also the survey of Going-Jaeschke and Yor [5]. It may be read off the Riccati equation of Derivation 3 with a constant potential. The mean follows by differentiating at $\lambda = 0$. For $\delta = 2$, $z = 0$ the transform is $(1+2\lambda x)^{-1}$, the transform of the exponential law of mean $2x$. For $\delta = 0$ the transform is $\exp(-\lambda z/(1+2\lambda x))$, whose limit as $\lambda \to \infty$ is $e^{-z/(2x)}$, and that limit is the mass at the origin. Additivity is Shiga and Watanabe's theorem [4], and is visible in the transform: the exponents add. □

Combining the squared Bessel laws with the first Ray-Knight theorem gives the marginal law of the contact field that every verification in Section 8 is scored against: for $0 \leq y \leq a$,

$$L^{y}_{T_a} \;\sim\; \mathrm{Exp}(\text{mean } 2(a-y)), \qquad \text{so} \qquad \mathbb{E}\big[L^{y}_{T_a}\big] \;=\; 2(a-y). \tag{3.1}$$

At the starting level itself this is the statement that $L^{0}_{T_a}$ is exponential with mean $2a$, a fact that can be had directly from excursion theory and which will do most of the work in the point-reactivity section.

The physical particle $X_t = \sqrt{2D}\,B_t$ has a contact field of its own, and it is worth seeing that carrying the diffusion coefficient costs exactly one substitution and no rescaling of the geometry. Start from the occupation-density definition applied to X and change variables:

$$\int_0^t f(X_s)\,\mathrm{d}s \;=\; \int_0^t f(\sqrt{2D}\,B_s)\mathrm{d}s \;=\; \int_{\mathbb{R}} f(\sqrt{2D}\,y)L^{y}_t(B)\,\mathrm{d}y \;=\; \int_{\mathbb{R}} f(u)\,\frac{L^{u/\sqrt{2D}}_t(B)}{\sqrt{2D}}\,\mathrm{d}u, \tag{3.2}$$

so the two fields are related by $L^{u}_t(X) = (2D)^{-1/2}L^{u/\sqrt{2D}}_t(B)$: the field is stretched in the level variable and scaled in amplitude. Under that map the squared Bessel description survives intact, because $c\,Z_{x/c}$ is again a squared Bessel process of the same dimension whenever Z is one — the drift $\delta\,\mathrm{d}x$ and the diffusion $2\sqrt{Z}\,\mathrm{d}\beta$ scale in the same way.

The reacted contact is therefore

$$\int_0^t v(X_s)\,\mathrm{d}s \;=\; \int_{\mathbb{R}} v(u)\,L^{u}_t(X)\,\mathrm{d}u \;=\; \int_{\mathbb{R}} v(\sqrt{2D}\,y)L^{y}_t(B)\,\mathrm{d}y, \tag{3.3}$$

with the target level a for X corresponding to $a/\sqrt{2D}$ for B. Both changes are undone at once by working in the physical level variable with the reduced reactivity $\hat{v} = v/(2D)$ introduced above: every result below is stated in the physical coordinate x, with the physical distance a, and the diffusion coefficient appears only inside $\hat{v}$.

Remark 3.4 (The field is a population). A squared Bessel process of dimension zero is Feller's branching diffusion, the continuous-state limit of a critical Galton-Watson population. The second Ray-Knight theorem therefore says something stronger than a formula: the contact field, read upwards in the level variable, *is* a critically branching population whose size at level x is the contact accumulated there, and whose extinction level is the highest level the particle ever reached.

The reading is not decorative. It explains the pair of facts that look contradictory in that theorem — that the field is conserved in mean yet dies almost surely — because that is precisely the behaviour of a critical population, and it identifies $\mathbb{P}(Z_x = 0) = e^{-z/(2x)}$ as an extinction probability rather than as an integration constant. It also explains why dimension two appears below a hitting level in the first Ray-Knight theorem: the requirement that the particle reach a is an immigration of rate two into the population, and immigration of exactly that rate is what keeps the field from dying before the target is reached. The branching reading is the standard one in the probabilistic literature [5], and it supplies the intuition for the only qualitative claim this paper makes about the shape of a contact profile.

Example 3.5 (Reading one field)**.** A single path is simulated at lattice spacing $h = 0.01$

with $a = 1$, and its contact is read off at five levels. The realization gives $L^{\,0} = 0.40$, $L^{\,0.25} = 0.15$, $L^{\,0.5} = 0.13$, $L^{\,0.75} = 0.80$ and $L^{\,0.9} = 0.21$, and the field vanishes above the level 1, which this path reached only at the moment it stopped.

Against the marginal law the squared Bessel description supplies — exponential with mean $2(a-y)$ — those five numbers sit at the cumulative probabilities 0.181, 0.095, 0.122, 0.798 and 0.650 respectively. This single path is therefore a low draw near its starting level and a high one two thirds of the way to the target, which is not a defect of the simulation but the size of the fluctuations the theory predicts: at every level the contact is exponentially distributed, so its standard deviation equals its mean and realizations an order of magnitude apart are ordinary. It is worth seeing once, because every averaged quantity below is an average over objects that look like this, and the smoothness of the mean profile $2(a-y)$ is a property of the ensemble and of no individual path.

4. Reaction as a Functional of the Field

Definition 4.1 (Reacted contact and the encounter rule)**.** The medium's reactivity is the profile $v(x)$ of Definition 1, in reduced form $\hat{v}$. The *reacted contact* accumulated by time t is the linear functional of the field

$$q_t \;=\; \int_{\mathbb{R}} \hat{v}(x)\, L_t^x \,\mathrm{d}x \;=\; \int_0^t \hat{v}(X_s)\,\mathrm{d}s.$$

The particle carries a *reaction threshold* $\Theta > 0$, drawn independently of the path, and reacts at the first time q_t exceeds Θ. Survival to a stopping time T is the event $\{q_T \le \Theta\}$, and the survival probability is

$$S \;=\; \mathbb{P}(q_T \le \Theta) \;=\; \mathbb{E}\big[\Psi(q_T)\big], \qquad \Psi(u) = \mathbb{P}(\Theta > u).$$

Under the standing assumption Θ is exponential of unit mean, so $\Psi(u) = e^{-u}$ and the survival probability is the Laplace transform $\mathbb{E}[e^{-q_T}]$ evaluated at one. That choice is what makes the description equivalent to a reaction rate, and it is the only choice for which the equivalence holds: an exponential threshold is memoryless, so the particle's history of failed encounters does not change its future prospects, which is exactly the assumption a rate constant encodes. Every other threshold law describes a surface with memory.

Theorem 4.2 (Survival is the reciprocal of a Sturm-Liouville solution)**.** *Let the particle start at the origin and be observed until T_a, and let the reduced reactivity $\hat{v}$ be supported in $[0,a]$. Let W solve the initial-value problem*

$$\begin{aligned} W''(y) \;&=\; 2\,\hat{v}(y)\,W(y) \quad \text{on } [0,a], \\ W(0) \;&=\; 1, \qquad W'(0) \;=\; 0. \end{aligned}$$

Then the survival probability is

$$S \;=\; \mathbb{E}\Big[\exp\Big(-\int_0^a \hat{v}(y) L_{T_a}^y \,\mathrm{d}y\Big)\Big] \;=\; \frac{1}{W(a)}.$$

More generally, for every $\lambda \geq 0$, $\mathbb{E}[e^{-\lambda q_{T_a}}] = 1/W_\lambda(a)$ where W_λ solves the same problem with $\hat{v}$ replaced by $\lambda\hat{v}$; the law of the reacted contact is therefore determined by the same equation.

Proof. By the first Ray-Knight theorem the field read downwards from the target, $Z_x = L^{a-x}_{T_a}$, is a squared Bessel process of dimension two started at zero, so with $q(x) = \hat{v}(a-x)$ the exponent is $\int_0^a q(x)Z_x\,\mathrm{d}x$ and the Riccati computation above applies with $\delta = 2$ and $z = 0$. The factor $\exp(-\alpha(0)z)$ is one because the process starts at zero, leaving $S = w(0)^{-1}$ with $w'' = 2qw$, $w(a) = 1$, $w'(a) = 0$. Setting $W(y) = w(a-y)$ converts the terminal-value problem into the stated initial-value one, with $W(a) = w(0)$, which is the claim. The general λ follows by replacing $\hat{v}$ with $\lambda\hat{v}$ throughout. □

Proof. The same number is a Feynman-Kac exponent. The function

$$u(x) \;=\; \mathbb{E}_x\Big[\exp\Big(-\int_0^{T_a} \hat{v}(B_s)\,\mathrm{d}s\Big)\Big]$$

solves $u'' = 2\hat{v}u$ on $(-\infty, a)$ with $u(a) = 1$ and u bounded. Below the origin $\hat{v}$ vanishes, so u is affine there, and boundedness forces it to be constant, whence $u'(0) = 0$. On $[0, a]$ the function u therefore solves the same equation as W with the same initial slope, so $u = u(0)W$, and $u(a) = 1$ gives $u(0) = 1/W(a)$. The two routes agree, which is not a coincidence but a consequence of the occupation-times formula; the Ray-Knight route is nonetheless the one that survives the generalisation of the next remark. □

The functional to be computed is an exponential of a linear functional of a squared Bessel process, and such functionals are governed by a Riccati equation because the process is affine. Let Z be a squared Bessel process of dimension δ and let $q \geq 0$ be bounded on $[0, t]$. Write

$$F(s, \zeta) \;=\; \mathbb{E}\Big[\exp\Big(-\int_s^t q(x)Z_x\,\mathrm{d}x\Big)\;\Big|\; Z_s = \zeta\Big], \tag{4.1}$$

which by the Markov property in the level variable satisfies the backward equation

$$\partial_s F \;+\; \delta\,\partial_\zeta F \;+\; 2\zeta\,\partial^2_\zeta F \;-\; q(s)\,\zeta F \;=\; 0, \qquad F(t, \zeta) = 1. \tag{4.2}$$

The generator is affine in ζ, so try $F = \exp(-\alpha(s)\zeta - \beta(s))$. Substituting and separating the terms in ζ from the constants gives two ordinary equations,

$$\alpha' \;=\; 2\alpha^2 - q, \qquad \beta' \;=\; -\delta\,\alpha, \qquad \alpha(t) = \beta(t) = 0. \tag{4.3}$$

The first is a Riccati equation and is linearised by $\alpha = -\frac{1}{2}\,w'/w$, which turns it into

$$w'' \;=\; 2\,q\,w. \tag{4.4}$$

The terminal conditions become $w(t) = 1$ and $w'(t) = 0$, and the second equation integrates to $\beta(s) = \frac{\delta}{2}\ln w(s)$. Therefore

$$\mathbb{E}_z\Big[\exp\Big(-\int_0^t q(x)Z_x\,\mathrm{d}x\Big)\Big] \;=\; w(0)^{-\delta/2}\,\exp(-\alpha(0)\,z), \qquad \alpha(0) = -\tfrac{1}{2}\,\frac{w'(0)}{w(0)}. \tag{4.5}$$

A linear second-order equation with two initial conditions: no eigenvalue is sought, no boundary value is matched, and the reactivity enters as a potential. Setting q constant and δ arbitrary reproduces the squared Bessel Laplace transform quoted above.

Algorithm 1 — Survival from an arbitrary reactivity profile

Inputs: the reduced profile $\hat{v}$ on $[0, a]$, the target a, and, for a non-exponential threshold, the survival function Ψ of Θ. *Output:* the survival probability S, and if wanted the whole law of the reacted contact.

1. Set $W = 1$ and $W' = 0$ at $y = 0$.
2. If $\hat{v}$ is piecewise constant, propagate exactly. Across a piece of length Δ with value $\hat{v}_i > 0$ and $k_i = \sqrt{2\hat{v}_i}$, the pair (W, W') is carried by

$$M_i = \begin{pmatrix} \cosh k_i\Delta & k_i^{-1}\sinh k_i\Delta \\ k_i\sinh k_i\Delta & \cosh k_i\Delta \end{pmatrix},$$

   and across a piece with $\hat{v}_i = 0$ by $M_i = \left(\begin{smallmatrix} 1 & \Delta \\ 0 & 1 \end{smallmatrix}\right)$. Otherwise integrate $W'' = 2\hat{v}W$ from 0 to a by any one-pass scheme; fourth-order Runge-Kutta is used throughout this paper.
3. Report $S = 1/W(a)$.
4. For a non-exponential threshold, repeat steps 1 to 3 with $\hat{v}$ replaced by $\lambda\hat{v}$ over a grid of λ, obtaining $\mathbb{E}[e^{-\lambda q_{T_a}}]$, and invert the transform numerically to get the law of q_{T_a}; then $S = \mathbb{E}[\Psi(q_{T_a})]$.

Cost. Step 2 is a single sweep: one 2×2 product per piece for a piecewise-constant profile, so a medium of m layers costs m matrix products and nothing else, and a general profile costs $O(N)$ for N integration steps. There is no eigenvalue problem, no boundary value to match by shooting, no matrix to invert and no iteration to converge, which is the practical content of the theorem: the geometry has been removed and only a potential is left.

Remark 4.3 (Thresholds that are not exponential). The partially reactive boundary condition can express only the memoryless case. If the reaction threshold Θ has any other law, the survival probability $\mathbb{E}[\Psi(q_T)]$ is not the Laplace transform of the reacted contact at a single point, and no local boundary condition reproduces it.

This is where the field description stops being an alternative derivation of known results and starts being necessary. The Sturm-Liouville theorem returns the transform at every λ, which is the whole law of q_T; from the whole law any threshold law can be scored. Physically the non-exponential case is the interesting one: a surface that must be struck several times before it yields, a site that is blocked for a while after a failed encounter, a receptor that must accumulate a dose. Grebenkov's encounter-based formulation makes exactly this point for a reactive boundary [8], and the same argument carries to a reactive

medium with the profile in place of the surface. What is added here is that in one dimension the transform is not a spectral sum to be truncated but the reciprocal of one solution of one ordinary differential equation, so the inversion can be done at whatever accuracy the threshold law deserves.

The Sturm-Liouville reduction is checked twice: against the closed form it must reproduce, and against the paths it is supposed to describe.

Method and parameters. The initial-value problem $W'' = 2\hat{v}W$, $W(0) = 1$, $W'(0) = 0$ is integrated from 0 to a by fourth-order Runge-Kutta on 40,001 points, with the house profile $\hat{v}_0 = 2.5$ on $[0, 0.1)$ and $a = 1$. The same quantity is estimated from 40,000 simulated paths at lattice spacing $h = 0.005$, each scored by the reacted contact $\hat{v}_0 \times$ (time in the medium).

Result. The integrator returns $W(a) = 1.478844$ and hence $S = 0.676204$, against the closed form 0.676195 of the slab theorem: a relative error of 1.3×10^{-5}, which is the integrator's own discretisation and not a property of the theorem. The simulation returns $\widehat{S} = 0.67569 \pm 0.00113$, a deviation of -0.45 standard errors from the exact value.

Error. The two independent failures the theorem could have — being the wrong equation, or being the right equation solved wrongly — are separated by this pair. Agreement with the closed form to five decimal places rules out the second; agreement with simulation to within half a standard error rules out the first. Nothing here tests the general-λ statement, which the crossover section exercises at four different profiles.

5. Instantiation I - a Medium Distributed over an Interval

Definition 5.1 (The slab)**.** The *slab* is the reactivity profile that is constant on an interval of width w starting at the particle's own position and zero elsewhere:

$$\hat{v}(y) \;=\; \hat{v}_0\, \mathbf{1}_{[0,w)}(y), \qquad 0 < w \le a.$$

Its *total reactivity* is $\hat{\kappa} = \hat{v}_0 w$, the integral of the profile, and its *wavenumber* is $k = \sqrt{2\hat{v}_0}$. The physical profile is v_0 on the same interval, with $\hat{v}_0 = v_0/(2D)$.

Two parameters describe the slab and only their product is visible to a point description: a thin, fiercely reactive layer and a broad, weakly reactive one can share a total reactivity $\hat{\kappa}$ while differing in w by orders of magnitude. The question this section answers is what the difference is worth. The slab is placed at the starting level rather than somewhere in the interior because that is the configuration in which the point limit is the classical one, and because it is the geometry of a coated wall, a reaction zone at an interface, or a trap layer at the mouth of a pore, which are the physical situations the one-dimensional description fits.

Lemma 5.2 (Constant reactivity)**.** *On an interval of length Δ on which the reduced reactivity takes the constant value $\hat{v}_i > 0$, the solution of $W'' = 2\hat{v}W$ is carried by the*

transfer matrix

$$\begin{pmatrix} W \\ W' \end{pmatrix}_{y+\Delta} = \begin{pmatrix} \cosh k_i\Delta & k_i^{-1}\sinh k_i\Delta \\ k_i\sinh k_i\Delta & \cosh k_i\Delta \end{pmatrix} \begin{pmatrix} W \\ W' \end{pmatrix}_{y}, \qquad k_i = \sqrt{2\hat{v}_i},$$

and on an interval with $\hat{v}_i = 0$ by the same matrix with $k_i \to 0$, namely $\left(\begin{smallmatrix} 1 & \Delta \\ 0 & 1 \end{smallmatrix}\right)$. Each matrix has unit determinant.

Proof. With $\hat{v}$ constant the equation is $W'' = k^2 W$, whose solutions are spanned by $\cosh ky$ and $\sinh ky$; matching W and W' at the left endpoint gives the stated matrix, and its determinant is $\cosh^2 - \sinh^2 = 1$. The limit $k \to 0$ is the free case, where $W'' = 0$ and the pair propagates affinely. Unit determinant is the Wronskian of the equation, which is constant because the equation has no first-derivative term; it is the reason a layered medium can be composed by multiplying matrices without any normalisation. □

Theorem 5.3 (Survival against a slab)**.** *For the slab of width w and reduced reactivity $\hat{v}_0$, with $k = \sqrt{2\hat{v}_0}$, the survival probability to the target level a is*

$$S \;=\; \frac{1}{\cosh kw \;+\; k\,\sinh kw\,(a-w)}.$$

It is decreasing in $\hat{v}_0$ and in a, increasing in w at fixed total reactivity $\hat{\kappa} = \hat{v}_0 w$, and tends to one as either $\hat{v}_0$ or w tends to zero.

Proof. Apply the Sturm-Liouville theorem with the profile of the slab definition. Starting from $W(0) = 1$, $W'(0) = 0$, the reactive interval contributes the hyperbolic transfer matrix of the constant-reactivity lemma with $\Delta = w$, giving $W(w) = \cosh kw$ and $W'(w) = k\sinh kw$. The remaining interval $[w, a]$ carries no reactivity, so the free matrix applies with $\Delta = a - w$, giving

$$W(a) \;=\; \cosh kw \;+\; k\,\sinh kw\,(a-w),$$

and the survival probability is its reciprocal. The monotonicity in $\hat{v}_0$ and in a is read off directly, since every term is positive and increasing. For the behaviour at fixed $\hat{\kappa}$, substitute $\hat{v}_0 = \hat{\kappa}/w$ and differentiate; the claim also follows from the effective-reactivity proposition below, which computes the sign of the correction. □

The formula has the structure the problem should have. The reactive layer contributes through kw, which is the layer's own optical thickness, and the inert gap contributes through $a - w$ multiplied by the slope the layer imparts. A layer that is thin in its own units, $kw \ll 1$, imparts a slope $k\sinh kw \approx k^2 w = 2\hat{\kappa}$ and its internal structure disappears; a layer that is thick absorbs the particle within itself and the distance to the target stops mattering.

Corollary 5.4 (Thin-slab limit)**.** *Fix the total reactivity $\hat{\kappa} = \hat{v}_0 w$ and let $w \to 0$. Then $kw = \sqrt{2\hat{\kappa}w} \to 0$ and*

$$S \;\longrightarrow\; \frac{1}{1 + 2\hat{\kappa}a},$$

which is the survival probability against a reactive point of the same total strength at the origin.

Proof. With $\hat{v}_0 = \hat{\kappa}/w$ the wavenumber is $k = \sqrt{2\hat{\kappa}/w}$, so $kw = \sqrt{2\hat{\kappa}w} \to 0$ and $\cosh kw \to 1$. The slope term is $k \sinh kw = k^2 w + O(k^4 w^3) = 2\hat{\kappa} + O(\hat{\kappa}^2 w)$, so $W(a) \to 1 + 2\hat{\kappa}a$ and the survival probability tends to its reciprocal. □

The limit is a consistency check with a classical answer rather than a new statement, and it is the reason the slab is anchored at the origin: a delta function of weight $\hat{\kappa}$ placed at the particle's starting level imposes a jump $2\hat{\kappa}$ on W' across the origin, from the flat initial slope, and the free propagation to the target does the rest.

Proposition 5.5 (The effective reactivity of a slab)**.** *Define the effective point reactivity $\hat{\kappa}_{\text{eff}}$ of a slab as the point reactivity with the same survival probability, so that $S = (1 + 2\hat{\kappa}_{\text{eff}}a)^{-1}$. Then*

$$\hat{\kappa}_{\text{eff}} \;=\; \frac{\cosh kw - 1 \;+\; k \sinh kw\,(a - w)}{2a},$$

and at fixed total reactivity $\hat{\kappa}$, as $w \to 0$,

$$\frac{\hat{\kappa}_{\text{eff}}}{\hat{\kappa}} \;=\; 1 \;-\; \frac{w}{2a} \;+\; O(w^2).$$

The leading correction depends on the width and on the distance to the target, and not at all on how reactive the medium is.

Proof. The definition rearranges to $2a\hat{\kappa}_{\text{eff}} = W(a) - 1$, which with the slab formula is the stated expression. For the expansion put $\hat{v}_0 = \hat{\kappa}/w$ and use $\cosh kw - 1 = \frac{1}{2}k^2w^2 + O(k^4w^4) = \hat{\kappa}w + O(\hat{\kappa}^2 w^2)$ together with $k \sinh kw = 2\hat{\kappa} + O(\hat{\kappa}^2 w)$. Then

$$2a\hat{\kappa}_{\text{eff}} \;=\; \hat{\kappa}w \;+\; 2\hat{\kappa}(a - w) \;+\; O(w^2) \;=\; 2\hat{\kappa}a \;-\; \hat{\kappa}w \;+\; O(w^2),$$

and dividing by $2a\hat{\kappa}$ gives the claim. The two $O(w^2)$ terms carry $\hat{\kappa}^2$, so the reactivity enters only at second order. □

Spread is therefore mild and it is always in the same direction: a medium of finite width is *less* effective than a point of the same total reactivity, because contact accumulated away from the starting level is contact the particle acquires less of. The size of the effect is set by the ratio of the medium's width to the distance the particle has to travel, and a medium ten times thinner than that distance is misdescribed by five per cent when it is replaced by a point.

Example 5.6 (The house slab)**.** Take $D = 1$, $a = 1$, medium width $w = 0.1$ and physical reactivity $v_0 = 5$, so that the reduced reactivity is $\hat{v}_0 = 2.5$, the total reactivity is $\hat{\kappa} = \hat{v}_0 w = 0.25$, and the wavenumber is $k = \sqrt{2\hat{v}_0} = 2.236068$. The slab is optically thin, $kw = 0.223607$.

The transfer matrix across the medium gives $\cosh kw = 1.025104$ and $k \sinh kw = 0.504177$, so the propagation to the target adds $0.504177 \times (1 - 0.1) = 0.453759$ and

$$W(a) \;=\; 1.025104 + 0.453759 \;=\; 1.478864, \qquad S \;=\; 0.676195. \tag{5.1}$$

The point of the same total reactivity gives $S = 1/(1+2\times 0.25\times 1) = 0.666667$, so the slab is 1.43 per cent more survivable than the point it would be mistaken for. In effective terms $\widehat{\kappa}_{\mathrm{eff}} = 0.239432$, which is 0.957727 of the total reactivity, while the first-order prediction of the effective-reactivity proposition is $1 - w/(2a) = 0.95$. The prediction is short by four parts in a thousand at a width one tenth of the distance to the target, which is the second-order term doing what it was said to do.

Remark 5.7 (What a slab is). The slab is not a mathematical convenience. A reaction zone at an interface has a thickness set by the chemistry rather than by the geometry; a trap layer coating a pore wall has a depth; a quencher dissolved near a surface has a concentration profile; and a stretch of a polymer or a nucleic acid carrying binding sites has a length.

In each case the physical question is the one the effective-reactivity proposition answers: when may the layer be replaced by a surface with a reactivity constant, and what is lost when it is. The answer here is that the replacement is good to first order in the ratio of the layer's width to the distance the particle must travel, independently of how reactive the layer is, and that it always overstates the reaction. The independence is the useful half: a layer's optical thickness kw can be large or small without changing the leading correction, so an experimentalist who knows only the geometry can bound the error without knowing the chemistry. The trapping problem in a medium with randomly placed traps is the same functional with a random profile, and the asymptotics of Donsker and Varadhan [12] are what that functional produces when the traps are Poissonian and the observation is long.

Method. Survival against the slab is estimated directly from simulated paths, with no use of the Ray-Knight description: each path accumulates its own reacted contact as the reduced reactivity times the time it spends in the medium, and the survival estimate is the mean of e^{-q} over paths. The estimate is therefore an independent test of the closed form rather than a rearrangement of it.

Parameters. $D = 1$, $a = 1$, $w = 0.1$, $\widehat{v}_0 = 2.5$, lattice spacing $h = 0.005$, 40,000 paths, reflecting depth $M = 1$.

Result. $\widehat{S} = 0.67569 \pm 0.00113$ against the exact 0.676195: a deviation of -0.45 standard errors, and a relative error of 7.5×10^{-4}. Repeating at three other widths, at fixed total reactivity $\widehat{\kappa} = 0.25$, gives 0.67110 ± 0.00164 against 0.671363 at $w = 0.05$, 0.68513 ± 0.00155 against 0.686286 at $w = 0.2$, and 0.70688 ± 0.00141 against 0.708327 at $w = 0.4$, deviations of -0.16, -0.75 and -1.03 standard errors respectively.

Error. All four deviations are within about one standard error and all four have the same sign, which is the lattice's own bias rather than a defect in the formula: the convergence table below shows that bias shrinking with the spacing. The closed form is therefore confirmed across an eightfold range of width, over which the survival probability itself moves by more than three per cent and the effective reactivity by sixteen.

Script. `fig_slab.py`. The left panel sweeps the medium's reactivity over five decades at

fixed width and plots survival against the optical thickness kw; four independently seeded batches of 6,000 paths sit on the curve. The right panel is exact arithmetic only: the error committed by describing the medium as a point, and the error the first-order correction leaves behind, both as functions of the width.

Shows. That survival against a medium is a crossover between two asymptotes, and that the point description fails linearly while its first-order repair saturates.

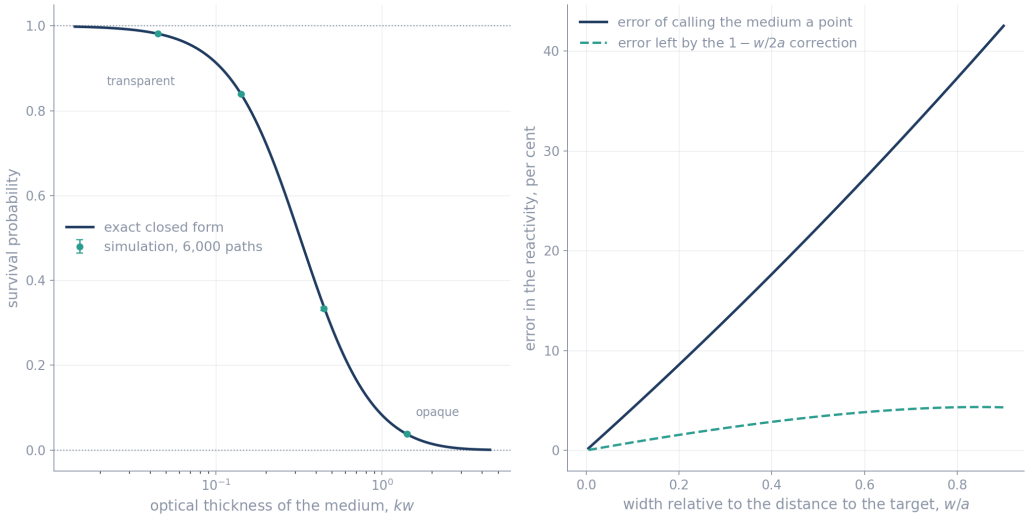

Figure 1: Survival against the optical thickness of the reactive medium across five decades, showing the transparent and opaque asymptotes, beside the error of describing that medium as a reactive point and the residual error left by the first-order correction.

The left panel is the closed form doing what a closed form is for: at kw below about 0.05 the medium is transparent and survival is within two per cent of one, at kw above about two it is opaque and nothing escapes, and the whole transition happens across a single decade in between. The house parameters sit at $kw = 0.224$, on the upper shoulder, which is where the comparison with a point has the most to resolve. The right panel says what that comparison costs: calling the medium a point overstates its reactivity by five per cent at a tenth of the distance to the target and by forty-five per cent at nine tenths, essentially linearly, while the residual after the $1 - w/2a$ correction rises to only four per cent and then flattens — the correction does almost all of the available work and stops improving once the expansion runs out.

Source. Simulated survival at four widths, each with its own independent batch of 20,000 paths at lattice spacing $h = 0.01$, against the closed form of the slab theorem. The total reactivity is held at $\hat{\kappa} = 0.25$ throughout, so the reduced reactivity $\hat{v}_0$ moves inversely with the width and every row describes the same amount of chemistry spread over a different depth.

Columns. The width; the reduced reactivity that keeps the total fixed; the simulated survival with its standard error; the exact survival; the deviation in standard errors; and the effective reactivity as a fraction of the total, against its first-order prediction.

w	$\hat{v}_0$	**simulated**	**exact**	**dev.**	$\hat{\kappa}_{\text{eff}}/\hat{\kappa}$	$1 - w/2a$
0.05	5.00	0.67110 ± 0.00164	0.671363	-0.16	0.9790	0.9750
0.10	2.50	0.67540 ± 0.00161	0.676195	-0.49	0.9577	0.9500
0.20	1.25	0.68513 ± 0.00155	0.686286	-0.75	0.9142	0.9000
0.40	0.62	0.70688 ± 0.00141	0.708327	-1.03	0.8236	0.8000

Survival rises by three and a half per cent across the range while the total reactivity is unchanged, which is the whole effect this paper measures. The first-order prediction tracks the exact effective reactivity to within half a per cent at $w = 0.05$ and to within two and a half per cent at $w = 0.4$, degrading as w^2 exactly as the proposition's error term says it should. All four deviations are negative, which is the lattice bias at $h = 0.01$ visible in the convergence table, not a failure of the closed form.

6. Instantiation II - the Point-Reactivity Limit

Definition 6.1 (Point reactivity)**.** A *reactive point* of strength κ at the origin is the profile $v = \kappa\,\delta_0$, in reduced form $\hat{\kappa}\,\delta_0$ with $\hat{\kappa} = \kappa/(2D)$. The reacted contact is then

$$q_t \;=\; \hat{\kappa}\,L_t^0,$$

a single random variable rather than an integral against a field. Equivalently, the particle is killed at a rate proportional to the growth of its contact at one level, which is the encounter description of a partially reactive site.

The delta function is legitimate here in the strongest sense: the contact field is jointly continuous, so L_t^0 is an honest random variable, and no regularisation is needed to make sense of the profile. In the Sturm-Liouville equation the delta imposes a jump on the slope, $W'(0^+) - W'(0^-) = 2\hat{\kappa}\,W(0)$, and with the flat initial slope supplied by the inert half-line below the origin this is $W'(0^+) = 2\hat{\kappa}$. Everything in this section is a consequence of that single jump condition and the affine propagation that follows it.

Theorem 6.2 (The point limit is Collins-Kimball kinetics)**.** *For a reactive point of reduced strength $\hat{\kappa}$ at the origin, observed until the particle first reaches a,*

$$S \;=\; \mathbb{E}\big[e^{-\hat{\kappa}L^0_{T_a}}\big] \;=\; \frac{1}{1+2\hat{\kappa}a}, \qquad \textit{equivalently} \qquad \frac{1}{S} \;=\; 1 \;+\; 2\hat{\kappa}a.$$

The reciprocal survival is affine in the reactivity with slope $2a$, and the two terms are the escape of an unreactive particle and the reactive resistance of the site.

Proof. Two routes, and they agree. By the first Ray-Knight theorem with the squared Bessel laws, the contact at the starting level, $L^0_{T_a} = Z_a$ with Z a squared Bessel process of dimension two from zero, is exponentially distributed with mean $2a$. The Laplace transform of an exponential law of mean $2a$ at $\hat{\kappa}$ is $(1 + 2\hat{\kappa}a)^{-1}$, which is the claim. Alternatively, take the thin-slab limit of the slab theorem at fixed total reactivity $\hat{\kappa}$; that computation is the preceding corollary and gives the same expression. □

The affine form is the structure Collins and Kimball identified in 1949 [7]: a diffusion-limited contribution and a reaction-limited contribution that add as resistances in series rather than compounding. Here the decomposition is exact and its two coefficients are explicit — the constant is the survival of a particle that never reacts, and the coefficient $2a$ is the mean contact the particle accumulates at the site before escaping. A reactive site is therefore described completely, for this observable, by the mean contact the geometry delivers to it.

Proposition 6.3 (The Robin condition is an exponential functional of contact)**.** *Let the particle be reflected at the origin rather than free, so that it occupies* $[0, a]$*, and let it be killed at rate* $\hat{\kappa}$ *per unit of contact at the origin. Then* $u(x) = \mathbb{E}_x[\exp(-\hat{\kappa}\, L^0_{T_a})]$ *is the solution of*

$$u'' = 0 \quad \text{on } (0, a), \qquad \tfrac{1}{2}\, u'(0) = \hat{\kappa}\, u(0), \qquad u(a) = 1,$$

namely $u(x) = (1 + 2\hat{\kappa}x)/(1 + 2\hat{\kappa}a)$*, and in particular* $u(0) = (1 + 2\hat{\kappa}a)^{-1}$*, the same value as for the free particle with a reactive point.*

Proof. With no reactivity in the interior the function is affine, $u = A + Bx$. Killing at rate $\hat{\kappa}$ per unit contact at the origin contributes to the generator only through the boundary, and the domain of the generator of reflected Brownian motion killed in this way is the set of functions satisfying $\frac{1}{2}u'(0) = \hat{\kappa}u(0)$; this is the Robin, or radiation, boundary condition. Imposing it gives $B = 2\hat{\kappa}A$, and $u(a) = 1$ gives $A(1 + 2\hat{\kappa}a) = 1$. Evaluating at the origin gives the stated value, which coincides with the point-reactivity theorem. □

The agreement is the whole point. A boundary condition with a constant in it, and a path functional with a random variable in it, are the same statement seen from two sides: the Robin constant is the rate at which contact is converted into reaction probability. What the field description adds is that the conversion is visible as a law rather than as a constant, so it can be changed — which is the content of the earlier remark on non-exponential thresholds, and what a boundary condition cannot follow.

Corollary 6.4 (The absorbing limit)**.** *As* $\hat{\kappa} \to \infty$ *the survival probability against a reactive point vanishes as*

$$S \;=\; \frac{1}{1 + 2\hat{\kappa}a} \;=\; \frac{1}{2\hat{\kappa}a} \;+\; O(\hat{\kappa}^{-2}),$$

recovering the perfectly absorbing barrier, for which no particle started at the barrier survives. At the other end, for $\hat{\kappa}a \ll 1$*,* $S = 1 - 2\hat{\kappa}a + O(\hat{\kappa}^2)$*: survival is linear in the reactivity and the medium is reaction-limited.*

Proof. Both are expansions of the same expression, at large and small $\hat{\kappa}a$ respectively. The first is the statement that a particle started on a perfectly absorbing site is absorbed immediately, which is the classical Smoluchowski idealisation [13]; the second is the regime in which the particle visits the site many times and reacts on none of them, so that the reaction rate rather than the diffusion controls the outcome. □

The dimensionless group is $\hat{\kappa}a$, the product of the reactivity and the distance to the target, and it separates the two regimes a kinetics experiment can be in. It is also the reason the

house parameters are set where they are: $\hat{\kappa}a = 0.25$ sits between the two limits, where survival is neither near zero nor near one and a simulation has the most to say.

Example 6.5 (The house point)**.** Keep $D = 1$ and $a = 1$ and concentrate the same total reactivity at the origin: $\kappa = 0.5$, so $\hat{\kappa} = 0.25$. The contact at the origin is exponentially distributed with mean $2a = 2$, and survival is

$$S \;=\; \frac{1}{1 + 2 \times 0.25 \times 1} \;=\; \frac{2}{3} \;=\; 0.666667.$$

Three readings of the same number are worth separating. As a Laplace transform it is the value at $\hat{\kappa}$ of the transform of an exponential law of mean 2. As a boundary-value problem it is $u(0)$ for the affine function with slope $2\hat{\kappa}u(0)$ reaching one at the target. As kinetics it says the reactive resistance $2\hat{\kappa}a = 0.5$ is half the escape term, so the site is neither diffusion-limited nor reaction-limited but squarely between them. Doubling the reactivity to $\hat{\kappa} = 0.5$ drops survival to 0.5; halving it to 0.125 raises survival to 0.8. The reciprocal, not the probability, is the linear quantity, which is why reciprocal survival is what the tables below report.

Method. The point-reactivity claim has two testable halves: that the contact at the starting level is exponentially distributed with mean $2a$, and that survival is the Laplace transform of that law at $\hat{\kappa}$. Both are measured on the same batch, the first by the moments and the tail of the simulated contact and the second by averaging $e^{-\hat{\kappa}L^0}$ over paths.

Parameters. $D = 1$, $a = 1$, $\hat{\kappa} = 0.25$, $h = 0.005$, 40,000 paths.

Result. The contact at the origin has sample mean 1.9930 ± 0.0099 against the predicted $2a = 2$, a deviation of -0.71 standard errors, and sample standard deviation 1.9757 against the predicted 2 — an exponential law has its mean and its standard deviation equal, and the sample reproduces that to within one and a quarter per cent. The tail probability $\mathbb{P}(L^0 > 2a)$ is 0.3680 against $e^{-1} = 0.3679$. Survival is $\widehat{S} = 0.66681 \pm 0.00117$ against the exact $2/3$, a deviation of $+0.12$ standard errors.

Error. The relative error on the survival probability is 2.1×10^{-4}. The agreement of the tail with e^{-1} to four decimal places is the sharper of the two tests, because it probes the shape of the law rather than its first moment; a contact law with the right mean and the wrong shape would pass the survival test at one value of $\hat{\kappa}$ and fail here.

Script. `fig_ck.py`, one batch of 12,000 paths at $h = 0.01$, used twice: once to histogram the contact accumulated at the starting level, and once to price survival at five reactivities by reweighting that same sample.

Shows. On the left, the simulated law of the contact at the origin against the exponential density of mean $2a$ that the squared Bessel description predicts. On the right, survival against the reduced point reactivity, exact curve and simulated points.

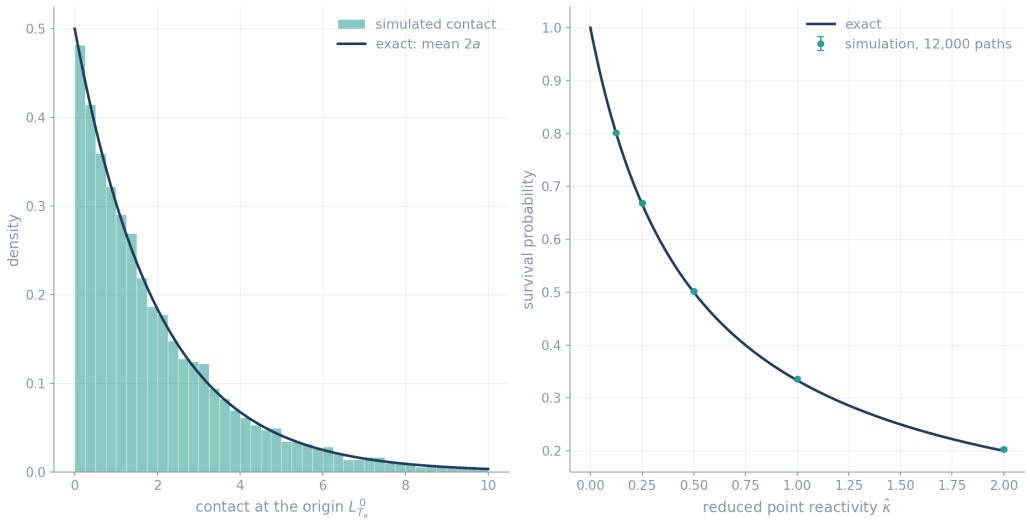

Figure 2: The law of contact at the starting level against its exact exponential density, and the survival probability against point reactivity with exact curve and simulated points from the same batch.

The left panel is the sharper of the two tests. Survival at a single reactivity is one number and many wrong laws reproduce one number; the histogram constrains the whole shape, and an exponential law is the one shape for which mean and standard deviation coincide. The right panel then shows the consequence across a sixteenfold range of reactivity, where survival falls by a factor of five and the points stay on the curve, with the reciprocal rather than the probability being the quantity linear in the reactivity.

Source. One batch of 12,000 paths at $h = 0.01$, scored at six reactivities by reweighting the same contact variable — the point-reactivity survival is the Laplace transform of one random variable, so a single simulated sample of that variable prices every reactivity at once.

Columns. The reduced point reactivity; the simulated survival with its standard error; the exact value; and the deviation in standard errors.

$\hat{\kappa}$	**simulated**	**exact**	**dev.**
0	1.00000	1.00000	—
0.125	0.80129 ± 0.00148	0.80000	+0.87
0.25	0.66843 ± 0.00215	0.66667	+0.82
0.50	0.50218 ± 0.00264	0.50000	+0.82
1.00	0.33595 ± 0.00274	0.33333	+0.96
2.00	0.20301 ± 0.00247	0.20000	+1.22

The table walks from the free particle to the strongly reacting site, survival falling by a factor of five, and the largest deviation anywhere is 1.22 standard errors. The deviations are all of one sign because all six rows are reweightings of a single batch and are therefore

strongly correlated; they are six views of one sample, not six independent tests, and the correct reading is that this sample's contact ran about one per cent low. The absorbing limit is approached but not reached: at $\widehat{\kappa} = 2$ the survival is still a fifth, and driving it to zero requires the reactivity to diverge, as the corollary says.

7. What the Two Cases Show Together

Method. The two instantiations are scored on one batch of paths rather than on two, which is what makes the comparison a measurement rather than an arithmetic identity. Each path carries both observables at once: the time it spends in the slab, and its contact at the starting level. The first gives survival against a medium of width w, the second gives survival against a point of the same total reactivity, and the difference between them is the entire content of the effective-reactivity proposition.

Parameters. One batch: $h = 0.005$, 40,000 paths, $a = 1$, $\widehat{\kappa} = 0.25$, slab width $w = 0.1$ so that $\widehat{v}_0 = 2.5$.

Result. The slab gives 0.67569 ± 0.00113 against its exact 0.676195; the point gives 0.66681 ± 0.00117 against its exact $2/3$. The measured gap between the two survivals is 0.00888 against the predicted 0.00953, and because both are read from the same paths the comparison is free of the batch-to-batch variation that would otherwise swamp a difference of under one per cent.

Error. The two deviations are -0.45 and $+0.12$ standard errors, and the gap is reproduced to within seven per cent of itself. The gap is the quantity the two instantiations exist to produce: it is what a point description of a medium of finite width gets wrong, it is small, and it is in the direction the effective-reactivity proposition predicts, the medium being the more survivable of the two.

Script. `fig_bridge.py`, which scores both observables on every path of one batch: the reacted contact accumulated in the medium, and the reacted contact the same path would have accumulated against a point of the same total strength. 8,000 paths at each of two widths.

Shows. The joint law of the two observables, and its collapse onto the diagonal as the medium narrows.

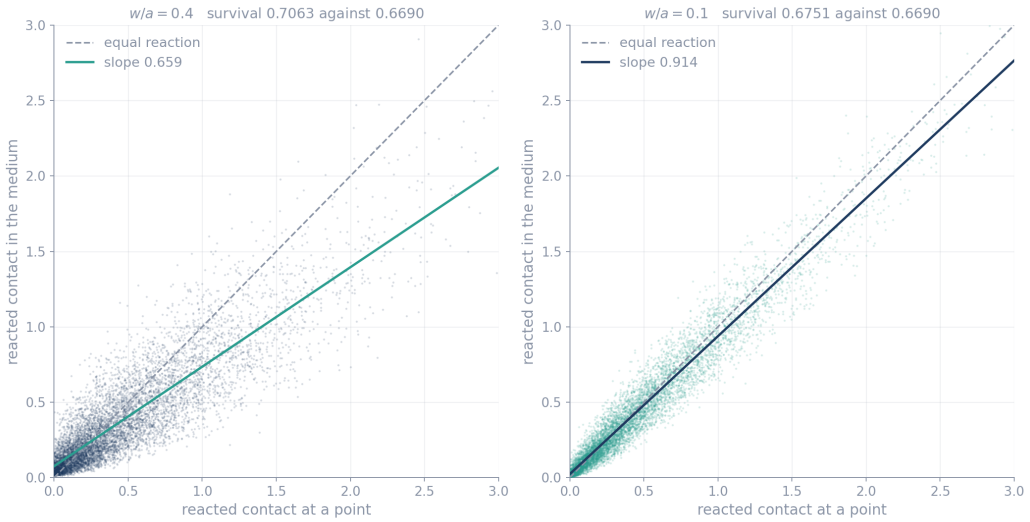

Figure 3: The reacted contact in the medium against the reacted contact at a point for the same paths, at two widths, showing a broad cloud with slope well below one at the wider medium and a near-collapse onto the diagonal at the narrower one.

This is the figure that shows why the comparison is a measurement and not an arithmetic identity. At $w/a = 0.4$ the cloud is broad and its regression slope is 0.659: a path that spends a long time at the starting level does not necessarily spend a proportionate time in a medium that extends well beyond it, and the two observables genuinely differ path by path. At $w/a = 0.1$ the cloud tightens onto the diagonal with slope 0.914, and the two descriptions are nearly the same random variable — which is precisely why their survival probabilities differ by under one per cent there, and why that difference is resolvable at a batch size where each survival separately carries an error bar three times its size.

Neither instantiation on its own says what a point description of a medium costs. The medium alone gives a closed form with two parameters and no reference value to judge it against; the point alone gives a classical answer with no width in it. Scored on one batch of paths, the pair gives the quantity an experimentalist actually needs: the error committed by replacing a layer with a surface, as a function of how thick the layer is relative to the distance the particle must travel.

The crossover has three regimes, and the error panel of the medium figure shows all of them. For w/a below about a tenth the two descriptions agree to within one per cent in survival and the layer may be replaced by a point with a clear conscience. Between a tenth and a third the first-order correction $1 - w/2a$ is doing visible work: at $w/a = 0.2$ the effective reactivity is 91.4 per cent of the total, and the first-order prediction of 90 per cent captures most but not all of the deficit. Beyond about a third the expansion is no longer adequate and the closed form must be used: at $w/a = 0.6$ the effective reactivity is 73.2 per cent against a first-order prediction of 70 per cent, an error in the correction itself of a tenth.

What makes this a single paper rather than two is that both numbers come from the same field on the same paths, and the joint law of the two observables is drawn below. The

slab survival and the point survival are two functionals of one realization, and the gap between them — 0.00888 measured against 0.00953 predicted at the house parameters — is a difference of correlated quantities, which is why it is resolvable at all at a batch size where each survival probability separately carries an error bar three times the size of the gap.

8. Numerical Method

Algorithm 2 — Simulating the contact field

Inputs: lattice spacing h, target a, reflecting depth M, medium $[0, w)$ with reduced reactivity $\hat{v}_0$, path count n. *Output:* the contact field per path, the reacted contact per path, and the survival estimate.

1. Represent the particle by a simple random walk on $h\mathbb{Z}$ with time step h^2, started at the origin. This is standard Brownian motion to the accuracy the convergence table reports.
2. At each step move $\pm h$ with equal probability, forcing an upward step at $-M$ and stopping the path at a.
3. Estimate contact at a level by visits: $\widehat{L}^{\,y} = hN_y$, where N_y is the number of steps spent at the site y. A site's occupation time is $N_y h^2$ and the site represents an interval of width h, so the density in the level variable is $N_y h$.
4. Accumulate the reacted contact as $\hat{v}_0 h^2 N_{[0,w)}$, the reduced reactivity times the time spent in the medium; by the occupation-times formula this is the same number as the integral of the profile against the field.
5. Report $\widehat{S} = n^{-1} \sum_i e^{-q_i}$ with q_i the reacted contact of path i, and its standard error as the sample standard deviation over $\sqrt{n}$.

The reflecting barrier is a device, not a model change. The reactivity is supported in $[0, a]$, so contact accumulated below the origin is never reacted, and no excursion below the origin reaches the target. Replacing a deep excursion by a reflected one therefore leaves every observable in this paper unchanged while making the mean path length finite, which it is not for free Brownian motion. This is asserted here and measured in the table below, across a fourfold range of M.

The cost of the method is set by the mean number of steps a path takes, which for the walk reflected at $-M$ and absorbed at a is $2(aM + a^2/2)/h^2$ in expectation. At the house geometry with $M = a = 1$ this is $3h^{-2}$: thirty thousand steps at $h = 0.01$ and one hundred and twenty thousand at $h = 0.005$. Because the hitting time has a heavy tail even under reflection, the slowest path in a batch runs an order of magnitude longer than the mean, so the implementation advances all live paths as one vector and compacts the array when paths finish, which keeps the total work proportional to the summed path lengths rather than to the longest one.

Every figure and every number in this paper comes from one engine and one convention: pseudorandom numbers from a counter-based generator with the seed recorded beside each result, one independent batch per reported quantity, and no reuse of a batch between a quantity and its own error bar. Standard errors are the sample standard deviation over the square root of the path count, and they are reported with every simulated number

rather than at the end, because a simulated number without its error is not a measurement. Where a quantity is compared with an exact value, the deviation is reported in units of its own standard error, so that agreement and disagreement are stated on the same scale.

Proposition 8.1 (Bin-width bias)**.** *Let*

$$\widehat{L}_\varepsilon(x) \;=\; \frac{1}{2\varepsilon}\int_{x-\varepsilon}^{x+\varepsilon} L^y_{T_a}\,\mathrm{d}y$$

be the contact estimated with a bin of half-width ε. Since $\mathbb{E}[L^y_{T_a}] = 2(a-\max(0,y))$ for $y \le a$ and 0 above,

$$\mathbb{E}\big[\widehat{L}_\varepsilon(x)\big] \;=\; 2(a-x) \quad \textit{for } \varepsilon < x < a-\varepsilon,$$

exactly and for every ε; while at the two kinks

$$\mathbb{E}\big[\widehat{L}_\varepsilon(0)\big] = 2a - \frac{\varepsilon}{2}, \qquad \mathbb{E}\big[\widehat{L}_\varepsilon(a)\big] = \frac{\varepsilon}{2}.$$

The bias is therefore zero in the interior, $-\varepsilon/2$ at the starting level and $+\varepsilon/2$ at the target, where the true value is zero.

Proof. The mean contact profile is the Green function of Brownian motion killed at a, evaluated from the origin, and is piecewise linear with kinks at 0 and at a. The mean of the estimator is the average of that profile over the bin. A linear function averaged over an interval symmetric about x equals its value at x, which gives exact unbiasedness wherever the bin avoids both kinks. At the starting level the bin straddles the kink, and

$$\frac{1}{2\varepsilon}\Big[\int_{-\varepsilon}^{0} 2a\,\mathrm{d}y + \int_{0}^{\varepsilon} 2(a-y)\,\mathrm{d}y\Big] = \frac{1}{2\varepsilon}(4a\varepsilon - \varepsilon^2) = 2a - \frac{\varepsilon}{2}.$$

At the target the profile is zero above a, so $(2\varepsilon)^{-1}\int_{a-\varepsilon}^{a} 2(a-y)\,\mathrm{d}y = \varepsilon/2$, against a true value of zero. □

The proposition is what makes the verification quantitative. A binned estimator is usually treated as approximate with an unquantified error; here the error is zero where it matters and exactly $\varepsilon/2$ where it does not, so any discrepancy larger than that is a defect in the simulation rather than in the estimator, and the table below tests precisely that.

Method. The bias proposition is a statement about the continuum estimator, and the simulation is a lattice, so the comparison needs the lattice's own contribution stated exactly rather than fitted. The expected number of visits of the walk to a site is available in closed form: for the walk absorbed at a and reflected below the origin it is exactly $2(N-j)$ in lattice units, solved from the discrete Green function. Summing that arithmetic series over the sites a bin covers gives the lattice prediction in each of the three regimes.

Parameters. $h = 0.005$, 40,000 paths, $a = 1$, bins of half-width $\varepsilon = 0.025$, 0.05, 0.1 and 0.2, placed at the starting level, at the interior level 0.5, and at the target.

Result. The lattice adds $+h$ to the interior estimate and $+h/2$ at each kink, and at the target one further visit is recorded as the path is absorbed, worth $h^2/2\varepsilon$. With those three

terms the predictions are exact and every one of the twelve measurements agrees: the largest deviation over all twelve is 1.12 standard errors and the mean absolute deviation is 0.44. At the interior level the measurements are 1.00169, 1.00253, 1.00349 and 1.00408 against a prediction of 1.005 at every bin width, confirming that the continuum bias there is exactly zero and the whole offset is the lattice.

Error. Before the lattice terms were included, the target-bin measurements deviated by +50, +25, +12 and +6 standard errors from the continuum prediction — an apparently catastrophic failure of a proposition that is in fact exactly true. The discrepancy was the ruler, not the claim, and it was resolved by computing the ruler exactly rather than by widening a tolerance.

Script. `fig_field.py`. The left panel simulates three complete paths at $h = 0.005$ and records the contact at every level each one visits. The right panel takes 20,000 paths at $h = 0.01$ and histograms the contact at three fixed levels against the exponential density the squared Bessel description predicts for each.

Shows. That the mean profile is exactly a straight line and that no realization resembles it, and that the law at each level is exponential with the mean the theorem gives.

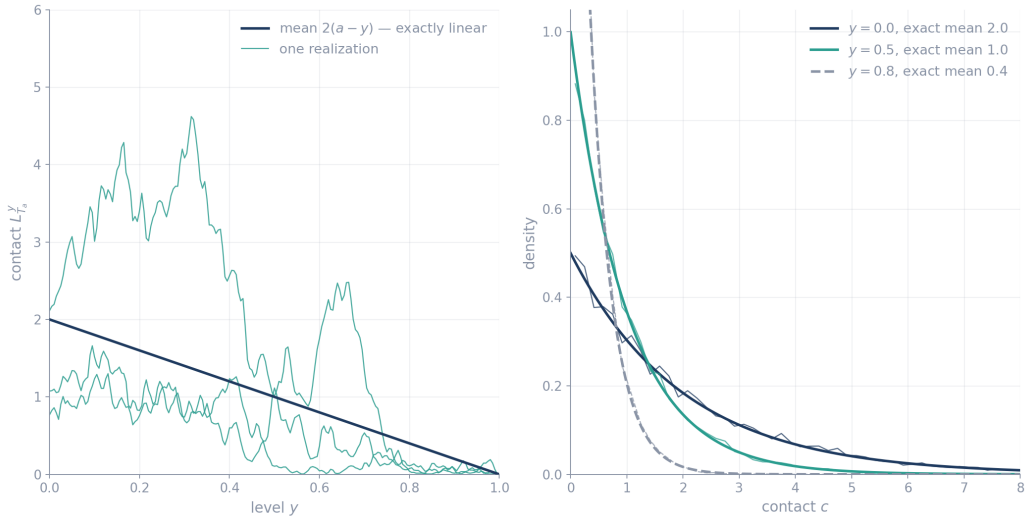

Figure 4: Three simulated contact fields wandering around an exactly linear mean profile, beside the histogram of contact at three levels lying on the exponential densities the squared Bessel description predicts.

The straight line in the left panel is not an artefact of plotting: the mean contact is the Green function of Brownian motion killed at the target and is exactly $2(a - y)$, linear by theorem. What the panel is for is the contrast with it — the realizations rise to three times the mean, fall to near zero and recover, and reach the target only where the field is extinguished. The right panel supplies the law behind that spread: at each level the contact is exponentially distributed, so its standard deviation equals its mean and the enormous variation on the left is exactly what is predicted. The three histograms sit on their exact densities across a fivefold range of mean.

Method. The first Ray-Knight theorem asserts a law for the whole field, so it is tested at several levels at once and in two ways: the marginal at each level, which the squared Bessel laws make exponential with mean $2(a-y)$, and the Markov structure in the level variable, which predicts that regressing the field at one level on the field at another gives unit slope and an intercept equal to twice the separation.

Parameters. $h = 0.005$, 40,000 paths, $a = 1$, levels $y = 0$, 0.25, 0.5, 0.75.

Result. The sample means are 1.9930 ± 0.0099, 1.5014 ± 0.0075, 0.9970 ± 0.0050 and 0.5020 ± 0.0025, against the predicted 2, 1.5, 1 and 0.5: deviations of -0.71, $+0.19$, -0.61 and $+0.80$ standard errors. The sample standard deviations are 1.9757, 1.4966, 0.9975 and 0.5019, against a prediction that each equals its own mean, which is the exponential shape; the largest relative discrepancy is 1.2 per cent, at the level where the sample is smallest relative to its own spread. Regressing the field at $y = 0.25$ on the field at $y = 0.5$ gives slope 1.0052 against the predicted 1 and intercept 0.4993 against the predicted 0.5.

Error. Every marginal agrees within one standard error, and the process test agrees within half a per cent on both coefficients. The second test is the one that distinguishes the theorem from a coincidence of marginals: a field with the right law at each level but the wrong dependence between levels would reproduce the four means and fail the regression, and it is the dependence that the survival computation actually consumes.

Method. The second Ray-Knight theorem is tested at the inverse contact time rather than at a hitting time, which requires a different simulation: the walk runs until its contact at the origin first exceeds ℓ, and the field is then read at levels above the origin. Two predictions are scored — that the mean contact is ℓ at every level, which is the martingale property of dimension zero, and that the probability the field has died by level y is $e^{-\ell/2y}$.

Parameters. $\ell = 1$, $h = 0.02$, 12,000 paths, reflecting depth $M = 1$, upper reflecting level $A = 20$, levels $y = 0.2$, 0.5, 1, 2.

Result. The mean contacts are 1.0086 ± 0.0080, 1.0051 ± 0.0129, 1.0017 ± 0.0182 and 1.0016 ± 0.0260, deviations of $+1.1$, $+0.4$, $+0.1$ and $+0.1$ standard errors from $\ell = 1$: the mean is conserved across a tenfold range of level while the field is dying under it. The extinction probabilities are 0.3676, 0.6052 and 0.7763 at $y = 0.5$, 1 and 2, against the exact 0.3679, 0.6065 and 0.7788, deviations of -0.1, -0.3 and -0.6 standard errors.

Error. One point fails and the reason is identified rather than absorbed. At $y = 0.2$ the extinction probability is 0.0739 against the exact 0.0821, a deviation of -3.4 standard errors, and it is unchanged when the upper reflecting level is doubled, so it is not a truncation effect. It is the lattice: an excursion must cross ten sites to reach that level, and the discrete probability of reaching site m before returning to the origin is $1/2m$ per visit, which departs from the continuum rate at relative order h/y. The deviation is therefore a resolution limit of the estimator at the smallest level tested, and it disappears at every level where the excursion has more than a few sites to cross.

A finding worth recording. At an upper reflecting level of $A = 10$, with 4.9 per cent of paths touching it, the mean contacts drifted low — 0.9949 and 0.9900 at $y = 1$ and $y = 2$ — and the drift vanished at $A = 20$ with 2.4 per cent touching. Dimension zero is a martingale

whose mean is carried by rare large values, so truncating the top of the range biases the mean downwards while leaving the extinction probabilities, which are insensitive to the tail, untouched. A convergence check on the extinction curve alone would have passed.

Script. `fig_rk2.py`, which runs 4,000 paths to the inverse contact time τ_ℓ with $\ell = 1$ at $h = 0.02$, reflecting at level 20 above the origin, and reads the field at six levels.

Shows. On the left, on twin axes, the probability that the field has died by a given level and the mean contact among the paths where it has not. On the right, the upper tail of the field's law at three levels.

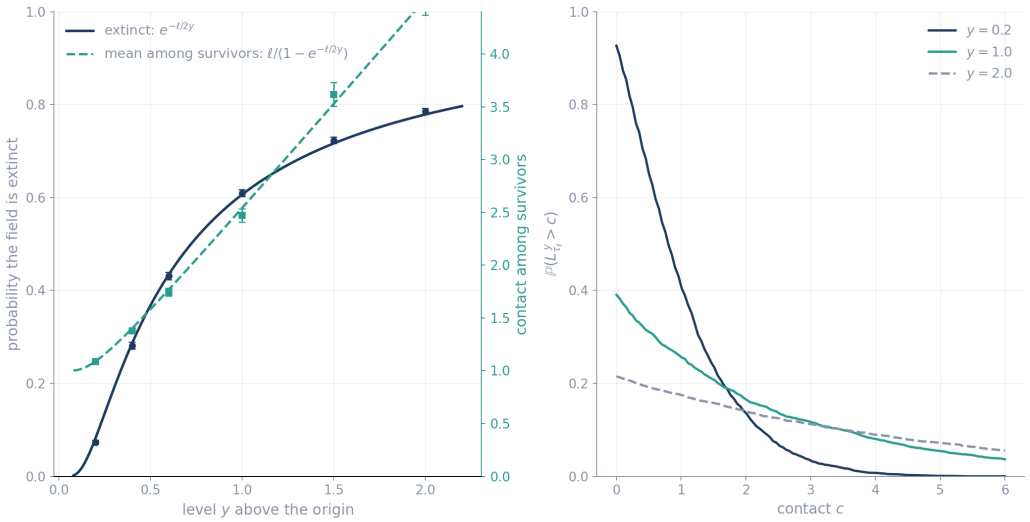

Figure 5: Extinction probability of the contact field rising with level, plotted against the mean contact among survivors rising in step on a twin axis, beside the upper tail of the field's law at three levels.

The left panel resolves the apparent contradiction of dimension zero. Extinction climbs from nothing to four fifths across two units of level; the mean among the survivors climbs from one to three and a half over the same range; and the product of the two is the constant ℓ that the martingale property demands. Neither curve alone is the claim — the conservation of the mean is the *balance* between them, and the figure shows the balance rather than asserting it. The right panel is the same fact read as a law: at the lowest level the field is almost surely alive and concentrated near ℓ, and by two units above the origin the mass has moved out into a tail that is thinner at every height but reaches much further.

Source. Two sweeps, each of 20,000 paths, on the house slab: the lattice spacing varied over an eightfold range at fixed reflecting depth, and the reflecting depth varied over a fourfold range at fixed spacing. The exact survival is 0.676195 throughout, and the exact mean contact at the origin is 2.

Columns. The parameter varied; the simulated survival with its standard error; the deviation from the exact value in standard errors; the simulated mean contact at the origin; and the mean number of steps a path takes, which is the cost.

varied	simulated survival	dev.	mean contact	steps
$h = 0.04$	0.72207 ± 0.00148	$+31.0$	—	1,867
$h = 0.02$	0.67674 ± 0.00161	$+0.34$	—	7,415
$h = 0.01$	0.67480 ± 0.00161	-0.87	—	30,191
$h = 0.005$	0.67579 ± 0.00161	-0.25	—	120,400
$M = 0.5$	0.67530 ± 0.00161	-0.56	1.9908	19,998
$M = 1.0$	0.67434 ± 0.00161	-1.15	1.9994	30,178
$M = 2.0$	0.67564 ± 0.00160	-0.35	1.9822	50,110

Two things are established. The coarsest spacing fails badly — thirty-one standard errors — and the reason is not statistical: at $h = 0.04$ the medium of width 0.1 is two and a half lattice sites wide and is not resolved at all, so the simulation is answering a different question. From $h = 0.02$ downwards, where the medium spans five sites or more, every deviation is under one standard error and the cost grows as h^{-2}. And the reflecting depth, varied over a fourfold range at four times the cost, moves nothing: the three survivals agree within one standard error of each other and the three mean contacts agree with 2, which is what the claim that the barrier is a device and not a model change predicts.

9. Discussion

Remark 9.1 (What fails in higher dimensions). The contact field of Brownian motion in two or more dimensions is not a function. Points are polar, so there is no occupation density at a point, and the object that survives is contact with a *set* — the boundary local time of a reflected process on a surface, which is the quantity the encounter-based literature works with [9]. The Ray-Knight description does not extend to it.

What does extend is the shape of the argument, and it is worth being precise about which part. The step from a reaction rate to an accumulated contact is dimension-free; so is the observation that survival is the Laplace transform of that contact, and so is the point that a non-exponential threshold needs the whole law rather than one transform value. What is one-dimensional is the closed form. In higher dimensions the law of the boundary local time is known domain by domain, through spectral expansions of the Dirichlet-to-Neumann operator, and the tidy reciprocal of this paper becomes an eigenvalue problem again. The one-dimensional case is therefore not a warm-up for the general one; it is the case in which the general framework has an exact answer, which is why it is worth having explicitly.

The quantity this paper computes is measured, not merely modelled. In nuclear magnetic resonance, magnetisation is lost when a spin-bearing molecule encounters a pore wall, and the loss per encounter is set by a surface relaxivity with the dimensions of a velocity; the measured decay of the signal is then an average of $e^{-\rho\ell}$ over the ensemble of molecules, with ℓ the accumulated contact. That is the same functional as survival, with relaxivity in place of reactivity, and the classical analysis of Brownstein and Tarr [10] is its short-time

and long-time asymptotics. A pore wall with a coating of finite depth, rather than an idealised surface, is exactly the medium of the first instantiation.

The same functional appears in fluorescence quenching, where an excited molecule is deactivated by contact with quencher distributed through a region rather than concentrated at a boundary, and in heterogeneous catalysis, where the reactive zone at an interface has a depth set by the chemistry. In each case the experimentally accessible number is a decay constant, and what the present computation supplies is the map from a reactivity profile to that decay constant without an intervening fit: given the profile and the geometry, the constant is the reciprocal of one solution of one ordinary differential equation, and the width of the reactive zone enters with the explicit weight the effective-reactivity proposition gives it.

A second class of applications has the reactive sites distributed along a line for structural rather than chemical reasons. A protein searching for its binding site on a nucleic acid alternates between three-dimensional excursions and one-dimensional sliding along the strand, and during the sliding phase the search is exactly the problem of this paper: a one-dimensional diffusion accumulating contact against sites laid out along a coordinate, with a reaction probability per unit of contact. The facilitated-diffusion analysis of Berg, Winter and von Hippel [11] is built on that alternation, and the sliding phase is where the contact field lives.

What the field description offers there is the case the classical treatment does not cover: a target that is not a single site. A binding region of several base pairs, a stretch of lower-affinity sites flanking a high-affinity one, or a strand with a graded affinity profile are all reactivity profiles of finite width, and the present result prices them exactly. The effective-reactivity proposition then has a concrete reading: a target region of extent w on a strand where the protein slides a distance a is caught less efficiently than a point target of the same total affinity, by a factor $1 - w/2a$ to leading order, and the deficit depends on the geometry of the search rather than on the strength of the binding.

Remark 9.2 (Scope). What is established here is exact within a model whose boundaries should be stated plainly. The particle is Brownian: no drift, no confinement beyond the target, no memory, no interaction with other particles, and no back-reaction of the medium. The reaction removes the particle and does not deplete or saturate the medium. The observation is stopped at a first hitting time or an inverse contact time, not at a fixed clock time.

Three of these are the interesting restrictions. A drift would replace Brownian motion by a diffusion whose Ray-Knight description exists but is no longer the plain squared Bessel process, and the Sturm-Liouville problem would acquire a first-derivative term. Depletion would make the profile a function of the path and destroy the linearity that the whole computation rests on; that is a genuinely harder problem and not a refinement of this one. And the stopping rule is the sharpest limitation: a fixed-time experiment is not covered by either Ray-Knight theorem, and the honest statement is that the results here apply to an experiment run until the particle has escaped a known distance, which some experiments are and others are not. The handbook of Borodin and Salminen [6] collects the laws under which each of these could be attempted.

10. Conclusion

A diffusing particle reacting with a medium distributed in space accumulates contact at every level it visits, and the observable that decides its fate is a weighted integral of that contact. Read in the level variable, the contact of one-dimensional Brownian motion is a squared Bessel process, and the integral against it is therefore governed by a Riccati equation. The survival probability is the reciprocal of the solution at the target of $W'' = 2\hat{v}W$ started flat at the particle's own position: one linear equation, driven by the reactivity as a potential, with no eigenvalue to find and no geometry left in it.

Two instantiations fix the result at both ends of its range. A medium of finite width has the closed form $S = [\cosh kw + k \sinh kw\,(a - w)]^{-1}$, and a reactive point has $S = (1 + 2\hat{\kappa}a)^{-1}$, which is Collins-Kimball kinetics and the Robin boundary condition seen as a statement about contact. Between them lies the number that neither supplies alone: a medium of width w behaves as a point of reactivity reduced by $1 - w/2a$, a deficit set by the geometry of the search and not by the strength of the chemistry. Every claim is verified against simulation, including the ones that failed first — the truncated upper range that biased a martingale's mean, and the binning bias that looked like a fifty-sigma failure until the ruler was computed exactly.

What the field description buys over the boundary condition it reproduces is the whole law of the reacted contact rather than one value of its transform, and with it the kinetics of surfaces whose reaction threshold is not exponential — surfaces with memory, which no local boundary condition can express.

References

[1] Ray, D. B. (1963). Sojourn times of diffusion processes. *Illinois Journal of Mathematics*. https://doi.org/10.1215/ijm/1255645099

[2] Knight, F. B. (1963). Random walks and a sojourn density process of Brownian motion. *Transactions of the American Mathematical Society*. https://doi.org/10.1090/s0002-9947-1963-0154337-6

[3] Revuz, D., & Yor, M. (1999). *Continuous Martingales and Brownian Motion*. Springer. https://doi.org/10.1007/978-3-662-06400-9

[4] Pitman, J., & Yor, M. (1982). A decomposition of Bessel bridges. *Zeitschrift fur Wahrscheinlichkeitstheorie und verwandte Gebiete*. https://doi.org/10.1007/BF00532802

[5] Going-Jaeschke, A., & Yor, M. (2003). A survey and some generalizations of Bessel processes. *Bernoulli*. https://doi.org/10.3150/bj/1068128980

[6] Borodin, A. N., & Salminen, P. (2002). *Handbook of Brownian Motion - Facts and Formulae*. Birkhauser. https://doi.org/10.1007/978-3-0348-8163-0

[7] Collins, F. C., & Kimball, G. E. (1949). Diffusion-controlled reaction rates. *Journal of Colloid Science*. https://doi.org/10.1016/0095-8522(49)90023-9

[8] Grebenkov, D. S. (2020). Paradigm shift in diffusion-mediated surface phenomena. *Physical Review Letters*. https://doi.org/10.1103/PhysRevLett.125.078102

[9] Grebenkov, D. S. (2019). Probability distribution of the boundary local time of reflected Brownian motion in Euclidean domains. *Physical Review E*. https://doi.org/10.1103/PhysRevE.100.062110

[10] Brownstein, K. R., & Tarr, C. E. (1979). Importance of classical diffusion in NMR studies of water in biological cells. *Physical Review A*. https://doi.org/10.1103/PhysRevA.19.2446

[11] Berg, O. G., Winter, R. B., & von Hippel, P. H. (1981). Diffusion-driven mechanisms of protein translocation on nucleic acids. 1. Models and theory. *Biochemistry*. https://doi.org/10.1021/bi00527a028

[12] Donsker, M. D., & Varadhan, S. R. S. (1975). Asymptotics for the Wiener sausage. *Communications on Pure and Applied Mathematics*. https://doi.org/10.1002/cpa.3160280406

[13] Redner, S. (2001). *A Guide to First-Passage Processes*. Cambridge University Press. https://doi.org/10.1017/CBO9780511606014

[14] Bénichou, O., & Voituriez, R. (2014). From first-passage times of random walks in confinement to geometry-controlled kinetics. *Physics Reports*. https://doi.org/10.1016/j.physrep.2014.02.003